



Compositional routes to (non)intersectivity

Citation

Martin, Joshua Robert. 2022. Compositional routes to (non)intersectivity. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

Permanent link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37372256>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

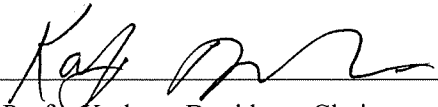


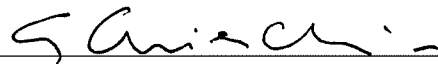
DISSERTATION ACCEPTANCE CERTIFICATE

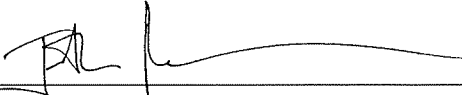
The undersigned, appointed by the
Department of Linguistics
have examined a dissertation entitled
"Compositional routes to (non)intersectivity"

presented by Joshua Martin

candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature 
Typed name: Prof. Kathryn Davidson, Chair

Signature 
Typed name: Prof. Gennaro Chierchia

Signature 
Typed name: Prof. Jonathan Bobaljik

Date: May 5, 2022

Compositional routes to (non)intersectivity

A DISSERTATION PRESENTED
BY
JOSHUA MARTIN
TO
THE DEPARTMENT OF LINGUISTICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
LINGUISTICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2022

2022 - JOSHUA MARTIN
ALL RIGHTS RESERVED.

Compositional routes to (non)intersectivity

ABSTRACT

This dissertation is concerned with a particular instance of ambiguity in natural language, between intersective and nonintersective interpretations of certain adjectives. Adjectives like *good* display a consistent alternation across many languages between a noun-dependent reading and a noun-independent reading, and this alternation is sensitive to various syntactic factors concerning word order and locality. In this dissertation, I develop a compositional account of the (non)intersective ambiguity. Central to this account is the idea that ambiguity is derived in the course of syntactic derivation, from uniform underlying adjective denotations, where different interpretations result from modification occurring in different domains of the nominal phrase. I argue that these interpretive domains are divided by the introduction of grammatical number to the noun, in the form of number marking or a classifier, and correspond to the ontological distinction between nouns-as-kinds and nouns-as-objects. Empirical motivation for this analysis comes from word order alternations in Italian, the interaction of plurality and kind-readings in Turkish, and focus-driven movement in Bangla, with a particular focus on the previously underemphasized behavior of privative adjectives. Maintaining uniform adjective denotations is enabled by positing covert structure in intersective readings of adjectives, motivated by visible morphophonological effects of that structure influencing suppletion across languages. I also extend the account by offering experimental evidence in favor of a fine-grained, semantically active lexical-conceptual structure for kinds, in order to capture compositionally the nonintersective readings of privative adjectives.

Contents

Abstract	iii
Acknowledgements	vi
1 Introduction	1
1.1 The study of meaning in language	1
1.2 Modification problems	5
1.2.1 The problem of ambiguity	6
1.2.2 The problem of privativity	10
1.3 Theoretical assumptions	13
1.4 Roadmap for the dissertation	16
1.4.1 Overview of Chapter 2	18
1.4.2 Overview of Chapter 3	20
1.4.3 Overview of Chapter 4	21
1.4.4 Overview of Chapter 5	23
2 The distribution of ambiguity	25
2.1 Past theories	25
2.1.1 Two types of adjectives (Siegel 1976)	25
2.1.2 Two parts of nouns (Larson 1998)	27
2.1.3 Pragmatic variable resolution (Maienborn 2020)	29
2.1.4 Syntactic ambiguity (Cinque 2010)	40
2.2 Privative disambiguation in Italian	51
2.2.1 Canaries in the mine	52
2.2.2 One-to-one syntax, one-to-many semantics	56
2.2.3 Contextually possible arguments	65
2.2.4 Economical asymmetries	66
2.3 Movement and ambiguity in Bangla	70
2.3.1 Focus movement in Bangla	70
2.3.2 A note on phase structure	73
2.3.3 Effects of movement on ambiguity	77
2.3.4 The irrelevance of focus	81
2.3.5 Privative disambiguation again	83
2.4 Interim summary	88
3 The morphosyntax of intersection	92
3.1 Suppletion in Serbian	92
3.1.1 The ban on intersective suppletives in Serbian	94
3.1.2 The diradical analysis	96
3.1.3 Issues with the diradical analysis	98
3.2 Suppletion across languages	100
3.2.1 The Nonintersective Suppletion Generalization	100
3.2.2 The Basque problem	106
3.2.3 A monoradical analysis	107

3.2.4	Lingering issues with the monoradical analysis	111
3.3	Interim summary	112
3.3.1	A note on compositionality and explanation	113
4	Number and kind modification	115
4.1	Modifying kinds and objects	116
4.1.1	Reference to kinds	116
4.1.2	Kinds, nouns, and number	120
4.1.3	Leffel (2014) on kind modification	123
4.2	Kind- and object-modification in Turkish	133
4.2.1	Number neutrality and modification	134
4.2.2	Singular kind terms	138
4.2.3	Applying the diagnostic	142
4.3	Nonintersectivity as kind modification	147
4.3.1	Revising the semantic analysis	147
4.3.2	Visibility problems	151
4.3.3	A genuinely nonintersective analysis	153
4.3.4	Grammatical number versus classifiers: the view from Bangla	158
4.4	Interim summary	162
5	Inside nonintersective modification	166
5.1	Returning to the problem of privativity	168
5.1.1	Privativity as a compositional problem	168
5.1.2	The coercion account	171
5.1.3	The limits of coercion	173
5.2	Experiment 1: shifting the default	175
5.2.1	Methods	176
5.2.2	Results	178
5.2.3	Discussion	180
5.3	Experiment 2: privative variation	182
5.3.1	Methods	184
5.3.2	Results	186
5.3.3	Discussion	194
5.4	Dual Content Semantics	200
5.4.1	Motivations for lexical enrichment	200
5.4.2	The semantic framework	203
5.4.3	Making privativity emergent	206
5.4.4	The role of extensional structure	211
5.5	Theorizing nominal concepts	216
5.5.1	Derivational, not simultaneous, duality	216
5.5.2	Roots identify concepts; kinds structure them	221
6	Conclusion	225
	References	233

Acknowledgements

It should come as no surprise to anyone who knows me well enough to show up in this section that I'm writing this at the absolute last possible moment, and so it's not going to be as poetic and elaborate as I had always envisioned it would be. But it would be painfully incomplete to finalize this document without some reference to all of the people without whom it would never have existed, even in a form much briefer than they deserve.

I may as well start at the beginning. To my parents, Rich and Jennifer: it should go without saying that nothing I've ever accomplished would have been possible without you two and your unending support, but that doesn't mean I shouldn't maybe say it more often anyway. Ours wasn't an 'academic' household in having generations of PhDs, but it was in the more important sense, how from the beginning you cultivated an environment where your kids were encouraged to ask questions, challenge ideas, think for themselves, and go off on our own intellectually even when it was confusing and conflicted with what you believed. You recognized a love for learning in me early on, and learned quickly, maybe by instinct, that the best way to cultivate that was not to push me to take school more seriously but to lean back and let me explore and engage in the way that worked for me. It's easy to imagine another path where that didn't happen, and I burned out on this work long before I could have written a dissertation. Add your dedication to supporting me through college and the way that you always made home a comfortable and exciting place to return to, and it's impossible to overstate how big of an impact you've had on helping me get through what's been a rocky path at the best of times; when so many people's parents are yet another source of stress on top of all of their work, it's been an absolute privilege to have parents who made things easier time and time again. I love you both very much, and thank you for everything.

To my younger brother and sister, Adam and Olivia, too - that family environment wouldn't have worked at all if it weren't for two brilliant, ambitious, impressive young people on the other side of me. It would've been really easy to have siblings who didn't care about the kind of weird, niche work that I do, but instead I had two cheerleaders who somehow

found it unfailingly cool (if I'm wrong about that, don't tell me) and were always excited to listen. I doubt that I would have found my love for teaching - the thing that motivated me through grad school - or been nearly as good at it, if I didn't have the practice of having you two listen to me talk the whole time that we were growing up. And while I've been on the other side of the country, it's been no small comfort to know that you were back home making sure things didn't fall apart in the chaos. I know that you'll both end up achieving anything you set your minds to; you're better at that than me, after all.

I had more than my fair share of good teachers in my life, and I'm sure that I'll miss out on some here, but I'll try to give credit where it's most due. My ability to do research, write, give talks, teach, evaluate arguments, anticipate responses, and everything else that makes writing a dissertation possible all has its roots in the time I spent in speech and debate, which has been such a critical piece of my life for so long. My coaches - McClay and Jonah in particular - were the first people to really teach me how to criticize, organize my thoughts, synthesize information, do research, and present arguments. It's also no surprise that I've ended up thinking so much about the interfaces between language and the mind, thanks to the teachers that got me excited about words (Mr. Riley) and brains (Mr. DeShazer).

I didn't plan to study linguistics when I started college; it was just one branch of cognitive science I had to cover in order to focus on what I wanted, which was neuroscience and philosophy. But Peter Jenks' introductory syntax class started with this beautiful description of how, because we are all language users who implicitly know these rules even if we can't describe them, linguistics is the act of 'wondering at the obvious'. At that moment, I was hooked, and I've repeated that description to countless people since. Peter gravitated to my inherent interest in the subject early on, and told me exactly what classes I should take next, and that if I wanted to do an honors thesis when I was a senior, he would advise it. When I said I wasn't actually a linguistics major, he just told me that we can change that too, and he was right. He also vouched for me to take graduate classes, where I met Line Mikkelsen, who maybe had the largest impact on my life with a single handwritten sentence

on the back of a problem set: ‘Have you considered going to graduate school for linguistics?’

I hadn’t considered it at the time, but I’m glad that I did, and I’m glad it was at Harvard, where I found a community more welcoming and supportive than I ever could have guessed. That’s more true of my committee than anywhere else. Gennaro: thank you for letting me sneak in as your student in my last year here, and for some of the most stimulating and complex conversations I had in my academic career. There’s no one else I would be so happy to unsettle. Jonathan: thank you for lending your expertise to a topic firmly outside your area, and yet always having the most insightful thoughts on even the most semantic aspects of the problem. More than that, thank you for being a rock (an unfailingly kind one) during some of the roughest parts of grad school for me. And Kate: I really can’t imagine designing an advisor that I would have enjoyed working with more than you. You’ve always been the perfect combination of empathetic and supportive, ambitious and driving, critical and challenging, and every bit of work I’ve done is better for it. You always understood where I was coming from both as a theorist and as a human being, and faced with the impossible balancing act of trying to encourage and cultivate both of them, walked the tightrope with grace. I know your other students agree, and I’m grateful to be both your student and friend.

Even with a stellar committee, I never would’ve gotten through without friends. To Niels, Ian, and Ethan, you’ve been an irreplaceable cohort, and I’ll always cherish our G1 homework days in the lounge. I won’t go on about every other friend I’ve had here, but to Gunnar, Tamisha, Zuzanna, and (in particular for the dissertation) Shannon, this would have been such a sadder five years without you all around. I’m especially grateful to Giuseppe and Dasha for collaboration, and to Ankana, Hande, Yağmur, and Saurov for data and insights on this project. And of course, so many friends outside of Harvard to keep me sane, supported, and happy through all of it, who I won’t try to list because inevitably I will forget more than I remember, except to say that I would never have written anything of substance without John; our friendship has both shaped how I think and kept me grounded more than anything else throughout the years, so maybe this should be dedicated to you most of all.

1 Introduction

1.1 The study of meaning in language

Unless you are rereading this dissertation, it is overwhelmingly likely that you have not encountered this particular sentence, with precisely these words in precisely this order, ever before in your life. Equally likely, I hope, is that you nonetheless had no trouble understanding it. Despite never seeing those words arranged together in that sequence, your mind will have processed them, understood their relation to each other, and extracted the overall meaning that I intended to convey when writing it. If this did not occur, not only would it be generally futile to read other people's writing, but more broadly the entire enterprise of language as a means of communicating novel ideas between members of a species would be bankrupt. In order for language to be a useful system worth dedicating valuable brain space to, we need to be able to use it to transmit new ideas to one another, not just things we both already know, and the listener in any particular exchange needs to be able to decode the meaning of the sentences they hear. This decoding process needs to happen for potentially infinite new combinations of words, on the fly, when listening to different speakers and in different contexts, and also sometimes needs to handle words that are themselves new, as well. Luckily for us, in most cases, our brains seem to do just that.

Because this process is largely subconscious, however, we rarely stop to think about just how astonishingly difficult it might be, and what is required for it to occur. Or, put another way, just how bad the situation would be if certain aspects were different. Our ability to interpret new sentences - and, therefore, to communicate - relies on at least two critical pieces of language being consistent across contexts. First, words need to mean roughly the same thing no matter which sentence they show up in. While there are subtle variations and exceptions, broadly, the meaning of a given word is stable across the various sentence structures that it can show up in. Imagine if this were not the case, and the meaning of a word was fully determined by, for example, how far into a sentence it was placed. The

dictionary entry for a simple noun like *dog* would be massively disjunctive, saying something like ‘if it’s the first word in the sentence, it means... if it’s the second word in the sentence, it means...’ and so on. Such a system would be atrociously painful to learn, though at least it would be predictable! Communication would, strictly speaking, be possible, so long as the mapping between words and meanings was not completely random. But it could get worse.

A system with massively variable word meaning would be at least stable because we would still know how to put words together to make phrases. Counting up each word in a sentence and looking through your mental dictionary to assign it the correct meaning for that position would be an inefficient and tedious process, but once you had the meanings selected, you could proceed by combining those meanings in the regular, predictable way. Is there a verb followed by a noun? Then treat that noun as the object of the verb, and the combination will mean something like ‘doing that verb to that noun’. Another noun in front of the verb? Make that the subject, and then we have ‘the subject does that verb to the object’, and so on. Even if meanings got very complex, the presence of stable rules for composing individual parts into bigger chunks of meaning, culminating in a sentence, would enable the successful communication of ideas across speakers. This is the second piece of the linguistic system that needs to remain constant across contexts for effective communication - the processes by which individual chunks of meaning are composed to make larger ones. If this component were unconstrained, even with semi-consistent word meanings, sentence meanings would be effectively impossible to derive; that is, it would be impossible to predict which possible meaning of a sentence was intended by the one who produced it, upon perceiving it. Considering a minimal example, if the two-part sequence of a verb followed by a noun occasionally meant to put the pieces together as ‘the act of doing this verb to this noun’, but also unpredictably could mean any conceptually possible relationship between those things - some linguistically familiar ones like ‘that noun doing the verb to someone else’, but also plenty of other logically definable possibilities like ‘doing that verb to everyone except that noun’ or ‘doing all actions except that verb to this noun’, invoking extra content

like ‘doing that verb to the noun’s best friend’, and so on - then comprehension would be near impossible. Even with just these examples, we see that meanings of the same phrases could alternate unpredictably, or meanings of a phrase might not even involve the meaning of one of its parts. It’s difficult to even lay out in plain language the kind of outlandish options for connecting meanings that we would have. Perhaps the discourse context could help decipher a given meaning-meaning pair, but given that an average-length sentence would contain dozens of these pairwise interactions, and we would have no reason to expect that only words next to each other should have their meanings interact, etc., the project would become nightmarish.

Given that communication with language does in general appear to actually function, we must not be in this situation. It does seem like there are powerful constraints over and above logical possibility on both what individual words may mean and the operations that may combine those meanings. Learning what those constraints are and formalizing them has been probably the primary function of the tradition of formal semantics, to which this dissertation is going to belong. The fact that neither component of the meaning system seems to be random has often been codified in the notion of linguistic compositionality. The principle of compositionality (Frege, 1914; Montague, 1970) is perhaps the fundamental driving force behind the project of formal semantics, which might be described as the joint project of determining the meanings of individual lexical items in natural language and determining the nature of the mechanisms combining them, such that we can predict the truth conditions and felicity conditions of all natural language expressions from the application of such combinatorial mechanisms to such lexical meanings. A rough statement of the principle, then, might be given as in (1).

- (1) *The Principle of Compositionality:*
The meaning of a complex expression is a function of the meanings of its constituents and the way they are combined.

As Szabó (2012) identifies, however, such a formulation is underspecified on multiple

fronts. What *aspects* of the meanings of the constituents are relevant? Are we concerned with the way that the constituents themselves are combined, or the way that their *meanings* are combined? Does ‘is a function of’ mean that the meaning of the complex expression is *determined* by that of its constituents, or simply that there is such a function? Some weaker readings of ex:principle are trivially satisfied, while the stronger readings are, as Szabó says, bold empirical hypotheses. The strongest such reading is (2).

- (2) *The Principle of Compositionality, revised:*
The meaning of a complex expression is determined by the meanings its constituents have individually and of the way those constituents are combined.

Such a formulation as (2) commits semanticists to significantly more than (1). One reason this is advantageous is that it makes compositionality an interesting empirical question for the semanticist to verify, rather than an *a priori* truth to define the boundaries of what semantic research may be. Another advantage is that it prevents compositionality from being vacuously true. Partee (1984) observes that ‘if the syntax is sufficiently unconstrained and the meanings sufficiently rich, there seems no doubt that natural languages can be described compositionally’, as clever maneuvers in an arbitrary theory of syntax or lexical meaning can always rescue compositionality (see Janssen 1983; Zadrozny 1994). Thus it is necessary that we evaluate the validity of (2) against a backdrop of ‘well-motivated constraints on syntax and/or on the mapping from syntax to meaning’ (Partee, 1984). Doing so allows compositionality to serve both as a meaningful benchmark for a semantic theory and a meaningful guide for conducting semantic research.

The suggestion that compositionality is meaningful only in the context of constrained, non-arbitrary theories of syntactic structure, lexical meaning, and the mapping between them implies that ‘we need to study lexical semantics and principles of semantic composition together; decisions about either may affect decisions about the other’ (Partee, 2009, 20), and similarly for syntax. Attending closely to the interplay of these three domains, with specific commitments to the nature of lexical semantic units which may enter into the compositional

process, opens up the space for potential solutions to linguistic phenomena which have been argued to pose a fatal threat to a strong compositional thesis like (2). In particular, (2) requires that the meaning of a complex expression be determined by the meanings that its constituent expressions have individually. The inputs to semantic composition, then, should be predetermined: for the compositional process to determine what the meaning of one of its constituents is, in this particular sentential context, would constitute a violation of strong compositionality, a kind of feedback problem in which the meaning of a constituent is simultaneously determined by the meaning of the complex expression of which it is a part.

Taking seriously the study of meaning in language therefore requires detailed attention to the meanings of individual pieces of language (often, but not exclusively, words), the architectural rules for building syntactic structures out of those pieces, the compositional rules for building complex meanings out of simpler ones, and the mapping principles that determine how compositional rules are applied over syntactic structures. In doing so, the semanticist is well-equipped to look at particular phenomena in natural language from a compositional perspective and, in defining each of these modular but interactive components for that phenomenon, attempt to build a formal model of how the relevant meanings are defined within the bounds of our assumed guiding principle for compositional theory-building, whatever that may end up being. But as described, that project of examining whether or not meaning can always be described compositionally implies that there are problem cases, phenomena that on the surface appear less than strictly compositional. This dissertation, like many other works in this tradition, is an investigation of some of those problem cases.

1.2 Modification problems

In introducing the notion of compositionality, I suggested that were it not to hold, understanding language would be a nightmarish project verging on impossibility. Meanings need to be stable, or else we are lost at sea! But as any seasoned user of a language knows, in fact, variation is frequent, and while it sometimes impedes perfect communication, it certainly

doesn't grind the entire infrastructure to a halt. In particular, I suggested with the small verb-noun pairing example that it would be disastrous if the meaning of even a two-word pair were allowed to vary between multiple options, or if the meaning of a two-word pair were allowed to exclude the meaning of one of its parts. The focus of this dissertation is on instances of adjectival modification - equally simple pairs of one adjective and one noun - where, in fact, both of these compositional nightmares are realized, quite regularly and productively. These problems, labeled *ambiguity* and *privativity* respectively, have challenged strictly compositional theories for decades and each motivated theorists to either rescue the principle of compositionality with the introduction of significant covert complexity to our syntactic or semantic representations, or to abandon the principle altogether.

1.2.1 The problem of ambiguity

Sometimes, the combination of a single noun and a single adjective can result in more than one possible output meaning. A canonical example of modification-induced ambiguity is given in (3).

- (3) She is a beautiful dancer.
- a. 'She dances beautifully'
 - b. 'She is a dancer and physically beautiful'

The most frequent and salient interpretation of (3) is the first one in (3a), where it appears that the adjective is functioning actually more like an adverb, modifying the activity or event that the noun performs. More generally, this reading involves the adjective modifying the actual semantic content of the noun, whatever that may be; there is not always such an obvious verbal component to the noun (e.g., in *just king* it is the act of being a king that is just, but there's certainly no activity of *kinging*). But there is also the interpretation in (3b), a less obvious but certainly present reading, on which the subject is labeled as both a dancer and a beautiful, physically attractive person. On this reading, how beautiful her dancing is is entirely irrelevant to whether you would agree or disagree with the statement.

As meanings in semantics are often modeled as sets of objects, it is easy to visualize the meaning of (3b) as the result of intersecting two sets: the set of *dancers* and the set of *beautiful* people, where the *beautiful dancers* are the people found in the Venn diagram overlap. As a result, the reading in (3b) is often called the *intersective reading* (and will be regularly abbreviated throughout as the I reading). In contrast, the reading in (3a) is the *nonintersective reading* (NI reading). The contrast displayed in (3) will throughout be referred to as the (non)intersective ambiguity, and adjectives which display this ambiguity as (non)intersective adjectives.

It is important to appreciate how, already, we have the makings of a compositional crisis. The meaning of the whole sentence appears to be determined by more than just the meaning of its individual parts and the way they are combined, but also some third factor that distinguishes the nonintersective from the intersective readings. If we want to accept that the solution is truly not compositional, then we have any number of ways to resolve this problem. If we want to retain a compositional system, however, we have effectively two options: we can play with the meanings of the parts, or we can play with the way in which they are combined.

The most immediate answer, and certainly the one pursued by the earliest formal theories of the (non)intersective ambiguity, is to play with the meaning of the adjective. Doublet theories like Siegel (1976), and many following, simply state that adjectives like *beautiful* in fact represent two different words, with two distinct meanings, in a similar way to how a classically polysemous word like *bank* does ('financial institution' vs. 'side of a river'). One of those words means something like 'doing the activity in a beautiful way' and the other means 'physically beautiful', and that's that for ambiguity. Nothing about compositionality requires that there be a one-to-one mapping between the phonological form of a word and a single meaning; in fact, we really can't have such a principle, because of cases like *bank*.

But we need to do more than just defining two meanings, because the pattern is more complex than just (3). For example, one individual might be a beautiful dancer in both

senses. In that case we could describe them with:

(4) She is a beautiful beautiful dancer.

Of course, one reading of (4) is just emphasis, doubling down on one of the single meanings. But in the right context - a casting agency finding a group of elegant dancers but then furthermore needing one who is very attractive - each instance of the adjective provides one of the distinct (non)intersective meanings. And if you are a native English speaker, I didn't even have to tell you which adjective was providing which meaning - we naturally get the inference that the first *beautiful* is contributing the intersective meaning and the second the nonintersective one, and swapping them around really isn't possible. This is a problem for doublet theories, which then need to define an additional principle explaining this ordering effect. And it's not just this order that matters:

(5) She is a dancer more beautiful than any other.

This kind of adjective can't be put after the noun in English without some comparative support like *more than any other*, but what's important is that (5) *only* allows the intersective interpretation: this is a dancer who is more physically beautiful than any other dancer, dancing skill notwithstanding. These kinds of ordering effects are found across the English lexicon (Bolinger, 1967) and across languages with more generally flexible word order as well, as we will see in more detail later on (Larson and Takahashi, 2004; Cinque, 2010). A doublet theory not only needs to provide principles explaining both of these kinds of word order effects, but needs to generalize them for any language's version, and also needs to accept that the consistent recurrence of these doubled lexical representations for many of the same adjectives across so many languages is just an accidental coincidence, with no specific motivation for why the nonintersective and intersective meanings should both end up linked to the form *beautiful*, rather than one of them ending up with a different phonological realization.

Other approaches, which will be reviewed in the next chapter, take the opposite compositionality-

preserving solution, in trying to make more complex the syntactic structure combining the adjective and the noun, perhaps adding covert structure, or adjusting some other aspect of the principles of interpretation allowing for flexibility in a single adjective's denotation. In the end, this is the kind of strategy that this dissertation will pursue. However, it is important to note why it is so hard to make such solutions genuinely 'compositionality-preserving'. The whole utility of having a compositional system, as suggested above, is that it enables the people perceiving language to reliably understand what those sentences mean. If compositionality is preserved in a way that is impossible to see in the actual linguistic strings that people need to understand, then that solution is perhaps not very satisfactory to the goals behind the principle. This argument, and its applicability to theories of adjectival ambiguity, is put forcefully by Bouchard (2002):

'Compositionality has some methodological value only if it relies on fairly directly accessible surface properties (Partee 1997, 61, Hausser 1984, Hintikka 1983). But previous analyses of adjectival modification all have in common that they add various covert elements and operations to the theory (new lexical categories of adjectives or nouns, new syntactic categories, multiple lexical entries for some adjectives, new movement operations triggered by new functional features, etc.). The added elements are not accountable to either interfaces, so that compositionality is only satisfied in a technical way. Compositionality that is merely technical loses the motivation for the principle, i.e., to explain how a human being understands a newly encountered sentence, since we are never sure what elements are present, or what their individual meaning is.' (Bouchard, 2002, 9)

This notion of interface-accountability is worth keeping in mind throughout the ensuing discussion. Inevitably, we will have to introduce some complexity to both the syntax and semantics of the theory in order to capture all of the patterns that together constitute the (non)intersective ambiguity cross-linguistically. But I hope to do so in a way that is motivated and interface-accountable. To do so, an investigation of the (non)intersective ambiguity will be supplemented with a simultaneous investigation of another pernicious compositional problem, with the idea that each will shed some light on the other in order to illuminate a genuinely, not merely technically, compositional solution.

1.2.2 The problem of privativity

That second problem is the problem of *privativity*. Generally, the composition of an adjective and a noun - whether interpreted nonintersectively or intersectively - results in a composite meaning that involves both of the meanings of its subparts. In the simplest case, a *brown dog* is both *brown* and a *dog*, and in similar ways a *beautiful dancer* is both *beautiful* and a *dancer*, just on one of two possible interpretations of the adjective. Classically, semantic work has divided adjectives into a taxonomic organization depending on the kind of inferences from the modified noun phrase to the individual adjective and noun that they license. Intersective adjectives, like *brown* and one interpretation of *beautiful*, allow you conclude from an individual being an ‘adjective noun’ that they are both an instance of the adjective and an instance of the noun, separately.

(6) Intersective inferences

- a. $\frac{X \text{ is an } A_I N.}{X \text{ is } A.}$
- b. $\frac{X \text{ is an } A_I N.}{\frac{X \text{ is an } N'}{X \text{ is an } A_I N'}}$

Subjective adjectives, like the nonintersective interpretation of *beautiful*, don’t necessarily allow you to conclude that the individual has the independent property denoted by the adjective outside of the context of the noun, but you can certainly conclude that they are an example of the noun.

(7) Subjective inference

- a. $\frac{X \text{ is an } A_S N.}{X \text{ is an } N.}$

Some adjectives, however, do not allow either inference. This is the class of privatives, including *fake*, *counterfeit*, *mock*, *artificial*, and perhaps many others (Nayak et al., 2014). Privative adjectives, by definition, license the opposite of the subjective inference: for a given privative adjective A_P , if X is an $A_P N$, X is not an N . Contrasting with subjectives, the

set denoted by a privatively modified noun phrase is disjoint from the set denoted by the unmodified noun, and no set-theoretic operation over the unmodified noun set can derive the noun phrase set.

(8) This is a fake gun. \rightarrow This is not a gun.

(9) Privative inference

a. $\frac{X \text{ is an } A_P \text{ N.}}{X \text{ is not an N.}}$

Hence, the problem of privativity. On a standard, strong interpretation of compositionality, the meaning of a complex expression like a modified noun phrase is to be fully determined by the meanings of its constituent parts, following (2). In the case of intersective adjectives, this is satisfied straightforwardly: intersect the meanings. For other subsectives, the process may be slightly more complex, but it is still easily computable, since we are selecting a subset of the noun's meaning according to a standard provided by the adjective's meaning. For privatives, however, it appears that we must pull from beyond the extensional denotations of both the noun and the adjective to generate the denotation of the noun phrase. Of course, this is not much of a problem for the cognitive system to accomplish. We all readily compute the set of *fake guns* one way or another; its membership is no mystery to a competent language user. But how such computation is achieved by the compositional semantic system is significantly more opaque, given that no meaning contributed by *gun* seems to be preserved in the end result, at least on a standard semantic conception of extensional, referential meaning.

As a result, privative adjectives have received substantial attention in the frameworks of cognitive or conceptual semantics (Franks, 1995; Coulson and Fauconnier, 1999), because these inference facts appear initially to be more easily captured through an account of shifting senses or conceptual blending rather than a strictly compositional one. However, compositional accounts have also been offered, beginning with Kamp and Partee (1995) and continuing in Partee (2007, 2009, 2010), though I will argue later on that those accounts fall

short of providing sufficient explanation. The compositional problem offered by privative adjectives is easy to appreciate on an extensional conception of meaning, and their difficult ‘edge case’ logical behavior - which makes them seem to belong to a different class from most adjectives - has caused them to be excluded from most paradigm-building for the analysis of other properties of adjectives. In most cases, theories focused on the syntax-semantics interface properties of adjectives in particular have footnoted or entirely excluded consideration of privative adjectives, because of the apparent ‘extra’ complexity that their unique semantics introduces. That includes theories of the (non)intersective ambiguity just described.

The core strategy of this dissertation is, rather than footnoting the peculiar behavior of privative adjectives, to put the spotlight directly on them, and center their behavior for building a syntactic-semantic theory of adjectives, precisely because their peculiarities can be illuminating in quite helpful, previously unrecognized ways. In particular, the important starting observation is that some privative adjectives display precisely the same (non)intersective ambiguity as their subsective cousins like *beautiful*. For example:

- (10) This is a fake painting.
- a. ‘This is not a painting at all (e.g., this image was printed, not painted, or this object is actually the door to a secret safe)’
 - b. ‘This is a painting but some aspect of it is fake (e.g., it is a counterfeit painted by someone other than the claimed painter)’

Immediately, examples like this show that the privative inference pattern described above is not a universal property of these adjectives. Moreover, it aligns their behavior with subsective adjectives in a way that allows us to ask whether or not this privative ambiguity behaves the same with respect to syntactic configurations. And eventually, I will argue, investigating the (non)intersective ambiguity by specifically focusing on how it is realized in privative adjectives is not only fruitful, but entirely necessary for building an accurate theory of the ambiguity.

1.3 Theoretical assumptions

This dissertation will assume a trivalent intensional semantics with application and intersection as modes of composition for building meanings, a standard X-bar-theoretic syntax for building phrase structures, and the basic principles of Distributed Morphology for building words.

Semantics. Morphemes, including but not limited to words, denote terms in a logical language that builds functions out of a set of basic types. The language used here is Ty_2 , a higher-order intensional logic as defined by Gallin (1975) with a set of three basic types: e , the type of individuals, t , the type of truth values (0 or 1), and s , the type of worlds. Additional types may be constructed out of these using the binary type-constructor $\langle -, - \rangle$, which creates a function from the first type to the second. So, $\langle e, t \rangle$ is a type which is a function from individuals to truth values, and $\langle \langle e, t \rangle, t \rangle$ is a type which is a function from those functions into truth values, and so on. Being a function is represented in a lambda calculus, where the term λ denotes the application of a function to a term via variable substitution. A lambda term takes a particular type as input on its left half, and then returns a term of another type as output on its right have, divided by a $.$ or $[]$. So a function might look like:

$$(11) \quad \lambda x_e.\text{dog}(x)$$

Which is equivalent to:

$$(12) \quad \lambda x_e[\text{dog}(x)]$$

And in both cases represents the $\langle e, t \rangle$ -type function that takes individuals and returns true if they are dogs and false otherwise.

There are two ways in which terms in a lambda calculus can compose with each other. The first is Function Application, where a functional type (one surrounded by $\langle \rangle$) composes with a term of the appropriate type to serve as its input argument (e.g., e , in the case of

$\langle e, t \rangle$), and returns its output (in this case, something of type t). This operation is defined as follows:

- (13) **Function Application**
 If α is a branching node and $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket$ is a function whose domain contains $\llbracket \gamma \rrbracket$, then $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket(\llbracket \gamma \rrbracket)$.

For example:

- (14) $\llbracket \mathbf{Toby\ is\ a\ dog} \rrbracket = \llbracket \mathbf{dog} \rrbracket(\llbracket \mathbf{Toby} \rrbracket)$
 $= \lambda x_e[\mathbf{dog}(x)](\mathbf{Toby})$
 $= \mathbf{dog}(\mathbf{Toby})$

The other compositional operation available in the grammar is Predicate Modification, which models the intersection of sets of the same type. This operation is defined as follows, where σ is a variable over any type.

- (15) **Predicate Modification**
 If α is a branching node and $\{\beta, \gamma\}$ is the set of α 's daughters, and $\llbracket \beta \rrbracket$ and $\llbracket \gamma \rrbracket$ are both in $D_{\langle \sigma, t \rangle}$, then $\llbracket \alpha \rrbracket = \lambda x_\sigma. \llbracket \beta \rrbracket(x) = 1 \wedge \llbracket \gamma \rrbracket(x) = 1$.

This allows us to conjoin same-typed predicates to form complex ones, such as:


- (16) $\llbracket \mathbf{Toby\ is\ a\ brown\ dog} \rrbracket = \llbracket \mathbf{brown} \rrbracket \cap \llbracket \mathbf{dog} \rrbracket(\llbracket \mathbf{Toby} \rrbracket)$
 $= [(\lambda x_e[\mathbf{brown}(x)]) \cap (\lambda x_e[\mathbf{dog}(x)])](\mathbf{Toby})$
 $= \lambda x_e[\mathbf{brown}(x) \wedge \mathbf{dog}(x)](\mathbf{Toby})$
 $= \mathbf{brown}(\mathbf{Toby}) \wedge \mathbf{dog}(\mathbf{Toby})$

The role of an intensional semantics is to allow these functional statements to be interpreted relative to particular worlds, so a more thorough denotation for such a predicate might be given as:

- (17) $\llbracket \mathbf{dog} \rrbracket = \lambda w_s \lambda x_e[\mathbf{dog}_w(x)] = \lambda w_s \lambda x_e[x \text{ is a dog in } w]$

Both our definitions for Function Application and Predicate Modification make reference to 'branching nodes' which implies that these rules operate over meanings that stand in structural, syntactic relations to one another. So we also need a framework for representing those structures.

Syntax. Syntactic structures are created via the application of an operation Merge to two terminal nodes (retrieved from the lexicon) or previously constructed substructures. The result is that syntactic trees will always be at most binary branching. Most of the syntactic claims made here should be compatible with a variety of theoretical implementations; so the explicit assumptions made here will be minimal and basic. The basic type of phrase structures assumed here are those of X-bar theory, where syntactic heads (X) may optionally combine with a complement before projecting to an intermediate level (X') and then optionally combine with a specifier before projecting to the phrasal level (XP). Phrases may also adjoin to other phrases.



It is also assumed that pieces of syntactic structure may undergo movement (formally a variation on Merge) to other, higher locations in the structure. Every instance of Merge at the syntactic level will correspond to an instance of one of our two composition rules at the semantic level.

I will also be assuming some of the basic principles of Distributed Morphology (Halle and Marantz, 1993), which argues that the processes for composing word-internal structures are one and the same processes for composing words together into sentence-level structures. Particularly relevant later on will be the hypothesis of Late Insertion, the idea that the phonological forms of syntactic elements are introduced only at the interface with Phonological Form, not present during the syntactic derivation. Pieces in the syntax are fully abstract collections of syntactic features, and only gain utterable form via Vocabulary Insertion which occurs after a syntactic structure is built, such that facts about the syntactic environment

of a morpheme may influence its phonological realization but not the other way around. Vocabulary Insertion rules define a relationship between a syntactic element, like a root, and its phonological realization, by stating the syntactic context in which that phonological form is inserted:

$$(19) \quad \sqrt{\text{root}} \rightarrow \text{form} / [_ X]$$

The left side of the slash defines the insertion: a syntactic item (before the arrow) is realized as a phonological form (after the arrow). The right side of the slash defines the context for the application of the insertion rule: the underscore indicates the location of the syntactic item pre-insertion, and relates it to other features of the context; in this example, this rule would apply, inserting the *form*, whenever the root is in the local environment of some other element X.

1.4 Roadmap for the dissertation

This dissertation develops a compositional theory of adjectival modification in the noun phrase which (i) explains the distribution of intersective and nonintersective meanings by appealing to minimal, motivated semantic and syntactic machinery, and (ii) traces the compositional process across the morphosyntactic, syntax-semantic and the conceptual-semantic interfaces. The goal is to characterize what we might call the full compositional ‘pipeline’: starting with prelinguistic concepts, how they are converted into linguistically relevant chunks of meaning at the conceptual-semantic interface, how those chunks are put into structures and determine the possible meanings of those structures at the syntax-semantics interface, and how that affects surface realization at the morphosyntactic interface. In terms of how the argument will proceed, the strategy will be to begin with the syntactic distribution and working ‘downwards’ to sublexical semantic structures, treating each interface in its own chapter.

In less technical terms, the argument will proceed as follows. First, where do ambiguous

meanings appear? Chapter 2 addresses this descriptive, surface-level location of ambiguity: where in sentences can adjectives have ambiguous meanings, and where do they have to be unambiguous (in one direction or another)? In this chapter, I propose that, semantically, adjectives always take a local nominal argument, even when the overt noun they modify is nonlocal. Chapter 3 contains empirical motivation for the assumption that this local argument is realized with its own syntactic projection, rather than being purely pragmatic. After establishing the location of different meanings, the next question is, why do meanings appear in those locations? Chapter 4 proposes an account of the syntax-semantics interface to explain why certain locations in sentences are consistently associated with one meaning or another, linking positions to underlying semantic properties. Finally, how are those meanings derived? The previous chapters establish where we see intersective versus nonintersective meanings, and why they appear in those locations, but Chapter 5 focuses in on nonintersectivity in particular and asks what those adjectives and nouns actually mean, and how those meanings combine to give nonintersective interpretations. Following this where-why-how structure, then, the dissertation begins with a descriptive characterization of where (non)intersective meanings occur, then develops an account of why they occur there, then an account of how they are actually composed. Chapter 6 concludes.

As per (i) above, the goal is to appeal to a minimal set of independently motivated semantic and syntactic tools to account for ambiguity. In particular, I am going to argue that all we need to characterize the (non)intersective ambiguity is a mapping at the syntax-semantics interface of the preexisting distinction between object (or token) meanings and kind (or type) meanings, and the assumption that adjectives always take a syntactically represented kind argument. The crucial observation is going to be that nonintersective readings that are identical on the surface are, in fact, not always the result of the same compositional operations - by characterizing the path that meanings take from their introduction into the derivation through their modification by adjectives, and identifying the places at which that path can split, we can properly describe the compositional routes to nonintersectivity using

only these minimal assumptions.

The next few sections give more extensive previews of what each chapter covers.

1.4.1 Overview of Chapter 2

The focus of Chapter 2, ‘The distribution of ambiguity’, is to provide a general characterization of how intersective and nonintersective meanings are distributed across the syntax, and develop an initial theory of the mapping between those syntactic positions and the compositional semantic operations governing modification in each position. The formal semantics in this chapter is minimal and the syntactic account centers on linking surface word orders to underlying syntactic configurations, rather than on a precise characterization of the syntactic process involved in building those configurations.

The chapter begins with a review of the four broad past varieties of approaches to analyzing the (non)intersective ambiguity: theories that posit multiple denotations for the adjective, theories that posit multiple subparts of the noun to be modified, theories that locate the ambiguity in purely pragmatic, post-compositional processes, and theories that distinguish multiple syntactic domains of modification. Criticisms of all of these approaches are developed, but the final type - structure-sensitive syntax domain theories - is taken to be the most viable starting point for an accurate theory; specifically, the one developed in Cinque (2010).

The core empirical strategy of the chapter is privative disambiguation. I show that privative adjectives, which display the same (non)intersective ambiguity as the more commonly considered subsective class but have been generally avoided in paradigm-building for syntactic analysis, in fact are the key piece to illustrate the distribution of (non)intersective meanings across syntactic positions. Because of their highly restrictive semantics, allowing less pragmatic flexibility and fewer options for composition, they function as canaries in the mine of derivations: highly sensitive to their compositional environment and more prone to crashing the derivation when their strict semantics are not satisfied, which reveals the

underlying structure at play. Applying this strategy to Italian, I show that Cinque's assumption of surface interpretive ambiguity corresponding to underlying syntactic ambiguity is not supported, and instead propose a 'one-to-one' mapping between surface word order and covert syntax in Italian, additionally a one-to-one mapping between syntactic configuration and semantic composition, but a 'one-to-many' mapping between semantic composition and interpretation.

The same privative disambiguation strategy is then applied to Bangla, a language which allows focus-driven movement of the adjective, noun, or noun phrase to various higher positions in the extended nominal projection. I show that this movement - which has only been investigated semantically with respect to definiteness, but not modification - has interpretive consequences with respect to the (non)intersective ambiguity, and again that privative adjectives show a less ambiguous version of the same pattern, revealing again the same underlying syntax-semantics mapping principles in play in Italian.

Drawing on both the Italian and Bangla data, the chapter develops the first component of the theory of the (non)intersective ambiguity: (non)intersective adjectives are not polysemous, but rather always a higher-type, two-place predicate that requires an initial nominal argument to determine the scale over which the adjective is interpreted. Intersective readings are not the result of a lexically-listed first-order property predicate version of the adjective, but rather the adjective first combining with a null argument to receive a noun-independent semantics before intersecting with the overt modified noun. Syntactically, this intersective composition process takes place when the adjective is non-local to the overt noun, and nonintersective modification takes place when the adjective is local to the noun. The fact that both non-local composition via intersection and local composition via application can, under the correct pragmatic circumstances, lead to nonintersective readings on the surface establishes that there are multiple compositional routes to nonintersectivity.

1.4.2 Overview of Chapter 3

The focus of Chapter 3, ‘The morphosyntax of intersection’, is to establish that the null argument involved in intersective readings of ambiguous adjectives is realized with a distinct syntactic projection, and the intersective adjective is therefore created compositionally in the course of the syntactic derivation, rather than being a merely pragmatic operation. Evidence for the syntactic reality of this argument is found in cross-linguistic patterns of suppletion.

Adjectives like *good* across languages display both the interpretive ambiguity of interest here - nonintersective ‘good at...’ vs. intersective ‘good person who...’ - and often the morphophonological phenomenon of suppletion, where the phonological form of their stem changes to an unpredictable alternate in certain morphological contexts. Specifically, the comparative and superlative forms of *good* often display suppletion, like English *better* and *best*. This chapter begins with an observation from Despić (2019) that this suppletion, despite being generally regarded as a semantically irrelevant process, has interpretive consequences in Serbian: adjectives which are ambiguous in their positive form actually lose their intersective meaning when they are in the comparative or superlative, but only if those forms are suppletive. I review Despić’s account of this pattern, which proposes two distinct roots for these ambiguous adjectives, and argue that it is untenable.

Specifically, the diradical account is unable to capture the fact that this pattern in Serbian in fact appears to be a substantial cross-linguistic generalization. Across a survey of sixteen languages from multiple language families, I find that in almost all cases, if a comparative/superlative form is suppletive, it loses the intersective reading, but if it is morphologically regular, it retains it. This pattern, termed the Nonintersective Suppletion Generalization, motivates a monoradical account on which the intersective reading is blocked for suppletive adjectives because it requires additional syntactic material intervening between the adjectival root and its suppletion trigger, the comparative affix, given that suppletion must be triggered in a local relationship. This accounts for the general unavailability of the intersective reading across languages without positing stipulative and arbitrary Vocabulary

Insertion rules, and provides independent phonological evidence that there is covert syntactic structure involved in the intersective reading, supporting the conclusion from Chapter 2.

The one exception to the Nonintersective Suppletion Generalization is found to be Basque. Basque, interestingly, is also the lone known exception to Bobaljik's (2012) *ABA generalization, where the comparative form is suppletive but the superlative is regular. While this chapter concludes without an explanation of either exceptional pattern, it does argue that the coexistence of these exceptions may provide additional evidence that Basque's comparative construction is fundamentally unlike other languages', and might be correctly treated as not a syntactic comparative in such a way that its abnormal behavior with respect to both generalizations is expected. Current accounts of the Basque morphological exceptionality, however, fail to predict its semantic exceptionality.

1.4.3 Overview of Chapter 4

The focus of Chapter 4, 'Number and kind modification', is on making precise the syntax-semantic interface properties that cause adjectives to display ambiguous meanings in different syntactic domains. Here, I develop the core semantic theory of the dissertation: the ambiguity between nonintersective and intersective meanings of adjectives is reducible to the distinction between kind-level and object-level denotations of the nouns that they are modifying.

The chapter begins by introducing the semantic notion of kinds, some background on how to maneuver with and between kinds and objects, and a particular approach relating kinds to grammatical number. On this view, nouns uniformly enter the derivation denoting kinds, and they are converted to object-level denotations, properties of individuals, when they compose with number (either morphological number marking or a classifier), which serves the semantic function of making them countable. This creates two distinct syntactic domains with specific semantic ontological properties - one corresponding to kinds and one to objects. I then pick up on an idea from Leffel (2014), who argued that adjectives may be

sensitive to this distinction and compose with the noun either before or after this number-driven object-conversion, which for Leffel marks the distinction between individual-level and stage-level interpretations of adjectives like *visible*. I draw out some criticism of Leffel's account, namely its continued reliance on polysemy and lack of empirical support, before pursuing a revision to resolve both problems.

The chapter then turns to Turkish, where bare singular (or unmarked for number) nouns show a particular interaction with kind/object modification - while these nouns are sometimes ambiguous between singular and plural interpretations, adding object-level modifiers forces singular interpretations while kind-level modifiers preserve the plural possibility. I argue that we can use this pattern as a diagnostic for kind-sensitivity, and apply it to both individual/stage contrasts and (non)intersective contrasts. I find that in both cases, the diagnostic reveals that these forms of modification are in fact sensitive to kind structures: individual-level and nonintersective interpretations behave like kind modifiers, while stage-level and intersective interpretations behave like object modifiers. This motivates assimilating the contrasts, which also finds support in the earlier Bangla data, where it is the classifier level that divides the relevant domains for nonintersective modification.

The analysis developed in this chapter, built off of Leffel's initial proposal, holds that these adjectives are uniformly functions over subkind predicates, taking one kind-level denotation and directly modifying the kind to create a new subkind. While this kind of modification is in effect always nonintersective, it can occur either with the overt modified noun prior to enumeration - to create surface nonintersective readings - or with the null argument, a pragmatically provided subkind predicate (whose syntactic existence was argued for in Chapter 3), which in combination with a linking operator both saturates the subkind argument of the adjective and converts it into an object-level predicate appropriate for intersection with the overt noun post-enumeration into an object-level property.

1.4.4 Overview of Chapter 5

The focus of Chapter 5, ‘Inside nonintersective modification’, is to elaborate on the process of modifying a kind that takes place inside Chapter 4’s proposed composition for nonintersective modification. It does so via a focus on the possible interpretations of privative adjectives and the variation they show across contexts and nouns.

The chapter begins with a reintroduction of the problem of privativity, and Partee’s classic account of privative interpretations as the result of coerced expansion of noun meaning triggered by violation of an interpretive principle that prohibits the otherwise contradictory result of privative composition. I argue that, though this account is fundamentally correct in assimilating privative adjectives to subsective adjectives in a basic sense, it fails to adequately predict privative meanings, in particular the wide variation that so-called privatives show in the actual intersectivity of their outputs.

The results of two experiments are then presented to illustrate the range of possible meanings for privative adjectives.

- Experiment 1 focuses on the way in which the meanings of nouns can change across contexts depending on the other nouns they compete with, via an image-selection paradigm modeled after ‘captcha’ questions. The results suggest that there is significant variation in how strictly people apply the categorization requirements of a noun, sometimes requiring a strict interpretation and other times loosening it to require only a single property of the noun’s kind, e.g., its perceptual form, but that such modulation is not unlimited and cannot involve the creation of novel categorization requirements, only weakening or strengthening of lexically encoded ones.
- Experiment 2 focuses on the variation in the meaning of different privative adjectives across nouns, and specifically what kind of inferences they license. It finds that adjectives like *fake* are not uniformly nonsubsective, but depending on the noun they compose with in fact regularly license subsective inferences, and this depends on the

particular interaction of the noun’s semantic content and the property modified by the adjective.

The experimental results are used to motivate a theory of privative modification under which nouns are imbued with a rich set of conceptual features as part of their semantic representation, and individual adjectives can modify distinct components of those conceptual representations. This Dual Content Semantics theory of Del Pinal (2018) enriches the lexicon with both conceptual structure and a modulation function that naturally captures the behavior observed in both experiments, with some proposed revisions. On this view, privative adjectives are not a uniform class at all, and labeling categories of adjectives based on the inference that they license is always going to be inaccurate because inferences are not lexical properties but emergent properties of composition based on the particular interaction of a noun’s conceptual content and the adjective’s targeted features to modify.

Finally, I propose the assimilation of Del Pinal’s enriched conceptual structure for nouns with the semantic notion of kinds, to treat kinds as sortal concepts composed of some grammatically-structured subset of the information found in a given concept. The end of the chapter sketches a loose proposal for how these compositional conceptual structures might be imposed onto underspecified conceptual semantics for syntactic roots. On this view, roots index concepts, identifying the locations in memory where the bundle of non-linguistic conceptual information can be found; nominal categorizers impose the grammatical structure of a kind onto those roots, retrieving a particular subset of that conceptual information and rendering it in linguistically interpretable form; grammatical number converts that kind structure into a set of objects.

2 The distribution of ambiguity

The goal of this chapter is to sketch the initial structure of a proposal for how intersective and nonintersective meanings of ambiguous adjectives are distributed across syntactic structures. It will be focused on the broad strokes in both syntax and semantics: distinguishing classes of approaches from each other, characterizing syntactic domains as ‘local’ vs. ‘non-local’, and employing basic schematic denotations for noun and adjective semantics. In Chapter 4, I will develop a specific account of what adjectives and nouns denote, and what constitutes the relevant senses of locality; here, I am focused entirely on contrasting one type of structural theory of ambiguity (surface-syntax) from another (semantic-interpretive). First, though, I will introduce some non-structural theories and cover why we not to pursue those avenues.

2.1 Past theories

This section introduces some examples of preexisting approaches to the ambiguity problem in the literature. This is only a very limited subset, but these examples have been selected because they are generally representative of what I see as four predominant strategies, of which many other theories are variations on a theme: adjectival polysemy, internal complexity of the noun, pragmatic resolution, and distinct syntactic domains of modification. In the end, I want to suggest that each of these approaches both fails to capture the empirical pattern and offers a bit of useful insight that the present proposal will draw from.

2.1.1 Two types of adjectives (Siegel 1976)

Probably the simplest and most intuitive account of the (non)intersective ambiguity in adjectives just holds that there are in fact two homophonous forms of the adjective, one for each reading. This style of account takes the ambiguity at face value and encodes it in the lexical representations of the adjective(s), and as a result has been called the ‘Blame the Adjective’ approach in the ensuing literature.

The best-known of these theories, and maybe the earliest to formalize this intuition in Montagovian formal semantics, is that of Siegel (1976). Siegel argues that the lexicon is fundamentally divided into two semantically and syntactically distinct categories, which have been falsely treated together under the umbrella of ‘adjective’.

Some adjectives are entirely predicative, entirely extensional, and only allow intersective readings. In Siegel’s early notation, these adjectives are $t///e$, taking an individual argument of type e and returning a truth value t ; in modern notation, $\langle e, t \rangle$ (I’ll be giving the examples below in more contemporary lambda notation, not Siegel’s original Montague notation). Whether syntactically serving as adnominal modifiers or predicates, they uniformly combine with nouns via intersection and therefore only involve extensional composition. This class includes adjectives like *sick*, *infinite*, *portable*, *nude*, *tall*, *aged*...

$$(20) \quad \begin{aligned} \llbracket \mathbf{sick} \rrbracket &= \lambda x.\text{sick}(x) \\ \llbracket \mathbf{sick\ friend} \rrbracket &= \lambda x.\text{sick}(x) \wedge \text{friend}(x) \end{aligned}$$

Other adjectives are fully attributive, semantically CN/CN : functions from ‘common noun’ denotations to another common noun denotation, or $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$. These adjectives are also intensional, creating a new property out of their input property rather than merely manipulating extensions. This class includes *alleged*, *former*, *veteran*, *rightful*, *chief*...

$$(21) \quad \llbracket \mathbf{alleged} \rrbracket = \lambda P.\lambda x.\text{alleged}(P)(x)$$

The ambiguity of adjectives like *beautiful* or *good*, on this account, is nothing more than homophony. When an adjective is ambiguous, it has two forms in the lexicon: one predicative and one attributive. For this reason, Siegel’s account has often been called a ‘doublet theory’.

$$(22) \quad \begin{aligned} \text{a. } \llbracket \mathbf{beautiful} \rrbracket_{t///e} &= \lambda x.\text{beautiful}(x) \\ \llbracket \mathbf{beautiful\ dancer} \rrbracket_I &= \lambda x.\text{beautiful}(x) \wedge \text{dancer}(x) \\ \text{b. } \llbracket \mathbf{beautiful} \rrbracket_{CN/CN} &= \lambda P.\lambda x.\text{beautiful}(P)(x) \\ \llbracket \mathbf{beautiful\ dancer} \rrbracket_{NI} &= \lambda x.\text{beautiful}(\text{dancer})(x) \end{aligned}$$

It is easy for a doublet theory like Siegel’s to capture variation in adjective meaning. Introducing lexical alternations is a powerful mechanism, and allows for any differences or

ambiguities to be encoded directly as another form of the adjective in the lexicon. However, constraining this powerful of a system requires especially stipulative restrictions on adjective distribution. Siegel discusses alternations between adjective meaning in adnominal and predicate position, but does not discuss the kind of adnominal word order variations that this dissertation focuses on, and I think that her account would struggle significantly with them. On her view, there are adjectives which are uniformly nonintersective (CN/CN), and as a result only appear in adnominal position and cannot be predicative. There are also adjectives which are uniformly intersective (t///e), and as a result only appear in predicative position and cannot be adnominal. Identifying these lexical types is straightforward, but accounting for their behavior with respect to different syntactic configurations requires significantly more stipulated restrictions, and accounting for their regular and consistent cross-linguistic ambiguity can only be accidental.

An updated version of Siegel’s account is offered more recently by Despić and Sharvit (2011) and Despić (2019), which will be discussed at length in Section 3.1. As they basically preserve the same core idea of having two distinct meaning options for an ambiguous adjective, however, many of the same objections apply.

2.1.2 Two parts of nouns (Larson 1998)

While Siegel’s undermining of the unified semantic category of ‘adjective’ does actually play nicely with some more contemporary syntactic lines of thought similarly questioning the utility of such a primitive category label (e.g., Mitrović and Panagiotidis 2020), the consequent lexical duplication that her implementation requires has been often seen as too severe of a downside for an adequately parsimonious theory. For every adjective like *beautiful* that displays both nonintersective and intersective readings, we need to posit that an accidental homophony exists with two distinct lexical entries corresponding to a single phonological surface form. Not only that, but this accident needs to occur repeatedly across languages, for a conveniently similar set of meanings in each language. Type-shifting rules could also

capture the alternation, but their range of application would need to be captured in an equally coincidental way, or would constitute abandoning Siegel’s account entirely. While possible, accepting this kind of large-scale conspiracy would be an unfortunate result from a theoretical perspective. And yet, upon rejecting it, we are still left with the reality that this wide array of adjectives systematically displays two distinct meanings.

One approach to resolving this is to instead locate the ambiguity in the meaning of the noun. Of course, done simply, this would duplicate the precise lexical redundancy problem that Siegel runs into for adjective-based ambiguity. If every noun that participates in these ambiguous constructions has two lexical entries, then what we have is likely an even more severe case of widespread accidental homophony. To avoid this, Larson (1998) develops an account on which the source of the ambiguity is the noun, but there is only one denotation of a given noun; that denotation is just internally complex and contains within it the seeds of both meanings.

The crucial move is to assume that nouns (at least deverbal nouns like *dancer*) have two argument slots - one for the individual, and one for the event, such that the individual is the agent of the event. The denotation of such a noun might be represented (in different notation from Larson’s original, but I hope the same spirit) as:

$$(23) \quad \llbracket \mathbf{dancer} \rrbracket = \lambda e. \lambda x. \text{dancing}(e) \wedge \text{agent}(e, x)$$

An adjective like *beautiful*, then, is optionally interpreted as modifying either the individual variable x or the event variable e of the noun, resulting in the intersective and nonintersective meanings, respectively. In either case, *beautiful* simply denotes that the modified variable is beautiful relative to some comparison class C , given by context.

$$(24) \quad \begin{array}{l} \text{a. } \llbracket \mathbf{beautiful dancer} \rrbracket_{\text{I}} = \lambda e. \lambda x. \text{dancing}(e) \wedge \text{agent}(e, x) \wedge \text{beautiful}(x, C) \\ \text{b. } \llbracket \mathbf{beautiful dancer} \rrbracket_{\text{NI}} = \lambda e. \lambda x. \text{dancing}(e) \wedge \text{agent}(e, x) \wedge \text{beautiful}(e, C) \end{array}$$

On this story, then, all of the meanings are really intersective. The only distinction is which aspect of the nominal semantics the adjective is intersecting with. This also preserves a unified semantic category of adjectives as predicates. However, a fundamental concern with

this type of account is a compositional problem: how does a single adjective compose with distinct parts of the nominal semantics? Or, more precisely, how does the adjective have compositional access to whatever the innermost variable of the noun is? In the examples above, I've ordered the variables such that the noun takes its event argument first, which matches up with standard syntactic accounts of such nominals (e.g., Alexiadou and Schäfer 2010), though in principle the compositional problem exists with either decision. Larson allows the adjective to modify each variable of the noun simply by defining distinct rules for each type of combination. But from a strictly compositional perspective, without recourse to individual meaning postulates, this approach reinforces the fundamental 'bracketing paradox' of nonintersective meanings: given that the adjective composes with a fully-formed noun, the internal event variable should be opaque and inaccessible for intersection.

The empirical coverage of Larson's account has also been challenged, with difficulties identified in extending it to a wider variety of nouns and adjectives; see Maienborn (2020) for such arguments.

Despite these concerns, the attractiveness of this account, which avoids the complications of multiple adjective denotations, has led to a number of subsequent expansions and reformulations (Winter and Zwarts, 2012; Egg, 2008). The denotations above, for example, are more in line with the version developed by Winter and Zwarts (2012), which tries to avoid the compositional problems by extrapolating the syntactic structure of the noun and allowing multiple attachment heights for the adjective. This basic idea of multiple nominal domains in which the adjective can reside will end up being quite valuable for the present proposal, in fact, though in a less event-oriented way. However, others have taken a more decisive stance against the noun-focused analysis.

2.1.3 Pragmatic variable resolution (Maienborn 2020)

Developing past 'Blame the Adjective' stories, Maienborn offers a contemporary account which locates the source of ambiguity within the adjective without positing actual polysemy

or multiple lexically listed denotations for the adjective. Maienborn tackles the compositional problem of nonintersective meanings by simply assuming that such meanings are, in a real sense, non-compositional. The adjective is still taken to be modifying the internal content of the noun, conceptually, but this modification occurs in entirely pragmatic processes and does not involve any actual interaction of the adjective and the nominal content in the course of the semantic derivation: ‘the ambiguity is not derived in the course of composition but only shows up at the semantics-pragmatics interface.’ (Maienborn 2020:71)

Mechanically, this is implemented through the use of *tropes*, a concept introduced to the philosophical literature by Williams (1953) and imported to formal semantics largely by Moltmann (2007, 2009, 2013), who describes them as follows:

Tropes are the particular manifestations of properties in individuals. Unlike properties conceived as universals, tropes depend on a particular individual and thus are generally taken to involve a particular spatio-temporal location. Tropes differ from properties also in that they are causally efficacious and can act as the objects of visual perception. Typical examples of tropes are ‘Genji’s handsomeness’ and ‘Socrates’ wisdom’. (Moltmann 2013:10)

In contrast to a state, which as an object is entirely abstract, a trope always involves a particular realization and is therefore fully dependent on the particular individual in which they are realized. A trope, therefore, intrinsically has a *bearer*. Maienborn’s account of the (non)intersective ambiguity hinges on this fact, and argues that it is always a trope to which the adjective attributes its description, which results indirectly in modifying the trope’s bearer. Exactly what trope (and therefore, what property), of the presumably many that any given bearer has, which is being modified is underspecified, and it is at this underspecification that the ambiguity is located. But this trope is not compositionally active, and underspecification is resolved entirely via pragmatics.

The denotation for an evaluative adjective like *beautiful* would then be:

$$(25) \quad \llbracket \mathbf{beautiful} \rrbracket = \lambda y_{\text{ENTITY}} [\text{bearer}(r_{\text{TROPE}}, y) \ \& \ \text{beautiful}(r)]$$

Per Maienborn, ‘*beautiful* qualifies its compositional target argument y as being the bearer

of some trope r that is judged as beautiful. Providing a contextually plausible specification for r is the task of pragmatics.’ (Maienborn 2020:72) Thus it is always the individual entity which is the compositional semantic argument of the adjective, but it is a particular property of that entity which is actually ascribed beauty, and so the notions of syntactic argument and semantic modification are both dissociated from a notion of conceptual modification.

This is the maximally underspecified version of a denotation for *beautiful*, but the range of possible subtypes of tropes that might exist is wide. This fully underspecified, maximal type of TROPE maybe more appropriate for an adjective like *good*, but we might want to narrow the possibilities for *beautiful* to only SENSORY-TROPE (though that might be too narrow, depending on whether we want to capture more abstract uses directly or via metaphor), while more specific adjectives target even more specific tropes: ‘expensive evaluates the price of its target (r_{PRICE}), elegant its style (r_{STYLE}), fragrant addresses olfactory properties ($r_{\text{OLFACTORY-TROPE}}$) and garish visual properties ($r_{\text{VISUAL-TROPE}}$), etc.’ (Maienborn 2020:72) The formal system relating these conceptual types to each other is to some extent independent of the proposal, and to an even greater extent not relevant to our current consideration of it, but for completeness of exposition it is intended to be couched in the Type Composition Logic of Asher (2011). It might be considered this type system’s responsibility to effectively constrain the possible trope subtypes that a given adjective can be realized with; Maienborn’s proposal does so in only a limited way but it is unclear empirically to what extent we would like stronger constraints, a question to which we will return at length in Chapter 5.

Composition in this account is straightforward and uniformly intersective, with the trade-off of introducing additional complexity to the nominal semantics in the form of *social roles* (Zobel, 2017). The kind of (frequently but not exclusively deverbal) nominals that generate nonintersective meanings with evaluative adjectives are argued to associate their individual entity arguments with particular social roles, culturally ascribed abstract entities that are manifested as characteristic activities that the individual performs. These *role nouns* in Maienborn’s framework take an individual argument and hold that it is the bearer of a cer-

tain social role, with social roles modeled as subtypes of tropes. A simple example is a noun like *king*, which denotes an entity that bears the ‘king’ social role (which, implicitly, it has a result of coronation and thus participates in characteristic activities like ruling a kingdom, but such properties are neither spelled out in the denotation nor compositionally active).

$$(26) \quad \llbracket \mathbf{king} \rrbracket = \lambda x_{\text{HUMAN}} \exists r'_{\text{SOCIAL-ROLE}} [\text{bearer}(r', x) \ \& \ \text{king}(r')]$$

The deverbal event nouns that Larson’s framework centralizes are represented similarly as role nouns, but with an eventive component: ‘They assign an individual a social role that manifests itself in those activities that are referred to by the verbal root.’ (Maienborn 2020:73)

$$(27) \quad \llbracket \mathbf{dancer} \rrbracket = \lambda x_{\text{HUMAN}} \exists r'_{\text{SOCIAL-ROLE}} \text{GEN}_{e_{\text{EVENT}}} [\text{bearer}(r', x) \ \& \ \text{manifest}(r', e) \ \& \ \text{dance}(e) \ \& \ \text{agent}(e, x)]$$

Denotations like (27) contain three variables: the individual referent x , their social role r' , and the verbal event e . However, only the first is compositionally active; both r' and e are bound internal to the nominal and inaccessible for any form of modification (in the syntactic or semantic sense). The fact that only the individual entity is active as an argument in both (25) and (27) allows them to combine via simple intersection to yield (28).

$$(28) \quad \llbracket \mathbf{beautiful dancer} \rrbracket = \lambda x_{\text{HUMAN}} \exists r'_{\text{SOCIAL-ROLE}} \text{GEN}_{e_{\text{EVENT}}} [\text{bearer}(r', x) \ \& \ \text{manifest}(r', e) \ \& \ \text{dance}(e) \ \& \ \text{agent}(e, x) \ \& \ \text{bearer}(r_{\text{TROPE}}, x) \ \& \ \text{beautiful}(r)]$$

The social role r' of being a dancer, defined as a trope that manifests in being the agent of dancing events, is bound by the existential operator, and so compositionally inactive. The unspecified trope r , which is being labeled as beautiful in some sense and of which the individual argument is the bearer, is not bound in any way, and so also compositionally inactive but still requiring further specification. The maximal trope, TROPE, or perhaps the slightly more constrained SENSORY-TROPE (if that is what we end up wanting for *beautiful*), are insufficiently precise for interpretation and thus ‘calls for pragmatic specification’ (Maienborn 2020:74). Exactly why something like SENSORY-TROPE forces further pragmatic specification but something like SOCIAL-ROLE or VISUAL-TROPE doesn’t is unclear - what

level of taxonomic specificity in the type hierarchy is enough to leave a trope be? This is likely to simply be a question of invoking appropriate world knowledge to come up with a valid reading, but it is worth noting that it is not spelled out in Maienborn’s account.

How does the pragmatic specification process for the trope r proceed? Certain constraints are placed on it in the course of composition. The adjective provides type restrictions, though in this case *beautiful* offers not much in the way of such restriction, given that the trope need only be a subtype of TROPE, allowing therefore any trope. The noun may also provide type restrictions, as less of a lexical specification, but more as a result of specifying the referential argument that is the bearer of r : the eventual value of r must naturally be the kind of property that can be instantiated in the kind of referent that x is. The two readings that we are interested in deriving - the intersective ‘physically beautiful & dancer’ and the nonintersective ‘dances beautifully’ - are on this account both the result of intersective modification, but arise from different trope specifications. The intersective reading results from specifying the trope r as the physical appearance of its individual bearer x (through a function ‘phys-appearance’ which maps objects to that trope), as in (29a). The nonintersective reading, in contrast, results from the even simpler specification that sets $r = r'$, possible because SOCIAL-ROLE \sqsubseteq TROPE, as in (29b).

- (29) a. $\llbracket \text{beautiful dancer} \rrbracket_{\text{I}} = \lambda x_{\text{HUMAN}} \exists r'_{\text{SOCIAL-ROLE}} \text{GEN}_{e_{\text{EVENT}}} [\text{bearer}(r', x) \ \& \ \text{manifest}(r', e) \ \& \ \text{dance}(e) \ \& \ \text{agent}(e, x) \ \& \ \text{bearer}(r_{\text{TROPE}}, x) \ \& \ \text{beautiful}(r) \ \& \ r = \text{phys-appearance}(x)]$
 b. $\llbracket \text{beautiful dancer} \rrbracket_{\text{NI}} = \lambda x_{\text{HUMAN}} \exists r'_{\text{SOCIAL-ROLE}} \text{GEN}_{e_{\text{EVENT}}} [\text{bearer}(r', x) \ \& \ \text{manifest}(r', e) \ \& \ \text{dance}(e) \ \& \ \text{agent}(e, x) \ \& \ \text{bearer}(r_{\text{TROPE}}, x) \ \& \ \text{beautiful}(r) \ \& \ r = r']$

The interpretation of (29a) should be straightforward: the individual x bears a social role of being a dancer, and also instantiate the property of their own physical appearance, which is described as beautiful. Interpreting (29b) requires somewhat more thought. The individual x bears the social role of being a dancer, and that role is also beautiful. What does it mean for the trope describing their social role to be beautiful? The social role manifests in being the agent of dancing events, and our intuitive description of the nonintersective

reading would be that those events are beautiful, but (29b) instead gives us that the role itself is beautiful. How does the latter end up being interpreted as the former? Maienborn elaborates:

‘What is then the link between being beautiful or elegant as a dancer and dancing beautifully/elegantly? On the present view, this link corresponds to an additional inference based on default knowledge. We know that qualities and judgments applying to social roles typically carry over to the activities by which they manifest themselves – always provided that the selectional restrictions of the pertinent predicate are met.’ (Maienborn 2020:75)

This intuition is implemented in Type Composition Logic with defeasible, weak conditional rules relating predicates of social roles to events manifesting those social roles:

$$(30) \quad \text{bearer}(r, x) \ \& \ \text{TYPE}(r) \sqsubseteq \text{SOCIAL-ROLE} \ \& \ \text{manifest}(r, e) \ \& \ \text{agent}(e, x) \ \& \ \text{TYPE}(e) \\ \sqsubseteq \text{ACTIVITY} \ \& \ Q(r) > Q(e)$$

This rule ‘accounts for inheriting predicates over social roles to their characteristic activities’, and ‘will eventually lead to the manner interpretation of an adnominal modifier’. (76) The term $Q(r) > Q(e)$ in particular ensures that ‘predicates Q that apply to r typically carry over to e’ (75). The rule in (30) is the most generic, but specific rules mapping specific social roles to their characterizing activities are also necessary, and together they will result in the mapping of predicates of social roles to predicates of specific characterizing activities for those roles. In this way, the term ‘beautiful(r')’ in (29b) transfers its meaning of ‘the social role of being a dancer is beautiful’ to ‘the characteristic activities of a dancer are performed beautifully’.

But the choice of $r = r'$ is only one of innumerable possible choices to specify the meaning of the trope variable. Not only do we have potential access to any subtype of the maximal type TROPE, but we also have the ability to set r equal to any other trope which is, like r' , defined in terms of other lexical items and available functions like ‘phys-appearance’. So why do these other options never seem to obtain in language use as natural interpretations of *beautiful dancer*? Interpretations like ‘dancer who smells beautiful’, which are readily definable given the existence of OLFACTORY-TROPE needed in this system for

words like *fragrant* and that OLFACTORY-TROPE \sqsubseteq TROPE, do not occur even with attempts at strong contextual support. Even the much more intuitively plausible interpretation which Maienborn labels the ‘actual nonintersective reading’, wherein an individual is currently engaged in an act of beautiful dancing (independently of generally having the dancer role or not), is strongly dispreferred compared to either the intersective reading or the ‘dispositional nonintersective reading’, which we have been considering as the standard nonintersective reading where someone is disposed to dance beautifully. This latter ‘actual NIR’ can occur, and can technically be achieved by defining r as the agent role for a salient event e' , but the former type (beautiful in some wholly other sense) does not¹. How do we constrain the pragmatic specification system to degrade and rule out these possibilities, respectively?

Maienborn appeals to a generally plausible principle of pragmatic interpretation:

- (31) Preference principle for the pragmatic specification of free variables:
Free variables are instantiated preferentially by linguistically introduced material, always provided that all the given requirements are met. (Maienborn 2020:78)

Per (31), the interpretation in (29b) is maximally parsimonious, since it references only the lexical content of the noun. The intersective reading in (29a) is less parsimonious because it also requires the invocation of a linguistically available but not linguistically introduced function, ‘phys-appearance’, from general world knowledge. Even less preferred are the ‘actual NIR’ and the ‘beautiful-smelling’-type readings, which require the contextual support of a particularly salient event which has not been linguistically introduced.

The upshot of this analysis is that a purely pragmatic disambiguation process, where all readings are compositionally identical throughout the derivation, avoids the compositional obstacles inherent in trying to let the adjective modify some internal component of the noun, and preserves a single lexical entry for the adjective without the need for type-shifting operations to achieve different meanings. All that is required is an intuitively plausible pragmatic

¹Whether such an interpretation exists or not is actually not clear - it certainly does for some ambiguous adjectives like *good*; see (71). For *beautiful* specifically, it’s not clear that the same flexibility is enabled, which may be due to its more limited range of application due to lexical constraints on the kind of scale it produces.

principle of interpretation to appropriately rank the empirically likely interpretations over the empirically unlikely ones. From a semantic perspective, intersection is the only required operation. In the face of the compositional difficulties that past semantic and syntactic analyses have encountered, this approach is undeniably elegant.

However, the source of its elegance - minimal compositional interaction of the adjective with its nominal argument - is also what makes it fundamentally untenable as a way to capture the empirically observed distribution of ambiguous meanings. There are any number of possible representations of intersective and nonintersective meanings. The core challenge, in my view, is to select the representation which correctly predicts not merely which interpretations are possible but where those interpretations will occur. As covered elsewhere in this chapter, the syntactic position of the adjective relative to the noun consistently has effects on the interpretations allowed in that position, across various languages. The intuition of Larson which causes the ‘bracketing paradox’ that this account attempts to resolve - namely, that the more noun-relative, nonintersective meanings occur when the adjective is in a more syntactically local relationship to the noun - is wholly opaque on this fully pragmatic view. How does a non-compositional account offer any insight into the effects of word order in Romance languages on the availability of different readings? This analysis specifically divorces the readings from any process of composing the adjective with the noun, which prevents that compositional interaction from determining where a given reading will be available. In terms of explaining why prenominal and postnominal adjectives behave differently with respect to (non)intersective readings, a fully post-compositional pragmatic account appears to be simply a non-starter.

For the sake of argument, let’s see what we would need to do to make the relevant syntactic predictions with this account. It is immediately clear that the current version of (31) is not going to be sufficient, with a preference for ‘linguistically introduced material’ as the only requirement. Whether an adjective is prenominal or postnominal, or syntactically more local or more distant from the noun, has no effect on the noun’s status as ‘linguistically

introduced'. In all possible syntactic configurations, the noun is going to qualify as linguistically introduced material, and so the pragmatic specification that makes reference to the lexical contribution of the noun - the dispositional NIR - is going to be the preferred option for trope variable resolution. This accurately predicts the general preference for the dispositional NIR, but nothing about the varying prominence of the intersective reading in different positions. For example, English postnominal adjectives are restricted to their intersective reading:

- (32) Olga is a dancer more beautiful than her instructor.
 a. 'Olga is more physically attractive than her instructor.'
 b. #'Olga dances more beautifully than her instructor.'

This restriction is unexplained by the pragmatic preference principle, as is the strict preference for the nonintersective reading of prenominal adjectives in Italian:

- (33) Un buon attaccante non farebbe mai una cosa del genere (Cinque 2010:10)
 a good forward not would-do never a thing of-the kind
 'A good forward would never do such a thing'
 a. 'A forward good at playing forward would never do such a thing' (nonintersective)
 b. #'A good-hearted forward would never do such a thing' (intersective)

In both cases, (31) should identify the noun as linguistically introduced material and preferentially choose its lexically encoded trope as the reference of the adjective's underspecified trope variable. The fact that all of the pieces of semantic machinery involved in determining the eventual reading are compositionally inactive makes it impossible to force one reading or another as a result of the way the adjective combines with the noun, including such facts as word order.

Of course, (31) does include an out that prevents the NIR from being preferred in all situations: 'provided that all the given requirements are met'. If we want to preserve the pragmatic strategy, we would then need to define - presumably in semantic terms - the requirements for the variable resolution options that lead to each reading, and link those requirements to specific syntactic relations such that they are only met in the appropriate

positions in a given language. For this to be explanatory, then, we would need the requirements to stem from facts about the syntactic relationship between the adjective and the noun in a given position. So, to make use of this provision, we need a fully developed theory of how syntactic position determines the interpretation - but, unless we want to abandon the whole purpose of having a post-compositional pragmatic mechanism, this theory needs to not involve the way in which the adjective and the noun combine with each other differently in different positions. If the theory of ‘given requirements’ appeals to anything about the noun’s semantic content being more available as a target for reference resolution in a given position, then we no longer access any of the parsimony benefits motivating the pragmatic strategy in the first place.

Rather than appeal to this provision, we could instead make an adjustment to the notion of ‘linguistically introduced material’. Perhaps, when resolving the reference of an underspecified variable, we first have a preference for material that is in a specific syntactic relationship - maybe a kind of locality, maybe a kind of argument relation, whatever the empirical syntactic investigation ends up determining - with the adjective. After that, the next most preferred option for determining reference would be linguistically overt material that is not in such a syntactic relationship, and then the worst option is contextually-supported but not linguistically present material. What this option would result in is a pragmatic theory that makes explicit reference to particular syntactic configurations, but in a sense leaps over semantics (because, again, involving a syntax-semantics interface property would make it compositional) - our pragmatic principle of interpretation looks directly at syntactic structures and effectively duplicates that information by reencoding the different positions in interpretive preferences. At this point, if we are defining a pragmatic principle in terms of precise syntactic factors, what is the point of including the pragmatic level at all? It seems to be a massive redundancy.

Beyond its struggle with syntactic effects, I want to also argue that the pragmatic account fails to distinguish adequately between the ‘actual NIR’ and the other contextual-support-

requiring interpretations, on its own terms. The goal of (31) is to establish a preference for readings that refer to only linguistically introduced material. But on what grounds does the dispositional NIR refer more to linguistically introduced material than the actual NIR? Both invoke different senses of the noun, e.g., *dancer*. For the dispositional reading, someone who dances generally or has the role of dancing; for the actual reading, someone who is currently dancing. Intuitively or technically, why is the former reading of *dancer* ‘linguistically introduced’ by the word, and the latter not? Given that the entire motivation for Maienborn’s account is to avoid compositional interaction with the particular semantic content of the noun, it seems difficult to distinguish these two on the basis of a principle like (31) without making some claim about what exactly the ‘material’ of the noun introduces.

Even more difficult to see is how (31) should rank the default intersective reading over the actual NIR. Surely the inference that someone is currently dancing is more supported by the linguistic material *dancer* than the default intersective inference that they are physically beautiful is, since physical beauty isn’t any particular part of the meaning of the adjective *beautiful* on this account, any more than any other kind of beauty is. The strong empirical preference for the default intersective reading over the actual NIR seems to run starkly contrary to the predictions of (31), unless we add some dimension of meaning to *beautiful* such that it constitutes linguistically introduced support for this intersective meaning over others.

And similarly, why should the default intersective reading be preferred over others that require contextual support on this pragmatic story? In both cases, no linguistically introduced material supports any particular option for variable resolution for the trope in *beautiful* - it’s not picking up on the noun *dancer*, but rather resolving to some other value not in the sentence. A distinct principle would be needed to explain why the selection of the physical appearance trope value is so strongly preferred over any other one, or in the case of *good* why the moral goodness resolution option is preferred over any other particular skill besides the one denoted by the noun. In general, it seems to me that (31) is substantially underspeci-

fied for the purpose of adequately ranking and distinguishing the possible readings, with its imprecise notion of ‘linguistically introduced’ and how to link linguistic material to variable resolution options that refer to that material.

In sum, a purely pragmatic approach struggles to capture the interaction of ambiguity with syntactic structure absent recourse to compositional explanations, and without any mechanism for syntax to determine meaning it also struggles to give substance to any principle of interpretation that tries to define a ‘material’ link between the adjective and the noun. While it may avoid both the lexical duplication of a Siegel-style analysis and the particular compositional puzzles raised by a Larson-style analysis by abandoning compositional machinery entirely, losing the explanatory power of compositional operations as a result ends up being an unwelcome tradeoff. What we need is a system that appeals to structure in a meaningful way.

2.1.4 Syntactic ambiguity (Cinque 2010)

One such account is developed by Cinque (2010, 2014) and centers around the idea of different structural positions for adjectives corresponding to different modes of interpretation, with an empirical focus on contrasts between Germanic (particularly English) and Romance (particularly Italian) word order options in the noun phrase. In both language families, adjectives sometimes have the option of appearing either prenominal or postnominal.

- (34) a. Olga is a more beautiful dancer than her instructor.
b. Olga is a dancer more beautiful than her instructor.

- (35) a. Olga è una bellissima ballerina.
Olga is a beautiful dancer
‘Olga is a beautiful dancer’
b. Olga è una ballerina bellissima.
Olga is a dancer beautiful
‘Olga is a beautiful dancer’

English postnominal adjectives tend to require an explicit complement or other overt comparative element (*as beautiful as any*, for another example), a restriction not present

in Italian. What Cinque claims is that these positions correspond in a regular manner to specific interpretive properties of the adjectives, in that the position to some extent determines which of the ambiguous meanings will surface. This claim goes well beyond the intersective vs. nonintersective contrast that is the focus of the present work, arguing that these positions are systematically associated with a wide variety of interpretive properties that cluster together in a uniform manner.

I will walk through each proposed property shortly, but the overarching claim is consistent: in both languages, there is a fully ambiguous position and an unambiguous position. The languages are the inverse of each other, however, in both which position is ambiguous and what meaning is present in the unambiguous position. In English, the prenominal position is ambiguous and the postnominal position unambiguous; in Italian, the prenominal position is unambiguous and the postnominal position ambiguous. All of the meanings that show up in the English postnominal position are impossible in the Italian prenominal position, and vice versa. We can illustrate with the (non)intersective contrast:

- (36) Olga is a more beautiful dancer than her instructor.
 a. ‘Olga dances more beautifully than her instructor.’ (nonintersective)
 b. ‘Olga is more physically beautiful than her instructor.’ (intersective)
- (37) Olga is a dancer more beautiful than her instructor.
 a. #‘Olga dances more beautifully than her instructor.’ (nonintersective)
 b. ‘Olga is more physically beautiful than her instructor.’ (intersective)

In English, prenominal *beautiful* allows both the nonintersective and intersective interpretations, while postnominally only the intersective is allowed.

- (38) Un buon attaccante non farebbe mai una cosa del genere (Cinque 2010:10)
 a good forward not would-do never a thing of-the kind
 ‘A good forward would never do such a thing’
 a. ‘A forward good at playing forward would never do such a thing’ (nonintersective)
 b. #‘A good-hearted forward would never do such a thing’ (intersective)

- (39) Un attaccante buono non farebbe mai una cosa del genere
 a forward good not would-do never a thing of-the kind
 ‘A good forward would never do such a thing’
- a. ‘A forward good at playing forward would never do such a thing’ (nonintersective)
 - b. ‘A good-hearted forward would never do such a thing’ (intersective)

In Italian, the situation is inverted. The postnominal position allows both meanings, while the prenominal position is unambiguous but allows only the nonintersective. This same pattern of asymmetrical ambiguities purportedly holds for other contrasts. For example, the contrast between individual- and stage-level meanings (as noted by Ferris 1993; Sadler and Arnold 1994; Svenonius 1994):

- (40) The visible stars include Aldebaran and Sirius.
- a. ‘The stars that are generally visible include Aldebaran and Sirius’ (individual-level)
 - b. ‘The stars that happen to be visible now include Aldebaran and Sirius’ (stage-level)
- (41) The stars visible are Aldebaran and Sirius.²
- a. #‘The stars that are generally visible include Aldebaran and Sirius’ (individual-level)
 - b. ‘The stars that happen to be visible now include Aldebaran and Sirius’ (stage-level)
- (42) Le invisibili stelle di Andromeda esercitano un grande fascino
 the invisible stars of Andromeda have a great fascination
 ‘The invisible stars of Andromeda have a great fascination’
- a. ‘Andromeda’s stars, which are generally invisible, have a great fascination’ (individual-level)
 - b. #‘Andromeda’s generally visible stars, which happen to be invisible now, have a great fascination’ (stage-level)
- (43) Le stelle invisibili di Andromeda sono moltissime
 the stars invisible of Andromeda are very many
 ‘The invisible stars of Andromeda are very many’

²This is the judgment from Cinque (2010); some English speakers including myself share the judgment that unmodified postnominal visible is somewhat degraded, compared to something like *The stars visible from Earth...*, or at least that the unmodified version needs to be embedded in substantial linguistic context to support this reading.

- a. ‘Andromeda’s stars, which are generally invisible, are very many’ (individual-level)
- b. ‘Andromeda’s generally visible stars, which happen to be invisible now, are very many’ (stage-level)

Again, English prenominal *visible* is ambiguous while postnominally it allows only the stage-level meaning; Italian is the inverse in both which position is ambiguous and what the unambiguous meaning is. Intuitively, the stage-level meaning (which is the unambiguous option for English) does seem to correspond conceptually to the the intersective meaning (available in the same position): they are both describing less essential, noun-centered properties, and so it may be unsurprising that they pattern together.

Another example that is stated to follow the same pattern is the restrictive versus non-restrictive contrast (Larson and Marušič, 2004). An adjective may be interpreted restrictively, to constrain the domain to a specific subset of the noun set, or nonrestrictively, to indicate descriptively that all of the members of the noun set share that adjectival property. For example:

- (44) Every blessed person was healed.
 - a. ‘All the people were healed’ (nonrestrictive)
 - b. ‘All the people that were blessed were healed’ (restrictive)

- (45) Every person blessed was healed.
 - a. #‘All the people were healed’ (nonrestrictive)
 - b. ‘All the people that were blessed were healed’ (restrictive)

- (46) Le noiose lezioni di Ferri se le ricordano tutti
 the boring classes of Ferri remember all
 ‘Everybody remembers Ferri’s boring classes’
 - a. ‘Everybody remembers Ferri’s classes, all of which were boring’ (nonrestrictive)
 - b. #‘Everybody remembers just those classes by Ferri that were boring’ (restrictive)

- (47) Le lezioni noiose di Ferri se le ricordano tutti
 the classes boring of Ferri remember all
 ‘Everybody remembers Ferri’s boring classes’
 - a. ‘Everybody remembers Ferri’s classes, all of which were boring’ (nonrestrictive)
 - b. ‘Everybody remembers just those classes by Ferri that were boring’ (restrictive)

Again, the English prenominal position is fully ambiguous while the postnominal position is unambiguous, here allowing only the restrictive meaning. The same inversion occurs in Italian as before, with the unambiguous prenominal allowing only the nonrestrictive meaning.

Besides these, Cinque lists a wide range of interpretive properties which he purports to be shared uniformly across these positions, all showing the same ambiguity asymmetry. Some additional properties are summarized in the table below.

(48)

Italian postnominal, English both	Italian both, English prenominal
intersective	nonintersective
stage-level	individual-level
restrictive	nonrestrictive
relative	absolute
epistemic ‘unknown’	evaluative ‘unknown’
discourse anaphoric ‘different’	NP-dependent ‘different’
literal	idiomatic
relative ‘possible’	modal ‘possible’
specific or nonspecific	specific

Cinque’s analysis of this pattern is that the interpretively ambiguous positions are syntactically ambiguous. There are two options for adjectival modification, which following Sproat & Shih (1994) are called direct and indirect modification. In a direct modification configuration, the adjective is phrasal and merges as a specifier of a functional head in the extended projection of the NP. In an indirect modification configuration, the adjective is in the predicate position within a reduced relative clause, and that clause merges also as a specifier of a functional projection (above NP and below DP), but one that is higher than the direct modification projections. So direct modification is distinguished from indirect by (i) the absence of a reduced relative clause containing the adjective and (ii) a closer merge position to the noun.

These syntactic options are associated uniformly with the interpretive properties listed above. Namely, those properties which appear unambiguously in the Italian prenominal position are properties of direct modification: adjectives in a direct modification configuration are always interpreted as nonintersective, individual-level, nonrestrictive, and so on.

By contrast, properties which appear unambiguously in the English postnominal position are of indirect modification: adjectives in an indirect modification configuration are always interpreted as intersective, stage-level, restrictive, and so on. This correspondence is strict, with few exceptions that have dedicated explanations (discussed below).


As a consequence of this one-to-one correspondence between syntactic configurations and interpretations, any time that more than one interpretation for an adjective is available in a particular (linear, surface) position, we are forced to conclude that that (linear, surface) position corresponds to more than one potential syntactic structure. As a result, the English prenominal position can allow both direct modification and indirect modification adjectives, as can the Italian postnominal position. But English postnominal adjectives are always only indirect (clausal), and Italian prenominal adjectives are always only direct (phrasal).

Cinque develops in detail a series of syntactic arguments for how the surface word orders are derived from these underlying positions; for example, since both indirect and direct modification adjectives merge above the noun initially, postnominal word orders must be derived through movement of the NP attracted to a null complementizer merged above the adjective-hosting functional projections. These arguments, and the debate between phrasal movement of the NP and head movement of the N, for example, are orthogonal to the debate here, and so I will not cover them. As far as I understand, this choice won't impact the basic claims that I want to make about different merge domains for adjectives; one way or another the correct word order must be derived after the adjectives are introduced to the derivation.

I'm also going to remain agnostic with respect to the need for dedicated functional projections for each type of adjective, as is advocated in the cartographic tradition of which Cinque's work is a part. What is important for present purposes is the distinction between the more highly-merged indirect modifiers and the closer-the-noun direct modifiers; whether you think that there are distinct functional projections hosting each class of adjective (Color, Nationality, Size, Shape, etc.) is not going to be a decision that features in the present discussion. That is largely a question of distinguishing certain direct modifiers from other direct

modifiers, and indirect from indirect.

What does matter here is the syntactic distinctions *between* direct and indirect modifiers. As just stated, there are two: the phrasal/clausal distinction, and the low/high merge distinction. A number of objections have been raised to Cinque's strong claim that indirect modifiers are always embedded in reduced relative clauses, which he attempts to extend to a wide range of languages, and in many cases the arguments for clausal structure in those languages have not appeared to hold up (e.g., Talić 2017). As far as the syntax-semantics mapping is concerned, however, this aspect is actually not so critical to Cinque's account: it appears that the merge height distinction is where he actually wants to draw the crucial dividing line for why direct and indirect modification result in different semantic outcomes, though admittedly the explication of this process is underspecified in what is intended to be a primarily syntactic work: 'The question of why direct and indirect modification adjectives have the cluster of interpretive properties that they have, rather than the opposite, is a deeper question, and one to which I cannot offer a definite answer.' (Cinque 2010:33). What is offered is the proposal that direct modifiers are merged below a 'small indefinite dP' which assigns 'some referential import (though not the uniquely individuating referential import of the higher D' (Cinque 2010:34):



As (49) illustrates, again, indirect modification adjectives (AP_{IM}) are (i) contained within a reduced relative CP and (ii) separated from the noun by this ‘small’ dP, while direct modification adjectives (AP_{DM}) are (i) independently phrasal and not contained in a CP and (ii) below dP with NP. There may be any number of FPs, each hosting a distinct class of adjective, with I suppose the implication that each adjectival functional projection in the cartographic view occurs twice, above and below dP.

At this point, I want to start developing some objections to this system. The most significant and novel motivation for overhauling Cinque’s account of (non)intersectivity specifically will be introduced in the remaining sections of this chapter, but prior to that we can sharpen some of the other problems it faces. It’s worth first noting the significant advantage that this proposal has over the others discussed so far this chapter: by attending precisely to syntactic structure, it allows at least the possibility of compositionally deriving the observed word order effects. There is good reason why I will be basing the present proposal on a reconfiguration of Cinque’s, rather than the others.

However, what we gain in syntactic precision - if the account is genuinely precise - we

lose out in semantic specification. What exactly is the semantic contribution of this dP layer? The explanation given is minimal, and so this is less an objection to the particular way it is spelled out and any potential consequences of that theory, and more an objection via omission - it's largely unspecified what the semantics of this layer is meant to be. We are told that it is indefinite, but given the need for it to be compatible with definite full Ds, what sense of 'indefinite' can it really be contributing? Presumably not anything like the semantics of a genuine indefinite article. Beyond that, all that we have is that d assigns *some* referential import which is less than that assigned by D. I think that, fundamentally, this direction is correct - we will end up wanting some syntactic layer that converts the lower-level predicative NP into something that is more suitable for reference. But this particular implementation lacks either sufficient development to really evaluate from a semantic perspective, or independent motivation for its existence beyond the need to distinguish indirect from direct modification. It also appears to somewhat duplicate the role of the reduced relative clause in distinguishing those types of modification.

While the intuition behind d, shared by plenty of prior semantic theories, is likely to turn out correct, an aspect of Cinque's theory that is less likely, in my view, to stand up to scrutiny is the strict association between the long list of interpretive properties that are purported to surface in direct and indirect modification positions. On this strict one-to-one syntax-semantics mapping, it is argued that all properties in the same columns of (48) should be necessarily coextensive in distribution - it should be impossible to find a sentence where a single adjective displays properties from both lists. However, they can in fact be dissociated.

- (50) Everyone respects this district's good lawyers.
- a. '...this district's lawyers, all of whom are skilled at law'
 - b. '...this district's lawyers, all of whom are moral people'
 - c. '...just those lawyers in this district who are skilled at law'
 - d. '...just those lawyers in this district who are moral people'

The sentence (50) displays four distinct interpretations, a two by two cross of the intersective/nonintersective contrast and the restrictive/nonrestrictive contrast. In the reading

(50a), *good* is interpreted both as nonintersective - good at being a lawyer - and nonrestrictive - all the district's lawyers are good, the adjective is merely a description. In (50b), *good* is read intersectively - a good person - but still nonrestrictive. The latter two (50c) and (50d) are both restrictive - it is only those specific district lawyers who are also good that receive respect - but *good* is interpreted nonintersectively in (50c), and intersectively in (50d). If the strict association given by Cinque's account were true, where one syntactic structure can only produce nonintersective and nonrestrictive semantics, and the other can only produce intersective and restrictive semantics, readings like (50b) and (50c) should be impossible. Supporting this conclusion, Leffel (2014) has also argued at length that the restrictivity contrast should be treated as distinct from the other ambiguities.

Not all of the properties in (48) are amenable to this kind of test. Since several of them are particular to one adjective - 'unknown', 'different', and 'possible' - and these adjectives do not display multiple ambiguities, it's not obviously possible to test whether those contrasts necessarily coincide with e.g., (non)intersectivity.

There is one instance where Cinque explicitly acknowledges that this strict correlation fails, as you can see in (48): while in general, direct modification adjectives are supposed to be interpreted as specificity-inducing and indirect modification adjectives as non-specific, there are some instances where indirect modification can be specific. The normal behavior is observable in the English prenominal position where both types of modification are possible:

- (51) John will burn a nearby house.
- a. 'John will burn some specific house that is near his' (specific)
 - b. 'John will burn some house or other among those that are near his' (nonspecific)

The abnormal behavior is that the same two interpretations are found in postnominal position, which should be exclusively indirect and therefore nonspecific.

- (52) John will burn a house nearby.³
- a. 'John will burn some specific house that is near his' (specific)

³*nearby* in this configuration also has a potential adverbial reading, but we can control for that by embedding, e.g., *John thinks that a house nearby will burn*, and the same ambiguity remains.

- b. ‘John will burn some house or other among those that are near his’ (nonspecific)

Cinque argues that this special case - the availability of reading (52a) - is due to the availability of that reading when *nearby* is in the predicate position of an overt relative clause, which again allows both.

- (53) John will burn a house that is nearby.
- a. ‘John will burn some specific house that is near his’ (specific)
 - b. ‘John will burn some house or other among those that are near his’ (nonspecific)

Since the indirect modification structure has the adjective in the predicate position of a reduced relative clause, it would follow that it should have the readings of the same position of an overt relative clause. However, this line of reasoning runs into two issues. First, more generally, it implies that adjectives in the predicate position of an overt relative clause should *always* display the same readings as those in a reduced relative clause. It is not at all clear that this is true. For example:

- (54) We need a surgeon who is good.

The indirect modification, reduced relative clause interpretation of *good surgeon* is intersective (‘morally good’), and so the overt relative clause in (54) should only allow that reading. But in fact, it seems like it even prefers the nonintersective reading (which is, of course, the more salient reading in general).

One source of confusion regarding Cinque’s predictions in this area stems again from the semantic role of dP in demarcating the zone of indirect modification. If it is really dP that matters for determining interpretation and not the relative clausal structure surrounding indirect modifying adjectives, then we shouldn’t expect any kind of meaningful correlation between reduced and overt relative clauses to be attributable to their being relative clauses. Of course, if overt relative clauses and reduced ones are always in the same position relative to dP, then we would - but then the availability of a specific reading in (53) isn’t an *explanation* for (52a). It’s just the same anomaly that requires an explanation: why does *nearby* allow a specific reading when it’s past dP? The reduction of the relative clause shouldn’t affect that

interpretive conundrum one way or another.

The second, more targeted objection to this explanation of the specificity contrast in particular is that the ambiguity doesn't seem to be located in the adjective itself anyway. The specific/nonspecific alternation for both (51) and (52) can be reduced straightforwardly to a standard scope ambiguity for the indefinite DP. The specific reading results from a wide scope interpretation of the indefinite; the nonspecific reading from narrow scope. It's not surprising that changing the position of the adjective within the DP would have no effect on the interpretation if it's the position of the whole DP that matters.

The goal of raising these particular criticisms regarding the correlation of interpretive properties is to narrow the desired empirical coverage of our theory. By drawing clear empirical lines distinguishing the restrictivity and specificity contrasts from (non)intersectivity, we can establish that a revision of Cinque's structural theory should actually not have as part of its mandate to account for those patterns. Additionally eliminating the patterns that are specific to a single adjective, we are left with the following contrasts: intersective vs. nonintersective, stage- vs. individual-level, absolute vs. relative, and literal vs. idiomatic. A single theory of the syntax-semantics mapping for modification may reasonably endeavor to capture these as manifestations of a single underlying alternation. While the remainder of the chapter will still focus on (non)intersectivity, and develop the more major objection motivating a reanalysis of this mapping, the theory developed later on in Chapters 4 and 5 may be applicable to these other contrasts as well.

2.2 Privative disambiguation in Italian

This section introduces what you might consider the core empirical strategy of the dissertation: taking advantage of the particular properties of privative nonsubsective adjectives in order to reveal the underlying semantic environment of different adjective positions. While these adjectives have generally been disregarded in the syntactic literature on these word order phenomena (often called Bolinger contrasts, after Bolinger 1967), they have the poten-

tial to be extremely illuminating. In particular, in this section I will argue that the behavior of privative adjectives in Italian constitutes strong evidence against Cinque’s theory of syntactic ambiguity described in the previous section, and motivates a reanalysis in terms of underlying semantic ambiguity.


Before getting into the weeds of privative adjective behavior, though, it is worth a quick note on what is meant by ‘privative’, which will be relevant at a number of points in the rest of the dissertation. As discussed in Section 1.2, historically, privative is one of the lexical classes of adjectives, a subtype of the nonsubsective adjectives where the denotation of the AN combination is nonoverlapping with the initial N denotation. This includes adjectives like *fake*, *counterfeit*, *mock*, *artificial*, and others. The question of whether our theory should actually include such a category will be a major focus of Chapter 5, and so I do not want to suggest by using the phrase ‘privative adjectives’ that I am committing to a taxonomy of modifiers that encodes such a class, lexically or otherwise. For terminological convenience, in this chapter and until otherwise changed, I will be using ‘privative adjectives’ uncritically to refer to those adjectives that have been canonically grouped together as such and, in their most common behavior with most nouns, have the property of nonoverlapping composition.

In fact, the original context in which Cinque (2014) discusses the behavior of privative adjectives in Italian - the data with which this chapter will be most concerned - is evaluating whether or not it is correct to say that there is a class of genuine privative adjectives, and the fact that their behavior is not uniformly privative is going to feature heavily in the present discussion. But I will hold off on drawing any conclusions for adjectival taxonomy until Chapter 5. For now, we will only be interested in the conclusions for the syntactic theory that we can draw from the unique behavior of this so-called privative class.

2.2.1 Canaries in the mine

Recall from the previous section the general pattern of ambiguity in Italian adjectives: prenominally they are uniformly interpreted nonintersectively, and postnominally they are

ambiguous between nonintersective and intersective interpretations. This asymmetry motivates Cinque’s theory that the prenominal position is syntactically unambiguous (solely NP-local, direct modification) and the postnominal position ambiguous (between direct modification and indirect modification from within a reduced relative clause). This pattern is schematically represented again below in (55).



Privative adjectives behave differently. To illustrate, we first need to observe that ambiguity is possible with privatives.

- (56) This painting is fake.
- a. Nonsubsective: ‘This is not a painting (e.g., a piece of plastic hiding a safe)’
 - b. Subsective: ‘This is a painting but in some respect faked (e.g., a forgery)’

This ambiguity is quite like the nonintersective/intersective ambiguity that is found in subsective adjectives like *beautiful*. For the (a) reading, the adjective modifies some core part of the nominal semantics, which here results in the painting itself being what’s fake, analogous to the ‘dances beautifully’ reading of *beautiful dancer*. With a privative adjective, this results in a nonsubsective reading, though we might also call it - correctly but less precisely - the nonintersective reading to group it with the other adjectives. For the (b) reading, the adjective modifies some other property, which here results in a painting which has had its painter faked, analogous to the ‘beautiful appearance’ reading of *beautiful dancer*. Similarly, this reading ends up being both subsective and intersective; the former is a more useful label to contrast with the specifically nonsubsective reading of the privative, but the latter is more useful to group it with other adjectives behaviorally. I will proceed by calling the (a) reading nonintersective and the (b) reading intersective, for more effective comparison.

The interesting behavior is that, in the case of privatives, the (non)intersective ambiguity is distributed differently from that of subsective adjectives. Where subsectives are asymmetrical, ambiguous in one position and unambiguous in the other, privatives are symmetrically unambiguous. The prenominal position allows only the nonintersective reading, and the postnominal position allows only the intersective reading.

- (57) Un falso quadro ornava la parete
 A fake painting adorns the wall
 ‘A fake painting was adorning the wall’
- a. Nonintersective: ‘not a real painting; something painted on the wall to resemble a real painting’
 b. #Intersective: ‘a forged painting; a copy of the original’
- (58) Un quadro falso ornava la parete
 A painting fake adorns the wall
 ‘A fake painting was adorning the wall’
- a. #Nonintersective: ‘not a real painting; something painted on the wall to resemble a real painting’
 b. Intersective: ‘a forged painting; a copy of the original’

And where normally adjectives in predicate position allow both interpretations, predicate position privatives allow only their intersective meaning:

- (59) Quel quadro è falso ‘That painting is fake’
- a. #Nonintersective: ‘That painting is not a real painting’
 b. Intersective: ‘That painting is forged’

If no plausible intersective meaning is available, they are simply infelicitous both postnominally and in predicate position:

- (60) a. Quello è un falso problema
 ‘That is a false problem’
 b. *Quello è un problema falso
 c. *Quel problema è falso

Cinque identifies this distinction between privatives and subsectives but offers no discussion of how it should be accounted for. How could it be explained in his system, the one represented in (55)? Given that there is a one-to-one correspondence between underlying

syntactic structure and semantic interpretation, the inability of a particular interpretation to hold in a given surface linear position necessarily suggests the impossibility of the corresponding covert structure in that position. Because privative adjectives lack the postnominal nonintersective interpretation that subsecutive adjectives have, our immediate conclusion should be that privative adjectives cannot occur in direct modification configurations postnominally. Absent introducing additional stipulations, this is the only way to exclude a particular interpretation on the (55) schema.

We also know, however, that privative adjectives can be in direct modification configurations in prenominal position, since they can receive nonintersective interpretations in that position and there is again a supposed one-to-one correspondence between nonintersectivity and direct modification. The difference in surface order is a difference of noun movement: recall that direct modification adjectives are argued to always Merge into the same functional projection, which is base-generated above the noun and can optionally end up realized postnominally as a result of noun raising. Therefore, in order to exclude postnominal direct modification privative adjectives, what we actually need is a restriction on noun raising across direct modification privative adjectives. I am not aware of any natural way to motivate such a restriction. Further, that restriction would not be sufficient on its own to exclude nonintersectively-interpreted privative adjectives entirely in predicate position.

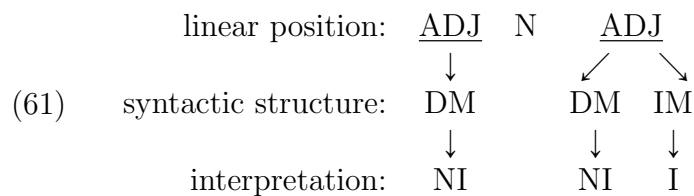
What I want to argue instead is that the unavailability of nonintersective readings for privatives in postnominal position is not a consequence of a syntactic structure being unavailable, but instead falls out from the particular compositional semantics inherent to the postnominal position and its interaction with privative lexical semantics. In short: standard subsecutive adjectives display ambiguity in the postnominal position not because of an underlying syntactic ambiguity, but instead because their lexical semantics is extremely flexible and - in a sense - that flexibility makes it resistant to constraints that the syntax may attempt to place on composition in a particular position. Privative adjectives, on the other hand, have a particularly inflexible meaning because of their non-overlapping semantics, and

this means that their interpretations transparently reflect the compositional constraints that the syntax places on them in any given position. Attempting to flexibly interpret a privative adjective in a way other than is furnished by the compositional mechanism in its place will lead to the derivation crashing in a way that subsecutive adjectives can avoid. This is what I mean by the section title: privative adjectives function for us as canaries in the mine, highly sensitive to their environment, and the loss of their interpretation in a given position can tell us useful information about what is going on there, in a way that subsecutive adjectives are too interpretively resilient to detect.

2.2.2 One-to-one syntax, one-to-many semantics

The data from privatives suggests that the ambiguity schema in (55) is not accurate. I want to argue that this is because an underlying syntactic ambiguity in postnominal position is not motivated. Importantly, in Cinque’s original formulation, the primary motivation for this syntactic ambiguity was the existence of an interpretive ambiguity in postnominal position. There are not, as far as I can tell, independent syntactic arguments for this structural ambiguity which then map conveniently onto the interpretive ambiguity and motivate treating the latter as due to syntax. The interpretive ambiguity is the sole argument for the structural one. Given that, in the face of the privative disambiguation, we might find ourselves skeptical of positing a postnominal structural ambiguity if we can formulate a theory that captures the privative vs. subsecutive alternation with fewer structural assumptions. In this section, I will describe such a theory.

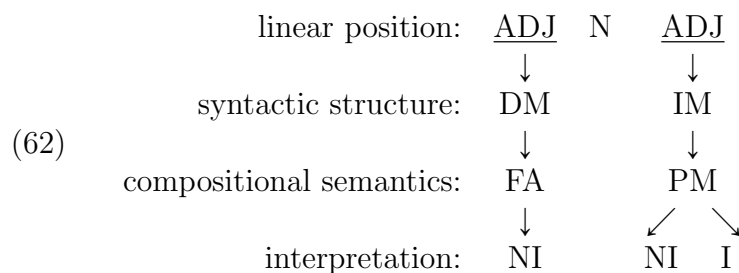
For comparison, I will reproduce (55) here as (61).



The critical features of this schema are (i) a one-to-many mapping between surface, linear

position and syntactic structure and (ii) a one-to-one mapping between syntactic structure and semantic interpretation. These assumptions run into clear problems with the privative pattern. Instead, I propose that we can get a clearer picture of the pattern by splitting apart *compositional semantics* from *interpretation*, two levels which the Cinque account implicitly collapses and treats together. Here, I will use ‘composition’ to specifically refer to the compositional semantic operation that takes place - namely, Function Application or Predicate Modification - and ‘interpretation’ to refer to the particular reading that obtains of the sentence or phrase - here, nonintersectivity or intersectivity. For this discussion, then, it would be incoherent to refer to ‘nonintersective semantics’, since that would conflate the compositional semantic level, which describes the logical properties of a position, with the surface interpretation, which describes the reading that speakers of the language have access to and provide in a judgment.

Against this backdrop, I want to argue that the privative adjectives data gives us good motivation for assuming that the actual ambiguity is in the mapping between composition and interpretation, which we can call one-to-many, while the mapping between linear position and syntactic structure is in fact one-to-one, as is the mapping between syntactic structure and compositional semantics.



On this schema, prenominal adjectives are always in a direct modification configuration, and postnominal adjectives are always in an indirect modification configuration. I am deliberately leaving this in terms of these direct vs. indirect descriptors since the particular implementation of ‘directness’ in syntactic terms is going to remain up in the air - Cinque instantiates indirect modification as an adjective phrase which is both in a reduced relative

clause structure and, independently, that relative clause is further away from the noun than the direct modification adjective phrase is. Here, I am at least for the present discussion not taking a stance on whether or not a relative clause is present sometimes or always, a question that is discussed briefly above in Section 2.1.4. What is important is that the postnominal position is uniformly associated with that indirect syntax, whatever it ends up precisely being. There is thus a one-to-one mapping between surface linear position and covert syntactic structure.

Similarly, adjectives in a direct modification configuration always compose with the noun via Function Application, and adjectives in an indirect modification configuration always compose with the noun via Predicate Modification. There is thus also a one-to-one mapping between covert syntactic structure and the compositional semantic operation that corresponds to that structure.

The ambiguity arises in the mapping between compositional semantics and interpretation, and even then only in half of the cases. Function Application is simple and restrictive: any time the adjective composes with the noun via Function Application, the result is a nonintersective interpretation (the ‘true privative’ reading for privatives, or the subsective reading for subsectives). The complexity is introduced by Predicate Modification: while indirect modification adjectives always compose with the noun via Predicate Modification, that operation is sufficiently flexible as to potentially result in both nonintersective and intersective interpretations, depending on the context.

How does this flexibility arise? To answer this, we need to come back to the question of adjectival polysemy. Recall that so-called ‘Blame the Adjective’ accounts have historically argued that there are two differently-typed forms of a given ambiguous adjective: one of a sufficiently high type, like $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ (abstracting away from intensionality), to take the noun as its argument in Function Application, and one of a lower type parallel to a common noun, like $\langle e, t \rangle$, appropriate to combine via Predicate Modification. Depending on the account, one of these may be basic and the other achieved through type-shifting,

or they may both be basic alternates listed alongside one another in the lexicon. One of the attractive components of a ‘Blame Structure’ account like Cinque’s and the one argued for in this dissertation, however, is the hypothetical ability to dispense with such polysemy either in lexical or type-shifting terms, and attribute all variation in meaning to (ideally independently motivated) variation in structure.

However, achieving this parsimonious goal is not as simple as just identifying differences in the structural relationship between the adjective and the noun corresponding to different interpretations. This alone does not allow us to dispense with type alternations, because if the adjective is intended to combine in a different way with its nominal argument according to its structural position, as in (62), then naturally it needs to have the appropriate type for those ways of composition when it is in different positions. And so we end up needing to posit a higher-type adjective for direct modification configurations, to allow Function Application, and a lower-type adjective for indirect modification configurations, to allow Predicate Modification, and in effect we have complicated the syntax without reducing complication in the semantics whatsoever. Not an ideal outcome. Nor does it achieve the asymmetrical flexibility of indirect modification that we set out to capture with (62) in the first place!

Instead, I propose that adjectives (of the relevant kind here) are uniformly of the higher type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ - again, abstracting away from intensionality for present purposes - and therefore always take a nominal argument first. In the case of Function Application composition in the direct modification position, this is straightforward - the noun is the adjective’s first semantic argument, outputting a noun phrase of type $\langle e, t \rangle$. In the case of indirect modification positions, which I have claimed are uniformly Predicate Modification, then, the situation is only slightly more complicated. What I mean is that the adjective always combines with the noun that is its surface argument via Predicate Modification. However, to do so, it must first combine with a covert nominal argument via Function Application, in order to output the correct $\langle e, t \rangle$ type to compose via Predicate Modification with the $\langle e, t \rangle$ noun. This covert nominal performs the exact same function, semantically, as the overt

noun in direct modification positions: it supplies the content or scale relative to which the adjective is interpreted. In the cases where there is then also a second noun, the overt one, the resulting adjective meaning with its implicit scale combines via Predicate Modification and therefore set intersection, giving the intersective interpretation.

Let's call the actual overt noun N_O , and the optionally present covert one N_C . For direct modification, the semantic composition is extremely straightforward:

$$(63) \quad \llbracket \mathbf{Adj\ Noun} \rrbracket_{\text{DIRECT}} = \llbracket \mathbf{Adj} \rrbracket(\llbracket N_O \rrbracket)$$

This results in the subjective interpretation for a phrase like *beautiful dancer*, where the semantic content of *dancer* is what is taken as the initial argument of *beautiful* and thus the relevant scale of beauty applied. Obviously, this begs the question of how that composition works and I am not at all giving a lexical entry for either the adjective or the noun here; the actual nature of the semantic composition is the topic of Chapters 4 and 5, and for now I am just focused on distinguishing the readings in broad strokes and establishing their relationship to the syntax. So for present purposes, I won't complicate the picture by saying any more than the intuitive description that the noun itself is taken to be the modificand of *beautiful*, in contrast to the intersective interpretation.

For indirect modification, the adjective first composes with the covert nominal via Function Application and then with the overt noun via Predicate Modification:

$$(64) \quad \llbracket \mathbf{Adj\ Noun} \rrbracket_{\text{INDIRECT}} = \llbracket \mathbf{Adj} \rrbracket(\llbracket N_C \rrbracket) \cap \llbracket N_O \rrbracket$$

As a result, the direct semantic modificand of the adjective is some other nominal element, which will provide a different scale of evaluation to form the $\langle e, t \rangle$ function that then intersects with the overt noun. In the case of the default intersective interpretation of *beautiful dancer*, that covert N_C is probably something like *person*, something quite semantically bleached and

minimally informative, which gives us *beautiful person* \cap *dancer*.⁴ This nominal element⁵ is taken from the context, some pragmatically salient set, and I assume here that some extremely minimally-contentful set like *person* might just always be available, or at least available in many more contexts than anything more specific. By appealing to a contextually-provided argument to saturate an always-present initial argument for the adjective, this approach allows us to preserve a single underlying semantics for adjectives while deriving both readings.

Now we are finally in a position to explain the ambiguity. Consider first just the case of subsecutive adjectives like *beautiful*, as we have been. Deriving the nonintersective/subsecutive reading through (63) is clear, and since on the schema in (62) this Function Application composition is the only method available in direct modification configurations like Italian prenominal position, it follows that only the nonintersective reading obtains in these positions. On the other side, deriving the intersective reading, which is available in indirect modification configurations like Italian prenominal position, via (64) is also simple. The wrinkle is the fact that the postnominal position is ambiguous between intersective and nonintersective readings. As discussed before, this ambiguity led Cinque to posit an underlying syntactic ambiguity in this position, to maintain a one-to-one mapping between syntax and interpretation. Here, I have argued for a simpler surface-to-syntax mapping and placed the ambiguity at the mapping between semantic composition and interpretation, with the claim that Predicate Modification - as in (64) - allows both intersective and nonintersective interpretations. How?

⁴In reality, the covert nominal likely can't correspond to the meaning of the actual word *person*, because *beautiful person* has additional meanings which aren't possible options for intersecting with *beautiful dancer*. A sentence like *Olga is a beautiful dancer* never means a dancer who has a beautiful soul, for example. I'll return to this issue, and more generally the compositional derivation of specific meanings, in Chapter 5, so for now we can put up with the imprecision of some *person*-like category.

⁵Whether we want to call it a 'noun' is also controversial at this point, assuming that 'noun' is a syntactic concept that picks out particular selectional features, perhaps provided by some nominalizing, category-determining head. For now, I'll use the term much more loosely, since whatever this element is certainly has basic noun-like semantics in denoting a set of individuals unified by some properties, but we'll return to the question of how syntactically realized this 'noun' is both in Chapter 3 and later on in Chapter 5, from different angles.

The solution is simple enough that it will probably come as no surprise: given that the initial covert nominal argument N_C in (64) is just some set provided by the context, simply set $N_C = N_O$. No principle of composition prohibits this kind of duplication, even if it is arguably sub-optimal from an economy perspective (which I will discuss shortly). Certainly, the overt noun is going to be sufficiently salient in the discourse to be an easy pull from the context, from a pragmatic point of view. In the default case, this is going to result in an output denotation that is identical to the direct modification, Function Application option, which is exactly what we want:

$$(65) \quad A(N_O) \cap N_O = A(N_O)$$

Let's work through the process of deriving the ambiguity using the *beautiful dancer* example. In postnominal position, as established, *beautiful* can be read intersectively or non-intersectively. By the one-to-one syntax-semantics mapping hypothesis illustrated in (62), the only composition available in that position is indirect, as in (64), which always involves eventually combining the adjective *beautiful* with the overt noun *dancer* via set intersection (Predicate Modification), after giving it another noun from the context as an argument to Function Application. The intersective reading comes from selecting a generic, default option like *person* for that contextual argument:

$$(66) \quad \begin{aligned} \llbracket \mathbf{beautiful\ dancer} \rrbracket_I &= \llbracket \mathbf{beautiful} \rrbracket (\llbracket \mathbf{person} \rrbracket) \cap \llbracket \mathbf{dancer} \rrbracket \\ &= [\lambda P. \lambda x. \text{beautiful } P(x)] (\lambda x. \text{person}(x)) \cap \lambda x. \text{dancer}(x) \\ &= \lambda x. \text{beautiful person}(x) \cap \lambda x. \text{dancer}(x) \\ &= \lambda x. \text{beautiful person}(x) \wedge \text{dancer}(x) \end{aligned}$$

Of course, the denotations here are radically simplified - again, at the moment, the point is to illustrate the differences between the readings in broad strokes. Doing so, the result is an $\langle e, t \rangle$ function that takes an individual and evaluates true if that individual is both a beautiful person and a dancer, with no interaction between the beauty and the dancing.

As stated above, the other half of the postnominal ambiguity, the nonintersective reading, can be derived from the same Predicate Modification compositional structure by simply

selecting the same noun as the contextual argument:

$$\begin{aligned}
(67) \quad \llbracket \mathbf{beautiful\ dancer} \rrbracket_{NI} &= \llbracket \mathbf{beautiful} \rrbracket (\llbracket \mathbf{dancer} \rrbracket) \cap \llbracket \mathbf{dancer} \rrbracket \\
&= [\lambda P. \lambda x. \mathbf{beautiful}\ P(x)](\lambda x. \mathbf{dancer}(x)) \cap \lambda x. \mathbf{dancer}(x) \\
&= \lambda x. \mathbf{beautiful\ dancer}(x) \cap \lambda x. \mathbf{dancer}(x) \\
&= \lambda x. \mathbf{beautiful\ dancer}(x) \wedge \mathbf{dancer}(x) \\
&= \lambda x. \mathbf{beautiful\ dancer}(x)
\end{aligned}$$

Because all *beautiful dancers* are, generally, going to be *dancers*, the actual step of set intersection via Predicate Modification ends up vacuous, and the second conjunct of the denotation can be harmlessly eliminated. By setting $N_C = N_O$, then, the Predicate Modification composition reduces to exactly the Function Application composition, deriving the ambiguity in postnominal position from a single underlying semantics.

Crucially, this flexibility is not available in the opposite direction. Because the Function Application step does not involve any open variable slot to be filled by pragmatics, the direct modification configuration will always result in the same nonintersective reading:

$$\begin{aligned}
(68) \quad \llbracket \mathbf{beautiful\ dancer} \rrbracket_{DM/NI} &= \llbracket \mathbf{beautiful} \rrbracket (\llbracket \mathbf{dancer} \rrbracket) \\
&= [\lambda P. \lambda x. \mathbf{beautiful}\ P(x)](\lambda x. \mathbf{dancer}(x)) \\
&= \lambda x. \mathbf{beautiful\ dancer}(x)
\end{aligned}$$

So the prenominal position, which is (on both proposals) uniformly direct modification, will remain unambiguously nonintersective.

If both proposals (55) and (62) derive the unambiguous prenominal position and the ambiguous postnominal position, what is the advantage of the semantically-based ambiguity of (62)? One point might simply be that it requires less covert syntactic complexity. But the much more forceful empirical argument is the privative pattern: recall that privative adjectives, unlike subsecutive adjectives, are unambiguous in the postnominal position as well as the prenominal, only intersective in the former. On the present proposal, this pattern actually falls out naturally from the compositional semantics without any additional stipulations on syntactic configuration or movement. The critical observation is that the last step of the composition in (67) - the reduction of the conjuncts that allows the indirect option to look exactly the same as the direct option - was only possible because *beautiful dancers*

are a proper subset of *dancers*. If this were not the case, this conjunction reduction would not be valid, since there may be some individuals who meet the *beautiful dancer* half but not the *dancer* half.⁶ So, in order for this $N_C = N_O$ trick to function and allow us to derive nonintersective meanings from intersection, a subset condition must be met.

$$(69) \quad A(N_O) \subseteq N_O \rightarrow A(N_O) \cap N_O = A(N_O)$$

Privative adjectives, by definition, do not satisfy the antecedent condition of (69). What happens if we try the same backdoor method to force nonintersectivity from intersection, with a privative adjective?

$$(70) \quad \begin{aligned} \llbracket \mathbf{fake\ gun} \rrbracket_{\text{IM/NI}} &= \llbracket \mathbf{fake} \rrbracket (\llbracket \mathbf{gun} \rrbracket) \cap \llbracket \mathbf{gun} \rrbracket \\ &= [\lambda P. \lambda x. \text{fake } P(x)] (\lambda x. \text{gun}(x)) \cap \lambda x. \text{gun}(x) \\ &= \lambda x. \text{fake gun}(x) \cap \lambda x. \text{gun}(x) \\ &= \lambda x. \text{fake gun}(x) \wedge \text{gun}(x) \\ &= \emptyset \end{aligned}$$

When dealing with privative adjectives, as defined in this section, attempting to intersect $A(N_O) \cap N_O$ will always result in an empty function with no range. Any number of basic principles of interpretation, like Partee's (2009) Non-Vacuity Principle, will object to this outcome and, one way or another, cause the derivation to crash. So nonintersective readings are not available for privative adjectives in indirect modification configurations, like the Italian postnominal position.

Recall that privative adjectives also only allow an intersective reading in predicate position, as in (59), and are simply ungrammatical in predicate position when that adjective-noun combination does not furnish an intersective reading at all, as in (60). By contrast, Italian subjective adjectives allow both readings in predicate position, making the predicate position fully parallel with postnominal position with respect to (non)intersective readings. This same pattern falls out from the above analysis: a predicate position adjective is, effectively, in an indirect modification configuration and so only has access to the set intersection

⁶Of course, this only works at all if we are keeping our interpretation of *dancer* constant - someone might be a *beautiful dancer* in the current activity sense but not a *dancer* in the generic habitual sense.

method of composition. In a sense, then, on this approach all of these predicative adjectives are covertly adnominal.

2.2.3 Contextually possible arguments

An additional piece of evidence for this proposed structure - where intersective readings always involve an additional, contextually-provided initial argument for the adjective, and this is only possible in postnominal position in Italian - comes from the availability of non-standard intersective interpretations. While an example like *beautiful* is highly restricted in the possible readings that it can have, something like *good* offers quite a few more possibilities depending on the context. Consider:

(71) Steven is a good violinist.

The most prominent reading is the nonintersective one, on which Steven plays the violin well. What we might call the ‘default’ intersective reading, analogous to the physical attractiveness reading of *beautiful*, is the ‘morally good person and also a violinist’ reading. The more flexible meaning of *good*, though, supports other intersective readings given sufficient contextual support. For example, imagine that we are discussing the chess club at a music school, and the relative skill (at chess) of different instrumentalists. Then (71) could naturally mean ‘good at chess, as far as violinists go’. Deriving this reading is straightforward under the current picture of intersectivity.

Crucially, this alternate contextually-supported intersective reading is only available in postnominal position in Italian, just like the more default intersective readings:

- (72) a. Abbiamo dei violinisti bravi e dei violinisti meno bravi (can be good or no good also as chess players)
We get some good violinists and some bad violinists
b. #Qui abbiamo solo dei bravi violinisti (cannot be good as chess players, but only as violinists)
Here we get only good violinists (Cinque 2014:25)

Cinque suggests from this example ‘that these adjectives have an implicit “as a N” adjunct

which in their direct modification usage cannot be anaphoric to the context but only to the N which they modify' (Cinque 2014:25), which descriptively is much like the analysis proposed here, except that it would be inaccurate to describe the direct modification usage as 'anaphoric' to the noun - it simply takes the noun as its initial argument.

Importantly, data like this motivates unifying the 'default' intersective reading with these other contextually-supported intersective readings, both requiring this implicit 'as a N' relativizing component, rather than treating the 'default' intersective reading as some lexically-listed absolute alternate meaning for the adjective.

2.2.4 Economical asymmetries

All of the above argumentation has been couched entirely in context of the Italian data. Attempting to copy it over as is to English, however, runs into a few problems.

The first of these is that the basic observation of privative disambiguation does not hold in the same way: in English, privatives are still fully ambiguous in the prenominal position. Given the schema proposed in (62), and the unavailability of the nonintersection-via-intersection mechanism for privatives, the only way to allow ambiguity for privatives is to genuinely have both direct and indirect modification in that position. On this proposed account, therefore, we are forced to conclude that there is a genuine covert syntactic ambiguity in the English prenominal position, breaking the mirror-image symmetry with Italian. This is the position that I will adopt in light of the privatives data: English and Italian do fundamentally differ in the nature of the syntactic options for adnominal modification, in that only English displays a surface ambiguity that actually corresponds to an ambiguity in underlying syntactic structure. Luckily, this move will end up having some independent motivation and be derivable from more basic typological differences between the languages, though I will not get into that discussion until later on, in Chapter 4. In this section, I want to focus on another problem for the present analysis that is more direct and not resolvable by positing this kind of syntactic difference: indirect modification, and therefore Predicate

Modification, does not appear to universally allow nonintersection in English for subjectives in the same way as I have just argued it does in Italian.

In Italian, the unambiguous position is restricted to nonintersective interpretations, which is unsurprising on the proposed account because Function Application cannot reproduce Predicate Modification, but Predicate Modification can reproduce Function Application. However, in English, the unambiguous (postnominal) position is supposedly restricted to intersective interpretations. Given the argument that set intersection (and thus PM) can freely invoke $N_C = N_C$, it is unexpected that any position should only allow intersective interpretations. PM is, on this view, the more powerful operation, since it involves both subsection and intersection. So even without considering the behavior of privatives in English, why would postnominal subjectives not have access to nonintersective interpretations? At least for some cases, this judgment does appear to be accurate:

- (73) She is a dancer more beautiful than her instructor.
- a. ‘She is more physically beautiful than her instructor’ (intersective)
 - b. *‘She dances more beautifully than her instructor does’ (nonintersective)

In theory, nothing should block the use of dancer as the contextually-provided set for intersection and thus the reading in (73b), even if only PM is available postnominally. While it is not out of the question to assume that English and Italian differ on whether there is covert syntactic ambiguity for a given position, it is extremely difficult to propose that the conditions on how Predicate Modification and the saturation of contextual variables work differs between the languages.

One possible line of reasoning to address this problem concerns economy. It is not at all unreasonable to assume that a language would disprefer this roundabout method of nonintersection via intersection, and defining some pragmatic principle to penalize the use of a null argument when there is an overt semantic equivalent in the same sentence seems feasible. Alternately, we could imagine an economy condition on semantic operations that penalizes a derivation when another derivation with the same output truth conditions with fewer

compositional ‘moves’ exists, and this would block the proposed artificial nonintersection process.

However, if we introduce these kinds of economy considerations, the question then becomes why the process is ever available at all in Italian. Let’s say that we do not want to posit them as language-specific considerations (or at least, that doing so would require substantial evidence that the economy principles are also active in other domains for one but not the other language, which we presently lack). How do we get the contrast between English and Italian with language-agnostic economy principles?

One option appeals to some notion of ‘canonicity’.⁷ In Italian, it is the postnominal position which is the standard, default, canonical adjectival position, and also the postnominal position which has the basic semantics of intersection. If we do the ‘backdoor’ trick in this position, it may violate our economy condition, which would prefer that we derive nonintersectivity by putting the adjective prenominal. But other conditions in Italian – either syntactic economy, if prenominal adjectives are derived via movement, or something about preferred linear orders – mitigate against this, and appear to win.

In English, however, the basically intersective postnominal position is noncanonical. Perhaps when we try to ‘backdoor’ nonintersectivity in English postnominals, the economy conditions still push for a prenominal strategy instead, but here the syntactic/word order conditions are aligned with the semantic economy and also would prefer prenominal nonintersectivity. So the difference between the languages is that in Italian, semantic economy competes with syntactic/order principles and the stalemate (or victory of the latter) allows for semantically uneconomical operations to proceed in postnominal position; while in English, semantic economy aligns with syntactic/order principles and both penalize (and therefore, together, exclude) the option to derive nonintersective meaning in the intersective position. The consequence of this would be that Italian postnominals allow the ambiguity

⁷Ideally, of course, this strategy would not require a basic notion of ‘canonical’ as an actual primitive of linguistic structure. I assume here that calling a position canonical is a convenient abbreviation for something about its syntax being more basic in that language - requiring fewer movement steps, being the base generated position, etc.

(at least for the standard subsective case) but English postnominals do not, as we appear to see in the data.

One piece of data that points in favor of this kind of account comes from modified postnominal adjectives. Recall that English postnominal adjectives are supposed to show exclusively stage-level readings:

- (74) The stars visible are Aldebaran and Sirius.
- a. #‘The stars that are generally visible include Aldebaran and Sirius’ (individual-level)
 - b. ‘The stars that happen to be visible now include Aldebaran and Sirius’ (stage-level)

But it appears that if you modify the adjective, specifying it with something like a prepositional phrase, the ambiguity reappears:

- (75) The stars visible from Earth are Aldebaran and Sirius.
- a. ‘The stars that are generally visible from Earth include Aldebaran and Sirius’ (individual-level)
 - b. ‘The stars that happen to be visible from Earth now include Aldebaran and Sirius’ (stage-level)

On this economy story, the reason why the individual-level (which I’m saying is analogous to nonintersective, as will be argued more extensively later in Chapter 4) reading is unavailable postnominally, even though in theory it could be produced via Predicate Modification with our backdoor nonintersectivity maneuver, is that an economy principle punishes using a more complex compositional operation to get the same reading that could be derived through simple Function Application in prenominal position. However, this ‘heavy’ modified version of the adjective is not available at all prenominally:

- (76) *The visible from Earth stars are Aldebaran and Sirius.

Potentially, then, this ungrammaticality of the prenominal position blocks the economy principle from ruling out the individual-level reading: because the compositionally simpler prenominal position is unavailable, the English grammar tolerates the nonintersective use of Predicate Modification for the postnominal position in the same way that it does for Italian.

Again, we may not need to resort to any of these economy-level considerations to derive the correct readings here, if our theory of the semantic differences between the languages that will be developed in Chapter 4 accomplishes its goals. Since I want to maintain that this chapter’s arguments for the revisions to Cinque’s picture of syntactic ambiguity stand on their own, independent of the particular semantic theory proposed later, the thoughts in this section just serve to demonstrate that we might not need anything but these syntactic considerations to do so.

2.3 Movement and ambiguity in Bangla

Privative disambiguation as an empirical strategy should have reach beyond Italian if it is going to be the focus of the dissertation, and if we want to explain the resulting patterns by appeal to fundamental logical properties of composition with privative adjectives which should not be language-specific. Luckily, it appears to. In this section, I will introduce the phenomenon of adjectival movement in Bangla and the basic facts of how it affects the realization of the nonintersective ambiguity, and apply the same privative disambiguation strategy to reveal that the pattern is strikingly like that in Italian.

As with the Italian data, I will refrain from giving an actual semantic characterization of the pattern until Chapter 4; the goal here will again be to describe the syntactic distribution of possible ambiguous meanings. Additional data motivating a particular semantic theory of Bangla will be introduced in Section 4.5.

2.3.1 Focus movement in Bangla

The canonical order of the nominal phrase in Bangla is given in (77), exemplified in (78):

(77) Demonstrative > Numeral > Classifier > Adjective > Noun (Bhattacharya, 1999)

(78) ei du ʈo lal boi (Syed, 2015)
 DEM two CL red book
 ‘these two red books’

However, this order is somewhat flexible. In particular, the adjective and noun may occur before the Numeral + Classifier sequence⁸, as in (79).

- (79) *ei lal boi du ʈo* (Syed, 2015)
 DEM red book two CL
 ‘these two red books’

Absent the demonstrative, these movements effect whether the noun phrase is read as definite or indefinite, with the canonical position (80) being indefinite and the moved position (81) definite (Bhattacharya, 1999; Dayal, 2012). However, this contrast will be orthogonal to the present discussion, and we will be more interested in examples with an explicit demonstrative anyway, so I will leave this fact to the side.

- (80) *du ʈo lal boi* (Syed 2015)
 two CL red book
 ‘two red books’

- (81) *lal boi du ʈo* (Syed 2015)
 red book two CL
 ‘the two red books’

What is of present interest is the fact that the adjective and noun are not required to move together. In particular, the adjective can precede the Numeral + Classifier sequence or the demonstrative on its own with the noun staying in its phrase-final base position.

- (82) *ei lal du ʈo boi* (Syed 2015)
 DEM red two CL book
 ‘these two red books’

- (83) *LAL ei du ʈo boi* (Syed 2015)
 red DEM two CL book
 ‘these two red books’

Syed (2014, 2015) reports that the pre-demonstrative position is not possible for the adjective without phonetic stress. I have found this judgment to be less consistent among Bangla

⁸In general, I will be using the numeral *du* ‘two’ or other larger numerals, since *ek* ‘one’ has been effectively grammaticalized as an indefinite determiner and therefore does not allow the occurrence of both an overt determiner and a distinct numeral within the same phrase, which is necessary to delineate the precise location of the adjective for the present arguments.

speakers: while the moved adjective in a sentence like (83) always bears the information-structural properties of focus, emphasizing the redness of the books as opposed to alternative colors, the phonetically unstressed equivalent is not always judged ungrammatical as Syed predicts. Regardless, this should be unimportant for Syed’s conclusion that the pre-demonstrative position is the specifier of a focus projection FocP: whether being in such a position (and bearing the focus feature that motivates the movement) strictly requires phonetic stress, or simply marks the adjective to be interpreted as focused, is a question of the PF-mapping algorithm more than a question for this narrow analysis. More importantly, we will see shortly that even the focus semantics is not relevant for the present argument.

What is important here is that, as a result, Bangla allows four distinct positions for the adjective and noun, relative to each other, in the full DP. They can both remain in base-generated position below the Numeral + Classifier complex; can both move past the Numeral + Classifier complex but not past the demonstrative; the adjective alone can move to that intermediate position; or the adjective alone may move all the way out past the demonstrative.

- (84) Bangla DP order options:
- a. DEM NUM-CL [ADJ NOUN]
 - b. DEM [ADJ NOUN] NUM-CL
 - c. DEM [ADJ] NUM-CL [NOUN]
 - d. [ADJ] DEM NUM-CL [NOUN]

I have been referring to the higher positions as the result of movement; Syed (2015) argues for such an account rather than different base-generation positions by pointing out that other intervening material in the specifier position of the Numeral phrase blocks the adjective and/or noun from appearing higher. Specifically, following the suggestion from Danon (2012) that numerals can alternate between being heads or specifiers of their projection even within a single language, Syed argues that Bangla numerals are split this way: *four* and below as heads, *five* and higher as specifiers of NumP, and shows that neither adjective nor noun can appear above the Numeral + Classifier complex when the numeral is in the larger category.

- (85) a. ami lal boi char Te kinechi
 I red book four Cl buy-pr perf-1p
 ‘I have bought the four red books’
 b. ?ami lal boi panc Ta kinechi
 I red book five Cl buy-pr perf-1p
 ‘I have bought the five red books’
 c. *ami lal boi choy Ta kinechi
 I red book six Cl buy-pr perf-1p
 Intended: ‘I have bought the six red books’

If the adjective or noun could be simply base-generated in the higher position, there is no reason it would be sensitive to the type of numeral lower down. But even being in the pre-demonstrative position is blocked by numeral six or higher, suggesting that it cannot be competition for the specifier of NumP position, and favoring a movement analysis.

Actually, nothing in the present argument hinges on the decision between movement and multiple base-generated positions. I will continue referring to this phenomenon as movement, but what matters here is simply that the adjective ends up in these distinct positions, because we can then investigate their effect on possible interpretations. I will move to that shortly, but first, a brief note on why Bangla in particular is an illuminating test case.

2.3.2 A note on phase structure

As discussed previously in this chapter, Cinque offers a structural account of the nonintersective ambiguity that ties intersective meanings to adjectives that are in reduced relative clauses and a certain distance from the noun. However, no syntax-semantics mapping is offered to explain why that position causes intersective meanings, and why the direct modification, non-clausal, local position causes nonintersective meanings. One attempt at such an explanation is offered by Reichard (2013) who ties the clausal nature of indirect modifiers to intersectivity by means of phases.

The grammatical notion of phases was introduced to linguistic thought in Chomsky (2001)), who identified that a computationally efficient language system, subject to external constraints such as memory limitations, would require some kind of ‘chunking’ algorithm: a

mechanism for packaging together certain complex structures and treating them as a single, simplified object so that the internal structure need no longer be remembered. A phase, then, is a particular structural boundary in the course of syntactic derivation at which, upon its completion, seals off its internal content from further manipulation or access by the syntax, in a sense flattening the structure of its complement. Various different syntactic structures may qualify as phases: Chomsky originally takes CP and vP to be phases, and so upon the Merge of the phase head (either C or v), the complement of that head becomes ‘impenetrable’ and invisible to further derivational operations. In Minimalist terms, the complement of a completed phase is packaged and transferred to both the conceptual-intentional and sensorimotor interfaces, where it is interpreted and spelled out phonetically, respectively, and so no further manipulation of its syntactic content can affect either the meaning or pronunciation of the phrase (and therefore, we can say, no further manipulation of its syntactic content can occur).

What Reichard argued is that, since indirect modifiers are CPs and CPs are phases, we may be able to attribute the intersective reading of indirect modifiers to the fact that they are contained in phases. By the time that the adjective modifies the noun, the CP containing it has already been completed and so that phase would be closed off and sent to the conceptual interface to be interpreted. Since intersective readings are, intuitively, the result of not interpreting the adjective relative to the noun, arguably that could be due to the phasal nature of the CP preventing any further adjustments to the adjective’s interpretation that would be required for the relativized nonintersective interpretation: ‘phases are, *inter alia*, taken to be units of semantic valuation. If indirect modifiers are clausal and thus CPs, they will, thus, be such units of valuation. In this case, the semantic value of an indirect modifier is already fixed when it modifies its host. The host can, thus, not relativize the modifier in any way. And if the modifier is not interpreted relative to its modifiee then intersectivity is the only possible interpretation left.’ (Reichard 2013:13)

The appeal of this kind of explanation is obvious. Reducing the nonintersective ambiguity

to a preexisting notion that is independently motivated by third factors like computational efficiency, while precisely characterizing its syntactic distribution, would be an elegant Minimalist analysis. However, the immediate concern with jumping to the phase conclusion is that there are multiple potential boundaries between an indirect modifier adjective and its modified noun. On Cinque's account, there is always a CP between the adjective and noun, but there are also a number of other projections, like NP and the posited dP, potentially also vP if the clause-internal adjective is in a predicative position. Any one of these barriers could be the critical one, since they are always all present, at least in Italian. Furthermore, even if we had evidence to zero in on CP specifically as the causal factor, that would not constitute evidence that it is the phasal nature of CP that matters. That explanation would be attractive, but parsimony is only one potential consideration. Some other feature of CP specifically might be in play. What we need is a diagnostic to distinguish the effect of a phase boundary from other potential boundaries, like clauses or definites.

Bangla potentially provides one. After Chomsky (2001) identified CP and vP as phases - with vP as the internal counterpart of CP, splitting the clausal spine into two phases - a body of subsequent work has argued that nominal phrases might show a similar phasal structure. A popular argument has been that DP also constitutes a phase (Svenonius, 2004; Želko Bošković, 2012), and Simpson and Syed (2016); Syed and Simpson (2017) have further argued that, just like how clauses have the biphasal structure of CP and vP, nominal phrases might also be biphasal with the external DP phase and an internal QP phase, which in Bangla hosts the numeral-classifier complex. I won't review Syed & Simpson's arguments for the phasal status of QP in Bangla, but will assume their conclusion here.

As a consequence, Bangla offers an ideal empirical testing ground for determining which specific kind of barrier might be the critical locality condition separating the intersective from nonintersective readings. As far as I know, there is no evidence that Bangla adjectives are clausal or need to be clauses to undergo focus movement as discussed in this section, so if it is specifically clausal structure that triggers indirect modification, we should not

see syntactically determined intersective interpretations. If it is definiteness in particular that matters, as Cinque posits, then we might expect the DP boundary to be the relevant locality factor, such that adjectives outside of DP are the indirect modifiers. DP would also be a phase, of course, but given the phasal status of QP inside DP, movement past QP but not past DP should be sufficient for indirect modification if the relevant notion of indirect is ‘separated by a phase’. The fact that Bangla offers a nominal phrase with two distinct phases, one independent of definiteness, is what allows it to be a potentially useful test for what notion of locality we should care about.

There are two caveats to this approach. First, it’s not clear what Reichard predicts when a phase *boundary* separates the adjective from the noun but the adjective itself is not contained *within* the phase. His (limited) exposition of the account describes the problem as the inability to adjust the meaning of the adjective after it has already been sealed off within a phase, and so taking that narrowly it might not matter if the noun is within a phase even if the phase boundary separates them. Interpreting his argument a little more widely, however, we could also predict that the process by which the adjective’s meaning becomes relativized to the noun requires some access to the semantic content of the noun in a way that would be blocked by the ‘impenetrable’ nature of the phase boundary. It’s entirely possible that mirrored constraints are at work in Italian and Bangla - in the former, the adjective is rendered semantically immutable before it encounters the noun, and in the latter, the noun semantically immutable before it encounters the adjective.

It’s also possible, however, that even if QP is a relevant barrier distinguishing direct from indirect modification in Bangla, it is not because of its phasal status. It’s possible that the act of movement away from the base-generated position at all blocks the process of direct modification somehow, and it’s also possible that - much like some inherent property of CP besides its phase status could determine nonlocality in Italian - some inherent property of QP specifically besides its phase status could be active in Bangla. Of course, if it’s shown reliably that Italian CPs and Bangla QPs are performing the same interpretive function,

then it would be reasonable to conclude that the property they share of being a phase is the relevant one, but that is not necessarily implied just by the QP boundary affecting interpretations.

This is all to say that Reichard’s proposal deserves serious consideration, and one motivation for investigating the nonintersective ambiguity in Bangla in particular is the perspective it might give us on the impact of phase boundaries on the ambiguity. That does not, however, allow us to necessarily conclude that the phase-based explanation would be the most cross-linguistically attractive analysis just from this particular test in Bangla, even if it scores quite highly on a Minimalist conception of parsimony. Reichard’s phasal proposal, as well, is seriously underspecified, and we would need to develop a much richer account of what kind of ‘relativizing’ occurs to give an adjective a nonintersective meaning, and how that is blocked by phase completion, so even though appeal to phases is attractive from a Minimalist perspective, the work required to account for the ambiguity in terms of phases would not be itself minimal. The pattern - as we will see - has a number of potential explanations.

2.3.3 Effects of movement on ambiguity

For Bangla adjectives which display the nonintersective ambiguity, this kind of movement has consequences for which interpretation surfaces in different positions. As we have established, there are basically four relevant possible positions for the adjective to appear in: below NUM+CL with N, above NUM+CL with N, above NUM+CL alone, and above DEM alone. In the first, base-generated position, ambiguous adjectives display a weak preference for their nonintersective, subsective interpretation, also allowing the intersective interpretation with the right contextual support.

- (86) oi du to b^halo ukil
 DEM two CL good lawyer
 ‘these two good lawyers’
- a. ‘these two legally skilled lawyers’
 - b. ?‘these two morally good lawyers’

- (87) oi du ʈo sundor dancer
 DEM two CL beautiful dancer
 ‘these two beautiful dancers’
- a. ‘these two dancers who dance beautifully’
 b. ?‘these two physically beautiful dancers’

- (88) oi du ʈo choto hati
 DEM two CL small elephant
 ‘these two small elephants’
- a. ‘these two elephants who are small for their species’
 b. ?‘these two elephants who are small for any animal’

As in every language, individual adjectives show subtleties in lexical meaning that constrain their range of application. For example, while *sundor* ‘beautiful’ does display the classical nonintersective ambiguity with nouns like *dancer* ‘dancer’ (an English loan) as in (87), describing as beautiful either the dancers themselves or their dancing, it appears to be restricted to describing beauty that can be visually perceived. Modifying a noun that describes a non-visual event, like *gayok* ‘singer’, allows only the intersective reading:

- (89) oi du ʈo sundor gayok
 DEM two CL beautiful singer
 ‘these two beautiful singers’
- a. *‘these two singers who sing beautifully’
 b. ‘these two physically beautiful singers’

This pattern in (89) persists regardless of how the adjective is moved. We might consider Bangla *sundor*, then, to fall somewhere on a scale of lexical flexibility between English *beautiful*, which robustly displays both meanings with a wide variety of nouns, and Russian *krasivyyj* ‘beautiful’, which allows only the intersective, physical beauty meaning (Martin and Bikina, 2022). This isn’t a problem with deriving ambiguity from the noun *gayok*, to be clear - it’s perfectly able to receive both meanings of *b^halo* ‘good’, for example, as in (90) below - but only with the particular interaction in *sundor gayok*.

- (90) oi du ʈo b^halo gayok
 DEM two CL good singer
 ‘these two good singers’

- a. ‘these two skilled singers’
- b. ?‘these two morally good singers’

For those adjectives that do show an ambiguity, movement of the adjective and noun together does not affect the possible interpretations.

- (91) oi [b^halo ukil] du ʈo
 DEM good lawyer two CL
 ‘these two good lawyers’
- a. ‘these two legally skilled lawyers’
 - b. ?‘these two morally good lawyers’

- (92) oi [sundor dancier] du ʈo
 DEM beautiful dancer two CL
 ‘these two beautiful dancers’
- a. ‘these two dancers who dance beautifully’
 - b. ?‘these two physically beautiful dancers’

- (93) oi [choto hati] du ʈo
 DEM small elephant two CL
 ‘these two small elephants’
- a. ‘these two elephants who are small for their species’
 - b. ?‘these two elephants who are small for any animal’

Both readings are still available, with a weak preference for the subjective reading. This indicates that there is no intrinsic effect of movement on the adjective’s interpretation: so long as it is moved alongside the noun as a constituent, interpretation is unaffected. As that implies, however, moving the adjective alone and leaving the noun in base-generated position does have an effect:

- (94) oi [b^halo] du ʈo ukil
 DEM good two CL lawyer
 ‘these two good lawyers’
- a. ?‘these two legally skilled lawyers’
 - b. ‘these two morally good lawyers’

- (95) oi [sundor] du ʈo dancier
 DEM beautiful two CL dancer
 ‘these two beautiful dancers’

- a. ?‘these two dancers who dance beautifully’
- b. ‘these two physically beautiful dancers’

(96) oi [choto] du ʈo hati
 DEM small two CL elephant
 ‘these two small elephants’

- a. ?‘these two elephants who are small for their species’
- b. ‘these two elephants who are small for any animal’

Moving the adjective across the NUM+CL boundary has a few effects, which our grammaticality judgment notation system is a little too coarse to capture. First, it does result in a word order that is somewhat degraded on its own, and requires more specific context to support. All of the sentences with the adjective moved separately from the noun are slightly dispreferred relative to the orders where the adjective and noun remain a constituent, but they are also all sufficiently acceptable to speak coherently of their available interpretations and differences between them. The second effect of movement across NUM+CL is that it inverts the preferences for the interpretations that is present in the base position. In (94)-(96), the intersective interpretation is preferred, and the subsective interpretation is possible though less salient. Third, the gap between the availability of the interpretations widens. The preference for the nonintersective over intersective meaning in the [adjective noun] orders is smaller than the preference for the intersective over nonintersective meaning in this order. None of these preferences are anything like absolute, however, and contextual support reliably enables each meaning in each word order; there is just a distinct difference in how strong the default preference is, or how salient the default interpretation is.

Finally, further movement of the adjective out past the demonstrative has no further effect on interpretation.

(97) [b^halo] oi du ʈo ukil
 good DEM two CL lawyer
 ‘these two good lawyers’

- a. ?‘these two legally skilled lawyers’
- b. ‘these two morally good lawyers’

- (98) [sundor] oi du ʈo dancier
 beautiful DEM two CL dancer
 ‘these two beautiful dancers’
 a. ?‘these two dancers who dance beautifully’
 b. ‘these two physically beautiful dancers’
- (99) [choto] oi du ʈo hati
 small DEM two CL elephant
 ‘these two small elephants’
 a. ?‘these two elephants who are small for their species’
 b. ‘these two elephants who are small for any animal’

In this order, the preference for the intersective over nonintersective meaning is retained.

To summarize, for those adjectives that display a nonintersective ambiguity, that ambiguity persists across different word order options, whether the adjective is moved alone or together with the noun. However, that movement can change the default ranking of preference for which of the ambiguous interpretations is more salient. Whenever the adjective is local to the noun, regardless of the position of the AN constituent, the nonintersective interpretation is slightly preferred over the intersective one. When the adjective moves past QP, the location of NUM+CL, without the noun, the intersective interpretation becomes preferred over the nonintersective one, and further movement past DP does not change this preference further. This pattern appears to be consistent across ambiguous subjective adjectives, and so we can outline the generalization so far like this:

Adjective Type:		Subjective	
Interpretation:		Nonintersective	Intersective
(100)	DEM NUM-CL [ADJ NOUN]	✓	?
	DEM [ADJ NOUN] NUM-CL	✓	?
	DEM [ADJ] NUM-CL [NOUN]	?	✓
	[ADJ] DEM NUM-CL [NOUN]	?	✓

2.3.4 The irrelevance of focus

Before moving to other adjective types, a quick note on focus. I have left any indication of phonetic stress off of the examples in this section, reflecting what I said above about the

potential irrelevance of stress to grammaticality, contra Syed (2015). What matters for the current argument is that it is not specifically focus causing this difference in interpretation, and we can establish that by identifying that the focus pattern is the same regardless of whether the noun is moved alongside the adjective or not.

- | | | | | |
|-------|----|--|---|--------|
| (101) | a. | oi B ^h ALO ukil du ʈo
DEM good lawyer two CL
'those two good lawyers' | → | NI > I |
| | b. | oi b ^h alo ukil du ʈo
DEM good lawyer two CL
'those two good lawyers' | → | NI > I |
| | c. | oi B ^h ALO du ʈo ukil
DEM good two CL lawyer
'those two good lawyers' | → | I > NI |
| | d. | oi b ^h alo du ʈo ukil
DEM good two CL lawyer
'those two good lawyers' | → | I > NI |

In (101) we have a two-by-two matrix of examples, crossing whether the noun is moved along with the adjective and whether the adjective bears focal stress. In both movement options, it is reported that phonetic stress on the adjective is more natural, though not necessarily required (phonetic stress on the moved noun is also an option, though again, less preferred relative to stressing the adjective). Critically, however, the presence or absence of overt focus has no effect on the ranking of preferred interpretation, but the movement of the noun does. In both of the configurations (101c) and (101d), with the adjective separated from the noun, the intersective interpretation is preferred regardless of the focus. By contrast, in both (101a) and (101b), with the adjective and the noun moved together, the nonintersective interpretation is preferred regardless of the focus. And recall that this is the same interpretive preference found when the adjective and noun are together unmoved, with neither overt phonetic stress nor any focus-motivated movement. Thus, it is clear that the adjective-noun locality relationship is what determines the interpretation, not the focus.

2.3.5 Privative disambiguation again

In Italian, we found that where subsecutive adjectives were ambiguous in certain syntactic positions, privative adjectives were in fact unambiguous, and this was revealing of the underlying form of composition available in those positions. So far, the behavior of subsecutive adjectives in Bangla is not unlike those in Italian: they prefer nonintersective meanings in the more noun-local syntactic position (which we can call direct modification in both cases), and intersective meanings in the non-local position (indirect modification), but with some ambiguous flexibility. Therefore, if the same kind of compositional account is to be given to Bangla as Italian, we would predict to see the same kind of disambiguating behavior for privative adjectives here as well.

First, we need to define the space of privative adjectives Bangla offers for this diagnostic. The most natural counterpart of English *fake* in Bangla is actually the loan word, *fake*, which is the default option chosen by Bangla speakers that I have worked with. There is also *nokol* ‘fake’, with slightly different meanings; I’ll cover both and illustrate the differences here, but let’s begin with the loan word as it is the primary option chosen by consultants.

In the base-generated position, Bangla *fake* only allows the nonintersective (truly privative) interpretation: for something like *fake bonduk* ‘fake gun’, the reading on which it is not a gun at all but some device made to look like one. The intersective reading, where it is truly a gun but faked in some other aspect (i.e., not really a valuable WWII antique) is possible but significantly degraded.

- (102) oi du ʈo fake bonduk
DEM two CL fake gun
‘these two fake guns’
- a. ‘these two objects which are pretending to be guns’
b. ??‘these two guns which are faked in some other aspect’

This pattern holds if both the adjective and noun are moved past QP together.

- (103) oi fake bonduk du ʈo
DEM fake gun two CL

‘these two fake guns’

- a. ‘these two objects which are pretending to be guns’
- b. ??‘these two guns which are faked in some other aspect’

If the adjective is moved on its own, either past QP or past DP, the readings flip and become stricter: now only the intersective reading is available, and the nonintersective reading is entirely gone.

(104) oi fake du tɔ bonduk
 DEM fake two CL gun
 ‘these two fake guns’

- a. *‘these two objects which are pretending to be guns’
- b. ‘these two guns which are faked in some other aspect’

(105) fake oi du tɔ bonduk
 fake DEM two CL gun
 ‘these two fake guns’

- a. *‘these two objects which are pretending to be guns’
- b. ‘these two guns which are faked in some other aspect’

Bangla *fake*, then, shows an asymmetrical and stricter version of the subjective pattern. In the direct modification position, where subjective adjectives show a slight preference for the nonintersective reading, *fake* shows a preference in the same direction but stronger. More interestingly, in the non-local indirect modification position, where subjective adjectives are ambiguous but have a slight preference for the intersective reading, *fake* is unambiguous and only allows the intersective reading.

Adjective:		<i>fake</i>	
Interpretation:		Nonintersective	Intersective
(106)	DEM NUM-CL [ADJ NOUN]	✓	??
	DEM [ADJ NOUN] NUM-CL	✓	??
	DEM [ADJ] NUM-CL [NOUN]	*	✓
	[ADJ] DEM NUM-CL [NOUN]	*	✓

This mirrors almost precisely the privative disambiguation effect in Italian. The privative adjective can be said to only allow readings that directly reflect the compositional operation(s) available in that position, and the same argument explaining why the non-local

indirect modification position does not allow the nonintersective meaning in Italian can be applied exactly to Bangla *fake*. Conversely, the same explanation can be given for why the non-local position is flexible for subsectives in both Italian and Bangla: while the composition in that location is only ever intersection via Predicate Modification, different choices of the pragmatically determined scale as the adjective's first argument can either create surface-level intersective readings or recreate the nonintersective reading via intersection. Bangla adds to the picture not only a second example of the privative disambiguation process indicating this underlying syntax-semantics mapping, but also a kind of non-locality that does not involve clausal structure for its equivalent to indirect modification, suggesting that the critical sense of locality must be something else. The fact that, just like the subsectives, it is the QP layer across which the meanings are divided, offers potential initial support for Reichard's phasal explanation, though as we covered earlier and will return to in Chapter 4, not conclusive evidence.

So the non-local position appears to function the same in both languages: only intersection for composition, which allows both readings with subsectives but only the intersective reading with privatives. The key distinction between Italian and Bangla - besides the nature of the syntactic structures - is the ambiguity in the local position. Recall that in Italian, the local direct modification position allowed only the nonintersective reading for both subsective and privative adjectives. In Bangla, we see that the local position does prefer the nonintersective reading, but does also allow the intersective reading (for both subsective and privative adjectives). How can we explain this flexibility, given the argument that privative adjectives are supposedly disambiguating and revealing the mode of composition possible in a given position? Following that logic, we simply have to conclude that both Function Application and Predicate Modification are possible compositional operations when the adjective is in base position in Bangla. In this sense, Bangla behaves more like English than Italian, since the English prenominal position is ambiguous for both subsective and privative adjectives. The important observation is that even in Bangla, the privative *fake* allows only

intersective readings when the adjective is non-local and therefore does not have access to the overt noun as its initial argument, forcing Predicate Modification. In principle, nothing should prevent the use of Predicate Modification and the choice of a distinct scale to which the adjective is relativized, even in local position. The behavior of the privative adjective is not what's at issue here: it is entirely unsurprising, if genuine nonintersective and intersective compositional processes are available in a position and that is reflected in the subjective ambiguity, that the privative adjective would also be ambiguous in that position.

The Bangla privative pattern is slightly complicated, though not in a way that undermines these conclusions, when we consider the alternate translation for 'fake', *nokol*. Where the loan word *fake* does allow both interpretations in distinct positions, *nokol* only allows the truly privative, nonintersective meaning. We see this in the base-generated position and when the adjective and noun are moved together:

- (107) oi du ʈo nokol bonduk
 DEM two CL fake gun
 'these two fake guns'
 a. 'these two objects which are pretending to be guns'
 b. *'these two guns which are faked in some other aspect'
- (108) oi nokol bonduk du ʈo
 DEM fake gun two CL
 'these two fake guns'
 a. 'these two objects which are pretending to be guns'
 b. *'these two guns which are faked in some other aspect'

However, in the other positions where the adjective and noun are non-local and *fake* only allows the intersective meaning, *nokol* is simply ungrammatical:

- (109) *oi nokol du ʈo bonduk
 DEM fake two CL gun
 Intended: 'these two fake guns'
- (110) *nokol oi du ʈo bonduk
 fake DEM two CL gun
 Intended: 'these two fake guns'

The unavailability of the intersective meaning with *nokol* leading to its ungrammaticality in these positions reinforces the conclusion that only intersective composition is available when the adjective is non-local. There are two possible explanations for why *nokol* cannot occur with this meaning: it may be a lexical alternation, much like Bangla *sundor* ‘beautiful’ is more limited than its English counterpart, or it may be due to a competition effect. The Bangla adjective lexicon also contains *jali* ‘counterfeit’, which uniformly displays the intersective reading of a forgery or counterfeit that is still an exemplar of the head noun, and it retains this meaning regardless of the position of the adjective.

- (111)
- a. oi du ʈo jali bonduk
DEM two CL counterfeit gun
‘these two counterfeit guns (= really guns, something else about them is faked)’
 - b. oi jali bonduk du ʈo
DEM counterfeit gun two CL
‘these two counterfeit guns (= really guns...)’
 - c. oi jali du ʈo bonduk
DEM counterfeit two CL gun
‘these two counterfeit guns (= really guns...)’
 - d. jali oi du ʈo bonduk
counterfeit DEM two CL gun
‘these two counterfeit guns (= really guns...)’

One possibility⁹ is that the availability of *jali* blocks the use of movement to signal the intersective meaning with *nokol*: since the movement is already always somewhat dispreferred, the fact that there is *jali* to unambiguously communicate the intersective meaning¹⁰ could be a strong enough secondary force that the movement option is no longer viable. This would also explain the even weaker availability of the intersective meaning of *nokol* in base position, relative to *fake*. The key observation would be that the native Bangla word *nokol* is in competition with *jali*, while the loaned *fake* is not.

⁹This line of argument was suggested by Ankana Saha (p.c.), who also provided some of the judgments for this section.

¹⁰While I am referring to the meaning of *jali bonduk* as ‘intersective’, this is not to say anything about the compositional nature of *jali* itself - solely that the resulting meaning is equivalent to the one that is achieved via intersection with *fake*. This example highlights the importance of distinguishing from intersectivity as a descriptor of interpretations and intersection as a descriptor of composition we can’t conclude anything here from just this data about what kind of compositional operation is occurring with *jali* in local position.

(112)

Adjective:	<i>nokol</i>		<i>jali</i>
Interpretation:	Nonintersective	Intersective	Intersective
DEM NUM-CL [ADJ NOUN]	✓	??	✓
DEM [ADJ NOUN] NUM-CL	✓	??	✓
DEM [ADJ] NUM-CL [NOUN]	*	*	✓
[ADJ] DEM NUM-CL [NOUN]	*	*	✓

Both Bangla *fake* and *nokol* suggest the same conclusion regarding the interaction of privativity and (non)intersectivity - nonintersective interpretations of privative adjectives are fully unavailable in syntactic positions that primarily support intersective interpretations, in contrast to subjective adjectives, which allow nonintersective interpretations in default-intersective positions with contextual support.

2.4 Interim summary

In this chapter, I have laid out the basic shape of a proposal for how nonintersective and intersective interpretations of ambiguous adjectives are distributed across different adnominal positions in syntactic structure. Adopting the notions of direct and indirect modification from Cinque (2010), I propose that the source of ambiguity in Italian is not covert syntactic alternations, nor lexical polysemy of the adjective, but rather arises in the course of semantic composition. These adjectives are uniformly two-place predicates of a type like $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$, and individual syntactic positions are uniformly associated with individual compositional operations in a one-to-one manner: adjectives in noun-local, direct modification positions always combine with the modified noun via Function Application (taking the noun as its first argument), while adjectives in non-local, indirect modification positions always combine with the modified noun via Predicate Modification. The syntactic and semantic complexity required of the proposal so far, therefore, is minimal.

The ambiguity arises entirely from the selection of the adjective's first argument when that argument is not the overt modified noun (call it N_O). In order for Predicate Modification with an $\langle e, t \rangle$ -type noun to be possible, the higher-type adjective must first have its

initial $\langle e, t \rangle$ -type argument saturated. This can happen pragmatically, with the selection of some contextually salient but phonologically unrepresented noun (call it N_C). This N_C combines with the adjective via Function Application, and in doing so serves as a kind of scale relative to which the adjective is interpreted, like the nonintersective reading. This is how the standard intersective reading of a phrase like *beautiful dancer* or *good thief* is derived - something like *person*, which due to its minimal semantic restriction is almost always sufficiently supported by the context, is chosen as N_C , and the resulting predicate of *beautiful person* or *good person* is intersected with *dancer* or *thief*, respectively.

But this process leaves open the possibility for ambiguity in the choice of N_C . Other contextually-salient options can be chosen to create more infrequent meanings, like the *violinist who is good at playing chess*, but one option for N_C that is always going to be contextually supported is N_O itself. The choice of N_O as N_C will create a redundant Predicate Modification step, and so duplicates the nonintersective meaning via intersection. However, this is only possible when the output of nonintersective modification is subsective; otherwise the second compositional step of PM would not be redundant but in fact return the empty set, an option ruled out by general principles of interpretation. This explains the ambiguity of subsective adjectives in indirect modification positions but the lack of ambiguity for privative adjectives, which is found in both Italian and Bangla.

This proposal has the advantages of (i) limiting redundancy in the lexicon, compared to ‘Blame the Adjective’ accounts, (ii) avoiding positing syntactic ambiguity where it would only be motivated by the surface interpretive ambiguity, compared to Cinque (2010), and (iii) deriving the ambiguity in the course of composition, and therefore allowing features of syntactic structure to affect the ambiguity, compared to the purely pragmatic, post-compositional Maienborn (2020) account, while preserving a role for pragmatic variable resolution that seems necessary given the context-sensitivity of certain intersective readings. But obviously, significant questions about the present account remain unanswered. In particular:

- *What is the syntactic status of N_C ?* I have argued here that, even in cases of intersective

modification, adjectives always enter the derivation as two-place predicates and always have their first argument supplied by the context. Unlike in a post-compositional pragmatic story like Maienborn (2020), where this is an unbound internal variable that is simply specified, I claimed that this is a compositional process: the adjective has an unsaturated lambda term that requires saturation by composing with this argument. But does this necessarily mean that there is a syntactic object denoting N_C which is actually projected in the structure? If so, what kind of evidence could support the syntactic reality of the always-null argument?

- *What is the locality condition that distinguishes direct and indirect modification?* I have characterized the compositional properties of adjectives in direct modification positions and indirect modification positions, and distinguished the two by reference to some generalized notion of locality: direct modification requires the adjective to be ‘local to the noun’. But how local? Having criticized both of Cinque’s proposals (that indirect modifiers are always clausal, and that indirect modifiers are past some underspecified semi-definite dP), it is now necessary to precisely characterize both what kind of syntactic boundary determines the relevant kind of locality for modification, and why that syntactic boundary has the semantic effect of separated Function Application from Predicate Modification.
- *What is the semantics of nonintersective composition?* So far, I have given the actual semantics of these adjectives a deliberately simplified treatment, in order to focus on their syntactic environments and characterize in broad strokes their different modes of composition. It is only in abstract terms that we have discussed the process of an adjective being interpreted ‘relative’ to a noun, how it comes to modify along a specific noun-determined scale. But assuming that the conditions are met for nonintersective modification, how does this actually occur? How does the same adjective end up with different meanings in nonintersective composition with different nouns? And how does

the context affect these differences in meaning? And in particular, how does privative composition work? For this chapter I have been treating them as simple, but the problems posed to compositionality by privative adjectives raised in the introduction still need answers.

Chapters 3, 4, and 5 are dedicated to answering these questions, respectively.

3 The morphosyntax of intersection

In Chapter 2, I argued that we should have a uniform two-place predicate analysis of ambiguous adjectives, where such adjectives always take a nominal argument that determines the scale or dimension of their evaluation prior to their individual argument. On this account, intersective readings always result from intersection of the modified noun with the $\langle e, t \rangle$ output of that two-place predicate after it composes with a phonologically null initial noun argument, supplied pragmatically by the context, defining the scale. But this account is, on its own, agnostic as to whether this contextual argument is realized in the syntax or if that saturation process takes place somehow pre-syntax.

The goal of this chapter is to argue for the syntactic reality of the contextually-supplied argument. The evidence comes from suppletive adjectives, where I argue that the unavailability of suppletive forms for certain adjectives with intersective meanings - across a wide variety of languages - suggests that some piece of syntactic structure is intervening between the adjectival root and the suppletion trigger, and that that piece of structure is this nominal argument. Having independent evidence for the syntactically-determined nature of the intersective reading bolsters the overall argument against lexical polysemy and in favor of a fully derivational, compositional analysis of the (non)intersective ambiguity. While no specific components of the analyses offered in other chapters depend on the conclusions argued for in this chapter, absent this syntactic evidence, the remainder of the account would lose a significant portion of its strength in parsimony.

3.1 Suppletion in Serbian

Suppletion is generally defined as the phenomenon where a single root is realized with different, unrelated phonological forms in different syntactic or morphological contexts, crucially without a corresponding change in semantic interpretation (Mel'čuk, 1994). In other words, suppletion occurs when the root that surfaces in a morphologically complex form is 'ety-

mologically unrelated' (Bobaljik 2012:72) to the root that surfaces in the base (simplex) form. Here, we will be concerned with the suppletion that can occur in the comparative and superlative forms of adjectives:

- (113) POSITIVE: good
 COMPARATIVE: better
 SUPERLATIVE: best

So in (113), the underlining represents the phonological realization of the adjective root in each case, and the root which appears in the comparative and superlative form, *be(t)*, is wholly unrelated to the positive form *good*. This kind of adjectival suppletion occurs widely cross-linguistically, but within languages has a very limited distribution, appearing with only a small and cross-linguistically consistent set of adjectives of which *good* and *bad* are the most frequent (Bobaljik 2012). Conveniently, those particular adjectives are paramount examples of the nonintersective ambiguity, and the interaction of suppletion and this ambiguity will turn out to constitute useful evidence for the syntactic structure of (non)intersective meanings.

The secondary claim of this section will be that accounting for this ambiguity does not require positing that Vocabulary Insertion rules in Distributed Morphology (Halle and Marantz, 1993) be sensitive to semantic diacritics on roots or category-assigning heads. The information that determines the (non)intersective reading, I will argue, is present in the syntax (in line with Despić (2019)), but not as a feature on roots (against Despić, and by extension, against this particular aspect of the ontology of roots in Harley (2005)). This doesn't constitute an argument that they are not, or can never be, but just that this particular set of data does not independently justify such a move. Along the way, we will also see that this data bears on some fundamental questions in the typology of suppletion (Bobaljik, 2012).

3.1.1 The ban on intersective suppletives in Serbian

Recall that the nonintersective ambiguity is illustrated by sentences like (114), where the adjective *good* can be alternately interpreted nonintersectively - ‘good at being the noun’ - or intersectively - ‘good as a person, good in general’.

- (114) Robin Hood is a good thief.
- a. Nonintersective: ‘Robin Hood is good as a thief’, ‘Robin Hood is a skilled thief’
 - b. Intersective: ‘Robin Hood is a thief and a good person’, ‘Robin Hood is a good person for a thief’

Despić and Sharvit (2011) and Despić (2019) observed that the (non)intersective ambiguity interacts in an unexpected way with suppletive morphology in Serbian (and other closely related languages). As stated above, suppletion is generally assumed to correspond to semantic consistency; that is, the meaning of the root which gets realized alternately as *good* or *be(t)* in (113) does not change, only its form. But in Serbian, this correspondence appears to be violated when it comes to (non)intersectivity.

- (115) On je dobar lopov (Despić, 2019, 511)
He is good thief
‘He is a good thief.’ (Int: ✓, NInt: ✓)

The same ambiguity from the English (114) is found in (115). It has both the nonintersective (NInt) reading, where the person being described is a skilled thief with any amount of moral character, and the intersective (Int) reading, where they are a morally good person and also a thief of irrelevant skill. But the picture becomes more complicated for the suppletive comparative and superlative forms.

- (116) a. On je bolji lopov (Despić, 2019, 511)
He is better thief
‘He is a better thief.’ (Int: ✗, NInt: ✓)
- b. On je najbolji lopov
He is best thief
‘He is the best thief.’ (Int: ✗, NInt: ✓)

Here, in both of the suppletive forms in (116), the intersective reading disappears. These

can only mean that the person is a more skilled, or the most skilled, thief; the interpretation where they are a thief and it is their moral character being compared to others is unavailable.

This effect is due to the suppletion, not simply the addition of the comparative or superlative morphology. For Serbian adjectives that show this ambiguity in the positive form and remain regular throughout the comparative and superlative, the ambiguity remains.

- (117) a. Petar je inteligentan teniser (Despić, 2019, 511)
 Peter is intelligent tennis-player
 ‘Peter is an intelligent tennis player.’ (Int: ✓, NInt: ✓)
- b. Petar je inteligentniji teniser
 Peter is more-intelligent tennis-player
 ‘Peter is a more intelligent tennis player.’ (Int: ✓, NInt: ✓)
- c. Petar je najinteligentniji teniser
 Peter is most-intelligent tennis-player
 ‘Peter is the most intelligent tennis player.’ (Int: ✓, NInt: ✓)

A noun phrase like *intelligent tennis player* in (117) has a nonintersective reading, where their tennis playing is particularly intelligent (a high ‘tennis IQ’) but their overall intelligence might not be high, and an intersective reading, where they are a generally intelligent person who also happens to play tennis, with no particular degree of insight or strategy. Both of these readings are available for the positive, comparative, and superlative forms. This means that our analysis of why the intersective reading disappears in (116) cannot appeal to some incompatibility between intersectivity and the semantics of the comparative or superlative morphemes, unless we find some particular friction between the semantics of *dobar* itself (and we will see later on that this is also not a viable path).

Given the canonical description of suppletion as requiring semantic regularity across morphological irregularity, there is of course an immediate question about whether data like this should cause us to expand our definition of suppletion to include semantic irregularity, or consider cases like (116) as something other than suppletion. Luckily, Despić (2019) provides an analysis on which we have to resort to neither of these options, by arguing that single roots always have a uniform semantics. The present paper’s analysis will do something similar, in the sense that it also sidesteps these large questions about the theory

of suppletion. Hopefully, from either perspective, we can demonstrate that this data can be considered genuinely suppletive without requiring any such radical rethinking. I now turn to reviewing the Despić (2019) analysis of this data. It is worth noting that it is informed by more than we have covered so far, including other contrasts with change-of-state verbs and definiteness, but for now we will focus on the comparative suppletion and return to those data wherever they might become a problem for the proposed counter-analysis down the line.

3.1.2 The diradical analysis

Most theories of the (non)intersective ambiguity make no prediction that it should interact with suppletion in the adjective. Despić and Sharvit (2011) rightly argue that Larson-style analyses, where the ambiguity arises from the existence of two distinct arguments in the noun for the adjective to choose from to modify, have difficulty accounting for this effect. We could conceive of a story for the Larson account where the presence of the comparative morphology disrupts a locality relationship between the adjective root and the noun which would be required in order to take the noun-internal event variable as the adjective's argument. But such a story not only requires providing a plausible mechanistic explanation for such a locality condition, but also seems *prima facie* implausible because it predicts the opposite of the facts. This kind of locality story would rule out the nonintersective reading in the comparative, which is the opposite of the observed pattern, and it has no way to distinguish suppletive from regular comparative forms, being only interested in the presence or absence of the comparative affix. In general, any account which takes the adjective to be structurally and semantically identical across its ambiguous uses and locates the ambiguity entirely within the noun, as Larson does, will struggle even at the conceptual level to make itself sensitive to suppletion.

Instead, the analysis of this nonintersective suppletion restriction offered by Despić and Sharvit (2011) and Despić (2019) is a diradical variety of 'Blame the Adjective'. They posit

two separate denotations for the intersective and nonintersective uses of *good*. I present a slightly simplified version of the semantics for these roots, hopefully without losing any relevant information:

- (118) a. The intersective root, *good**: (adapted from (Despić, 2019, 513))
 $\llbracket \mathbf{good}^*-d_i \rrbracket^{w,C} = [\lambda x \in D : \text{context } C \text{ supplies an assignment, } g_C, \text{ and a scale, } GOOD_{C,w}, \text{ that ranks people according to their 'goodness' in } w . \text{ the ranking of } x \text{ on } GOOD_{C,w} \text{ is at least } g_C(i)]$
- b. The nonintersective root, *good***:
 $\llbracket \mathbf{good}^{**}-d_i \rrbracket^{w,C} = [\lambda P \in D_{\langle s, \langle e, t \rangle \rangle} . \lambda x \in D : \text{context } C \text{ supplies an assignment, } g_C, \text{ and a scale, } S_{P,w}, \text{ that ranks individuals by their } P\text{-skills in } w . \text{ the ranking of } x \text{ in } w \text{ on } S_{P,w} \text{ is at least } g_C(i)]$

So, the surface form *dobar* corresponds to two roots, one of which is always intersective and has the semantics in (118a), and the other of which is always nonintersective and has the semantics in (118b)

- (119) a. $\sqrt{DOBR^I} = \llbracket \mathbf{good}^*-d_i \rrbracket^{w,C}$
b. $\sqrt{DOBR^{NI}} = \llbracket \mathbf{good}^{**}-d_i \rrbracket^{w,C}$

These diacritic markers I and NI on the roots are taken as ‘markers of the two relevant semantic sub-types ($\langle e, t \rangle$ and $\langle \langle s, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$, respectively). They are visible to morpho-phonology and may be referred to by morpho-phonological rules.’ (Despić, 2019, 517) Here, then, the semantic subtypes I and NI are properties that roots can have, following the root ontology of Harley (2005). Despić also offers an alternate account, where the diacritics are located on the adjectival category-determining head, though tentatively concludes against this option. What is relevant is that Vocabulary Insertion rules which determine suppletion can see and make reference to these diacritics. Thus, a Vocabulary Insertion rule realizes the nonintersective root (119b) as the a suppletive form *bol-* in the comparative environment. On the other hand, the intersective root (119a) lacks a specific rule, and so will only ever be realized as *dobar* as a result of the general rule targeting the semantically-unspecified root.

- (120) a. $\sqrt{DOBR^{NI}} \rightarrow \text{bol- / -] CMPR}$
b. $\sqrt{DOBR} \rightarrow \text{dobar}$

The use of both a diacritic-specific rule (120a), applying to only the NI form, and a root-general rule (120b), applying to both the I and NI form, reinforces that this analysis requires not just positing two root forms for *dobar*, but also encodes a relationship between those forms as instances of one more general root. Thus, the I/NI distinction is not reducible to simple homophony, but requires additional theoretical machinery.

So, on this account, the fact that the suppletive forms only ever allow the nonintersective reading is simply because only the nonintersective root is allowed to be realized as the suppletive form. The positive *dobar* is homophonous, but the suppletive *bol-* is not. The comparative rule (120a) also covers the superlative case, given the Containment Hypothesis:

(121) *The Containment Hypothesis* (Bobaljik, 2012, 31)
 The representation of the superlative properly contains that of the comparative.

(122) *Structure of a superlative on the Containment Hypothesis*
 [[[ROOT] -CMPR] -SPRL]

We'll return to discussing the motivation behind the Containment Hypothesis, and the relevance of the current data, below. For now, just note that the rule (120a) will apply any time the comparative affix is present, regardless of whether the superlative morpheme is also present, and so the correct suppletive form with the nonintersective interpretation will be derived for both *bolji* 'better' and *najbolji* 'best'.

This analysis captures the suppletion pattern in a way that a Larson-style 'Blame the Noun' analysis cannot. However, it faces some issues of its own.

3.1.3 Issues with the diradical analysis

To be clear, not all of the problems identified here are avoided by the analysis proposed later on in Section 3.2.3. However, they are all concerns for a proponent of the diradical option.

First, general parsimony conditions should penalize homophony in a theory, and in particular accidental homophony. Given that neither root semantics in (118) is derived from the other, we have no explanation for why the single word *dobar* receives both of those

interpretations; or, stated in the inverse, no explanation for why those two interpretations which are intuitively but not derivationally related both happen to get realized as *dobar*. Of course, this isn't just a problem for *dobar*. When we consider nonsuppletive but ambiguous words like *intelligentan* 'intelligent', it raises the question of whether we also need to posit two roots here to get the intersective and nonintersective readings - I would assume that, extending the present analysis, we would - and in doing so, why the Vocabulary Insertion rules for this root family never make reference to the I/NI diacritics. This leads us to one of two problematic conclusions: either we have widespread homophony across all adjectives that show this ambiguity and for nonsuppletive adjectives that homophony is never broken (the two roots are always realized with the same phonological form), or the ambiguity for *dobar* is due to split roots and the ambiguity for nonsuppletive words like *intelligentan* is due to some other, unknown factor. Neither conclusion is too theoretically satisfying.

Second, nothing in this analysis explains why it is the NI root that suppletes and not the I root; this is true both for *dobar* and for *zao* 'bad', that shows the same pattern with respect to suppletion. Nothing seems to logically prevent the opposite situation, where the I root has a special Vocabulary Insertion rule. Despić tentatively hypothesizes:

'In the case of NI-comparison, however, the adjective and the noun both matter; i.e. the [AP+NP] combination establishes the relevant scale/degrees (the scale/degree of "successfulness" in stealing, playing a cello etc.) Thus, it might not be unexpected that the NI-reading would prevail in suppletive comparatives, given that on many semantic analyses of comparatives (e.g. Heim (2000)), the comparative morpheme (or DegP containing it) scopes over the adjective and the noun modified by the adjective.' (Despić, 2019, 519)

However, I don't see how this proposal functions. Following this line of logic, it makes sense that the NI reading might prevail in comparatives in general, if the relevant factor is that the comparative morpheme scopes over both the adjective and the noun in combination, and thus the meaning to which the comparative applies is one where the adjective is interpreted relative to the noun. But this makes no predictions regarding suppletion. As we have seen with (117), the intersective reading is perfectly compatible with the comparative

when it is not suppletive, which should not be the case if it's just about the comparative containing the noun as well. This argument would require a structural difference in terms of the DegP scope between suppletive and regular comparatives, at which point we could exploit that difference to explain the intersectivity contrast, rather than posit it in Vocabulary Insertion rules.

Third, why doesn't the intersective root have a comparative or superlative form at all? The rules in (120) should block the nonintersective root from getting realized as *dobar* in the comparative environment, but they shouldn't block the intersective root from the same thing. Absent any separate constraint, $\sqrt{DOBR^I}$ + CMPR should be realized as *dobriji*, with the intersective interpretation. But this doesn't occur. Despić attributes this to a language-specific constraint: 'Serbian resists comparative/superlative doublets; i.e. *dobar* (or any other adjective) in Serbian cannot have more than one morphological comparative/superlative form.' (Despić, 2019, 518) Setting aside that a theory would ideally predict this parameter from something independent about the language, it is at least reasonable to say that's not the job of a theory of intersectivity.

3.2 Suppletion across languages

None of the objections just developed are fatal. They are simply deficits in parsimony or aspects of the pattern which are left unexplained, not direct empirical challenges. As such, if this nonintersective suppletion pattern - and therefore its analysis - were limited to Serbian and its close relatives, it would be tolerable to attribute these deficits to language-specific accidents or constraints. But this is not the case.

3.2.1 The Nonintersective Suppletion Generalization

The pattern observed in Serbian - that suppletive morphological comparative and superlative forms disallow the intersective reading - holds up across a significantly wider sample of

languages. Across a sample survey of sixteen languages from multiple languages families¹¹, only one language violates this:

Do the CMPR/SPRL forms have:	Serbian	Russian	Polish	Swedish	German	Dutch	English	Persian	Italian	Basque	Romanian	Lithuanian	Turkish	Hebrew	Korean	Mandarin
...suppletion?	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓						
...Int reading?										✓	✓	✓	✓	✓	✓	✓

Table 1: Availability of the intersective reading across comparative and superlative forms

Not every language in the sample allows the nonintersective reading even in the positive form - the Romanian, Korean, and Mandarin *good* words are uniformly intersective across all uses, and so perhaps they make less of an impact in this sample. However, leaving Basque aside for a moment, we can see the emergence of a potential generalization:

- (123) THE NONINTERSECTIVE SUPPLETION GENERALIZATION (NSG): If the comparative or superlative form of an adjective is suppletive, then it lacks the intersective reading.

We have already seen that Serbian obeys this generalization, but the NSG also holds in English:

- (124) a. Robin Hood is a good thief. (Int: ✓, NInt: ✓)
 b. Robin Hood is a better thief. (Int: ✗, NInt: ✓)
 c. Robin Hood is the best thief. (Int: ✗, NInt: ✓)

The majority of the languages in the sample show this same behavior: Russian, Polish, Swedish, Dutch, Italian, German, and Persian all have suppletive comparatives and superlatives, and similarly disallow the intersective reading for those suppletive forms.

- (125) Russian
 a. On xoroshiy vor
 he good thief
 ‘He is a good thief’ (Int: ✓, NInt: ✓)

¹¹This is an obviously preliminary sample for making broad claims about typological universals, and the data collection itself reflects that. For most languages, data was gathered from one to three speakers (except for Persian, where the pattern was verified for twenty speakers).

- b. On luchshiy vor
he better thief
'He is a better thief' (Int: ✗, NInt: ✓)
- c. On luchshiy vor
he best thief
'He is the best thief' (Int: ✗, NInt: ✓)

(126) Polish

- a. On jest dobrym złodziejem
he is good thief
'He is a good thief' (Int: ✓, NInt: ✓)
- b. On jest lepszym złodziejem
he is better thief
'He is a better thief' (Int: ✗, NInt: ✓)
- c. On jest najlepszym złodziejem
he is best thief
'He is the best thief' (Int: ✗, NInt: ✓)

(127) Swedish

- a. Han är en bra tjuv
he is a good thief
'He is a good thief' (Int: ✓, NInt: ?¹²)
- b. Han är en bättre tjuv
he is a better thief
'He is a better thief' (Int: ✗, NInt: ✓)
- c. Han är den bästa tjuven
he is the best thief
'He is the best thief' (Int: ✗, NInt: ✓)

(128) Dutch

- a. Hij is een goede dief
he is a good thief
'He is a good thief' (Int: ✓, NInt: ✓)
- b. Hij is een betere dief
he is a better thief
'He is a better thief' (Int: ✗, NInt: ✓)
- c. Hij is de beste dief
he is the best thief

¹²Swedish consultants reported some degradation for the nonintersective meaning of *bra* 'good' in the positive form, attributed to competition with the unambiguously nonintersective *duktig* 'skilled'. This doesn't constitute an issue for the NSG, given that the intersective meaning preferred in the positive form still disappears in the suppletive forms - arguably, it is an even more striking piece of support for the pattern that it overrides the preference.

‘He is the best thief’ (Int: ✗, NInt: ✓)

(129) Italian

- a. *È un ladro buono*
he a thief good
‘He is a good thief’ (Int: ✓, NInt: ✓)
- b. *È un ladro migliore*
he a thief better
‘He is a better thief’ (Int: ✗, NInt: ✓)
- c. *È il miglior ladro*
he the best thief
‘He is the best thief’ (Int: ✗, NInt: ✓)

German is particularly illustrative, as *gut* ‘good’ is suppletive but *schlecht* ‘bad’ is regular:

- (130)
- a. *Er ist ein guter Dieb*
He is a good thief
‘He is a good thief’ (Int: ✓, NInt: ✓)
 - b. *Er ist ein besserer Dieb*
He is a better thief
‘He is a better thief’ (Int: ✗, NInt: ✓)

- (131)
- a. *Er ist ein schlechter Dieb*
He is a bad thief
‘He is a bad thief’ (Int: ✓, NInt: ✓)
 - b. *Er ist ein schlechterer Dieb*
He is a worse thief
‘He is a worse thief’ (Int: ✓, NInt: ✓)

The regular comparative *schlechterer* allows the intersective reading, while the suppletive *besserer* does not. Thus, even within a single language, the alternation between suppletive and nonsuppletive forms appears to affect the possible interpretations. An even more striking example comes from Persian, which has two comparative series for *khub* ‘good’: one regular, and one suppletive.

- (132) *Dozd-e khub*
thief-EZ good
‘He is a good thief’ (Int: ✓, NInt: ✓)

- (133) a. *Dozd-e beh-tar*
thief-EZ good-CMPR

- ‘He is a better thief’ (Int: ✗, NInt: ✓)
 b. Dozd-e khub-tar
 thief-EZ good-CMPR
 ‘He is a better thief’ (Int: ✓, NInt: ✗)

Here, not only is the suppletive form *behtar* strictly nonintersective, the regular form *khubtar* is strictly intersective and disallows the nonintersective reading. Without going into too much detail on Persian specifically, we can likely attribute this to a pragmatic blocking effect: the availability of unambiguously nonintersective *behtar* blocks that use of potentially ambiguous *khubtar*. Justifying this kind of approach is Persian *bad* ‘bad’, which has only a regular comparative and both readings.

- (134) Dozd-e bad-tar (regular)
 thief-EZ bad-CMPR
 ‘He is a worse thief’ (Int: ✓, NInt: ✓)

The unavailability of the nonintersective reading is not typical of regular comparatives: Lithuanian, for example, which has fully regular morphological comparatives, allows both readings in the comparative and superlative to the extent that they are available in the positive form.

- (135) a. Jis yra geras vagis
 He is good thief
 ‘He is a good thief’ (Int: ?, NInt: ✓)
 b. Jonas yra geresnis vagis negu Marija
 Jonas is better thief than Marija
 ‘Jonas is a better thief than Marija’ (Int: ?, NInt: ✓)
- (136) a. Jis yra geras burtininkas.
 He is good wizard
 ‘He is a good wizard’ (Int: ✓, NInt: ✓)
 b. Jonas yra geresnis burtininkas negu Marija
 Jonas is better wizard than Marija
 ‘Jonas is a better wizard than Marija’ (Int: ✓, NInt: ✓)

The contrast between (135) and (136) illustrates that the availability of these readings is also dependent on context and the particular noun being used; the Lithuanian speaker consulted found it somewhat degraded to get the intersective reading at all with *vagis* ‘thief’,

but the crucial point is that it was no less available in the comparative. Changing the noun to one that was less inherently negatively valued, like *burtininkas* ‘wizard’, made the intersective reading available, and again in both the positive and comparative forms.

Turkish also lacks suppletion, but as a result of lacking morphological comparatives and superlatives altogether. Instead, these are formed analytically, by modifying the positive form with distinct comparative and superlative adjectives (like the English pattern for *more beautiful*, *most beautiful*, as opposed to **beautifuler*, **beautifullest*).

(137) Turkish

- a. Ali iyi hırsız.
Ali good thief
‘Ali is a good thief.’ (Int: ✓, NInt: ✓)
- b. Ali daha iyi hırsız.
Ali more good thief
‘Ali is a better thief.’ (Int: ✓, NInt: ✓)
- c. Ali en iyi hırsız.
Ali most good thief
‘Ali is the best thief.’ (Int: ✓, NInt: ✓)

Hebrew falls into this class as well. Some other languages in the sample putatively satisfy the NSG by not allowing the nonintersective reading at all, having no suppletive form and only intersective meanings for *good* anyway. For example, Romanian is like Turkish in having only analytic comparison, but is unlike Turkish in having no ambiguity in the positive form.

(138) Romanian

- a. El e un hot bun.
he is a good thief
‘He is a good thief’ (Int: ✓, NInt: ✗)
- b. El e un hot mai bun.
he is a good more thief
‘He is a better thief’ (Int: ✓, NInt: ✗)
- c. El e cel mai hot bun.
he is the more good thief
‘He is the best thief’ (Int: ✓, NInt: ✗)

Expressing the subjective meaning in Romanian requires an entirely different adjective, e.g., *capabil* ‘capable’ or *priceput* ‘able’, and these adjectives are similarly unambiguous. In

this sample, Korean and Mandarin form the remainder of the unambiguous analytic class alongside Romanian.

3.2.2 The Basque problem

In this limited sample, the lone exception to the NSG is Basque, which has a suppletive comparative form, and then both a regular and suppletive option for the superlative. Both the intersective and nonintersective readings are available across all of the forms, suppletive or regular.

- (139)
- a. Lapur ona da
thief good is
'They are a good thief' (Int: ✓, NInt: ✓)
 - b. Lapur hobea da
thief better is
'They are a better thief' (Int: ✓, NInt: ✓)
 - c. Lapur onena da
thief best is
'They are the best thief' (Int: ✓, NInt: ✓)
 - d. Lapur hoberena da
thief best is
'They are the best thief' (Int: ✓, NInt: ✓)

While I won't offer a full analysis of the Basque case here, it is worth noting that this is not the only way in which the Basque suppletive paradigm is exceptional. Recall the Containment Hypothesis from (121). This hypothesis was motivated by another generalization:

- (140) THE COMPARATIVE-SUPERLATIVE GENERALIZATION (Bobaljik, 2012, 2)
If the comparative degree of an adjective is suppletive, then the superlative is also suppletive (i.e. with respect to the positive).

What (140) prevents is so-called ABA constructions, where the comparative is suppletive but the superlative is regular. This *ABA generalization is near-universally held, but 'the Basque pattern presents an apparent ABA pattern... a lone adjectival counterexample' (Bobaljik, 2012, 112) to (140). The comparative *hobea* is suppletive, but the superlative *onena* is regular (even though there is also a suppletive comparative - this contrasts with

Persian where the existence of both a suppletive comparative and a regular superlative is not problematic because there also exists a regular comparative, on top of which the regular superlative can be built unproblematically). So, though it's not clear how, perhaps we now have converging evidence from both the CSG and the NSG that the Basque paradigm is in some way structurally non-standard.

What is clear is that Bobaljik's proposed explanation for Basque's *ABA will not work to explain its abnormality with respect to the NSG. In short, Bobaljik proposes analyzing the *-(r)en* superlative suffix in *on-en/hobe-ren* as a conventionalized genitive, not a true superlative marker, i.e., 'the good of the good' rather than literally 'the best'. This would obviate the *ABA problem, because there is no underlying CMPR morpheme in these genitive constructions. However, it would not explain the NSG violation: here, the primary issue is with the suppletive comparative itself, *hobe*, and the fact that it receives the intersective interpretation. Reanalyzing the superlative constructions is irrelevant to this. The present data is thus a double-edged sword for the *ABA generalization: it both provides novel evidence for treating Basque as a legitimate outlier and simultaneously undermines the viability of the preexisting explanation for that outlier status.

3.2.3 A monoradical analysis

The existence of the NSG, even if it does not end up robust as an ironclad generalization across a larger sample size, makes the problems highlighted above for the Despić (2019) analysis all the more pressing. Recall that we had no explanation for why it was the NI root that suppleted, or why the I root didn't form its own regular comparative. Had the pattern been Serbian-specific, we might have accepted these as random variation. But given the persistence of the pattern across language families, an analysis which has some explanatory power for why the suppletion facts play out in this way, rather than the other logical possibilities, is significantly more desirable. Before making it explicit, however, I want to also note the status of this analysis in context of the rest of the dissertation - the core compo-

nents that will be taken into the subsequent chapters are (i) the use of only a single root for the adjective and (ii) the presence of intervening structure in the intersective reading. The particular semantic implementation here is phrased to vary minimally from Despić’s (2019) account, to illustrate the key differences; the actual denotations for *good*, for example, will be different from those developed later on. With that in mind, let’s proceed with the analysis.

I propose that the unavailability of the intersective reading in suppletive forms is due to the interaction of two factors. First, suppletion is local (Bobaljik and Harley, 2017). Vocabulary Insertion rules which determine suppletion may only reference adjacent environments, and so the presence of any intervening structure between the root and the suppletion-triggering morpheme (here, the comparative suffix) would block the insertion of the suppletive form. Second, picking up on a proposal from Arregi and Nevins (2014), I posit that the intersective reading is derived from the insertion of a phonologically null but syntactically represented argument, which is semantically broad and made available by context, and serves as the first argument to a uniform higher-type denotation for the adjectival root. Therefore, the intersective reading will be unavailable for suppletive comparatives because it requires structure which disrupts the local adjacency requirement for the insertion of the suppletive form in the first place.

On the proposed analysis, we require only a single root for *good*, equivalent to the higher-typed nonintersective root *good*** from Despić and Sharvit (2011):

$$(141) \quad \llbracket \mathbf{good-d}_i \rrbracket^{w,C} = [\lambda P \in D_{\langle s, \langle e, t \rangle \rangle} . \lambda x \in D : \text{context } C \text{ supplies an assignment, } g_C, \text{ and a scale, } S_{P,w}, \text{ that ranks individuals by their } P\text{-skills in } w . \text{ the ranking of } x \text{ in } w \text{ on } S_{P,w} \text{ is at least } g_C(i)]$$

The nonintersective reading is derived in the same way: the noun, e.g., *thief*, saturates the λP argument of (141) and provides the type of skill that the individual argument will be measured on. The intersective reading is derived by the insertion of a relatively semantically bleached, null argument, e.g., PERSON, into the adjective’s structure before it combines with the noun. That argument saturates λP and gives a scale of default ‘goodness’, like in

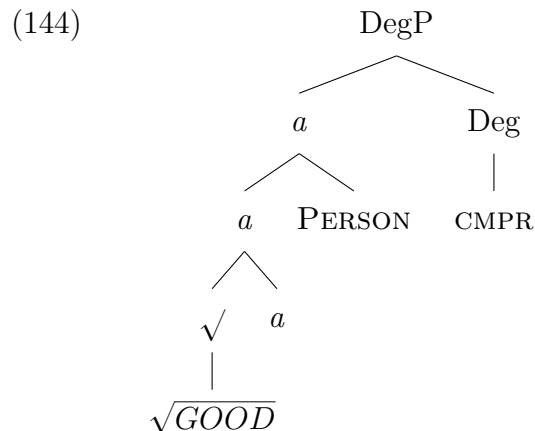
Despić and Sharvit’s intersective *good**. Then, the remaining $\langle e, t \rangle$ type adjective composes with the noun via Predicate Modification to derive the intersective reading. Simplifying the denotation in (141) considerably for space, we can give sketches of the compositional derivations:

- (142) Nonintersective reading:

$$\begin{aligned} \llbracket \text{good thief} \rrbracket &= \llbracket \text{good} \rrbracket (\llbracket \text{thief} \rrbracket) \\ &= [\lambda P_{\langle s, \langle e, t \rangle \rangle} \lambda x_e . x \text{'s goodness as } P \text{ exceeds } d] (\lambda x_e . x \text{ is a thief}) \\ &= \lambda x_e . x \text{'s goodness as a thief exceeds } d \end{aligned}$$
- (143) Intersective reading:

$$\begin{aligned} \llbracket \text{good thief} \rrbracket &= \llbracket \text{good} \rrbracket (\llbracket \text{person} \rrbracket) \cap \llbracket \text{thief} \rrbracket \\ &= [\lambda P_{\langle s, \langle e, t \rangle \rangle} \lambda x_e . x \text{'s goodness as } P \text{ exceeds } d] (\lambda x_e . x \text{ is a person}) \\ &\quad \cap [\lambda x_e . x \text{ is a thief}] \\ &= [\lambda x_e . x \text{'s goodness as a person exceeds } d] \cap [\lambda x_e . x \text{ is a thief}] \\ &= \lambda x_e . x \text{'s goodness as a person exceeds } d \wedge x \text{ is a thief} \end{aligned}$$

The syntactic structure for the intersective is thus like (144). I intend to be noncommittal many parts of this implementation: for one, does PERSON merge above or below the category-determining head *a*? We might look to work like Anagnostopoulou and Samioti (2013) or Marantz (2013) on the function of material in and out of those domains. For now it suffices that it intervenes between the root and CMPR; I’ll return to this question at the very end of the dissertation when we consider the role of categorizing heads in structuring meaning. It should also be unimportant whether the noun occurs below or above CMPR - I leave it out of (144) for simplicity, and assume that this is a question outside of our scope entirely, but see Heim (2000) for relevant arguments.



This intervening argument is only going to be present in intersective readings¹³. With this compositional setup, we require only a single Vocabulary Insertion rule for the comparative environment, and no semantic diacritics. The root for *good* is realized as *bet* in the environment of the comparative, and otherwise simply as *good*.

- (145)
- a. $\sqrt{GOOD} \rightarrow \text{bet- / -] CMPR}$
 - b. $\sqrt{GOOD} \rightarrow \text{good}$

This analysis has significant conceptual advantages over that of Despić (2019). First, it avoids positing homophony of any kind, deriving both readings from the same root in a fully compositional manner. Second, it explains why the nonintersective reading is the one available in suppletive forms. Since that is the reading in which the adjective is interpreted relative to the noun, it makes sense that the noun saturates the property argument of the adjective, while in the intersective reading, the adjective is interpreted relative to a default property retrieved from the context. Assuming that that default property is syntactically represented allows us to derive the suppletion facts from a preexisting locality condition, in a way that is correctly asymmetrical rather than logically reversible.

¹³One of the arguments from Chapter 2 is that nonintersective readings *can* be derived using this structural argument, even if this is the standard way to get intersective readings. However, this doesn't predict that nonintersective readings should also cause intervention, because that is only one of two ways to reach nonintersectivity - it can also be composed without such an intervener, in the local Function Application context. Intersective readings, on the other hand, cannot be duplicated via Function Application, so they will always be impossible with a suppletive comparative form.

3.2.4 Lingering issues with the monoradical analysis

However, it does have remaining problems. It does not explain the impossibility of *dobriji*, i.e., why the intersective reading does not get realized with a regular comparative. It's not the case that this never happens - Persian regular *khubtar* is one counterexample, and in English *gooder* is attested if non-standard, and when attested it does refer to this broader intersective notion of 'goodness'. But there is a general prohibition or at least degradation of these forms, and we cannot appeal to language-specific constraints like 'Serbian disallows comparative doublets' to explain it. This problem exists for both analyses.

One final problem is the existence of intersective readings where the relevant property is something other than moral 'goodness'. For example, you might imagine a conversation at the chess club of a music school, talking about the levels of chess skill that various instrumentalists have. Applying our present, compositional analysis, you might expect this to be derived similarly to the moral intersective reading, by the insertion of some null contextually-supported argument like *CHESSPLAYER*. However, this would predict that such readings would be unavailable with suppletive forms, and this is incorrect:

(146) She is a better violinist than average. (Moral Int: ✗, NInt: ✓, Chess Int: ✓)

The pattern in (146) holds up across all suppletive languages in the sample. Rescuing the present analysis this requires treating these cases as not at all parallel to (144) (perhaps locating the problem in the noun entirely, e.g., reinterpreting *violinist* to mean *chess-playing violinist* and treating the adjective as nonintersective). However, given that they require significantly more contextual support than the moral intersective reading, perhaps it is not so unreasonable to derive them from distinct pragmatic mechanisms.

I won't offer a solution to any of these issues here, but hopefully it is still convincing that a monoradical approach is overall more satisfying than a more stipulative diradical one which struggles with all the same concerns, among others. Zooming out somewhat, I want to highlight the difficulty that competing theories of the (non)intersective ambiguity

would have in accounting for its sensitivity to morphological factors like suppletion. The diradical theory considered in this section is effectively a version of Siegel’s doublet theory, and so the latter certainly could be made to derive morphological differences - in fact, it might be more satisfactory for a lexical doublet theory if there were differences in the phonological realization of the alternates - but it would be subject to all of the same criticisms levied at Despic’s account here. The data in this section poses a much more significant obstacle, I think, to a Maienborn-style post-compositional account as discussed in Section 2.1.3. Given that the pragmatic variable resolution process occurs entirely after the syntactic derivation has generated and shipped the structure to the conceptual interface (in fact, even after any narrow semantic-specific operations at that interface take place), there is no plausible way for variable resolution of this type to affect Vocabulary Insertion or vice versa. Given that, the suppletion pattern constitutes an argument both in favor of this particular syntactic implementation of a null argument, and also generally in favor of a structure-sensitive derivational approach to ambiguity.

3.3 Interim summary

This section has motivated the syntactic realization of the intersective nominal argument with the observation that across languages, suppletive comparative and superlative forms fail to show the intersective reading: the Nonintersective Suppletion Generalization. This is derived by locating the source of the ambiguity in syntactic structure, positing the existence of a null argument for the intersective reading which blocks the necessary locality for suppletion. This fully compositional analysis is shown to be more explanatory and parsimonious than a diradical alternative, and refutes the need for semantic diacritics subtyping for (non)intersectivity in Distributed Morphology. Finally, the identification of Basque as the sole exception to the NSG potentially provides semantic evidence that its suppletive paradigm is structurally special, in a way that converges with its status as the sole *ABA violation.

3.3.1 A note on compositionality and explanation

Finally, I want to end this chapter with a small but critical comment on the role that morphosyntactic evidence like this plays in motivating a compositional analysis of this kind of phenomenon. In Chapter 1, I identified the goal of the dissertation as developing a theory of ambiguity that is fully compositional, and discussed some of the criteria for meeting this, the standards that a compositional theory of this kind might be held to. Recall that with sufficient freedom in manipulating either the lexicon, or syntactic structure, achieving compositionality is trivial (Partee, 1984; Zadrozny, 1994), Zadrozny 1994. I have levied one half of this criticism against various preexisting theories so far in this dissertation: theories that rely on duplicating lexical entries for adjectives as a solution to ambiguity, either in the form of words like Siegel’s doublet theory or roots like Despić’s diradical theory, trivialize accomplishing compositionality to the extent that they manipulate the basic building blocks in a way that loses out on explanation. And I do think this is true.

But it would be just as easy to levy the inverse criticism at the theory that I lay out in Chapter 2: extracting the required complexity from the lexicon by assuming that there’s only one underlying adjective form and then introducing that complexity into the syntax, by assuming that any time we need the other form an argument conveniently appears to saturate the adjective and create the needed form before it ever interacts with any other syntactic elements, functionally engages in the exact same kind of artificially compositional hacking, just on the syntax side. This seems to be exactly one of the ‘clever maneuvers in an arbitrary theory of syntax’ that Partee (1984) cautioned us against. This kind of objection was put even more forcefully, and in a very useful way in the specific context of adjectival modification, by Bouchard: ‘previous analyses of adjectival modification all have in common that they add various covert elements and operations to the theory (new lexical categories of adjectives or nouns, new syntactic categories, multiple lexical entries for some adjectives, new movement operations triggered by new functional features, etc.) The added elements are not accountable to either interfaces, so that compositionality is only satisfied

in a technical way. Compositionality that is merely technical loses the motivation for the principle' (Bouchard 2002:9) And again, what I have been proposing in Chapter 2 is exactly this kind of solution: it adds a covert element to the syntactic theory in order to satisfy compositionality, which threatens to be a merely technical satisfaction. Bouchard's notion of 'accountability to the interfaces' - forcing the theorist to make only moves which have empirical grounding in visible linguistic effects, preventing maneuvers that begin and end entirely within the mechanistic ivory tower of syntax and do not have surface reflexes in either form or meaning - is an enormously valuable one here to maintain a strict commitment to non-technical compositional solutions.

It is precisely because of this that the suppletion data is so important to the present proposal. Stipulating the syntactic status of the contextually supplied nominal argument is perfectly fine insofar as we are concerned with building a compositional theory that *works*. But insofar as we are concerned with building a compositional theory that is *accountable*, what we need is a reflex of that decision that is visible at the interfaces. Of course, it is obvious that this argument would have semantic consequences, but the ambiguity itself is the reflex at the conceptual-intensional interface, and thus cannot be reasonably used to justify a particular syntactic implementation of ambiguity. The fact that the morphophonological form of the adjective changes in a way that is sensitive to the ambiguity, and explained most cleanly (by reference to strong syntactic principles that we have substantial preexisting support for, like locality) by the assumption of this intervening argument, is what makes the theory interface-accountable. It is in this sense that while nothing technical in the remainder of the dissertation relies on the arguments in this chapter, it is crucial for the account's ability to satisfy a strong standard of compositionality in a genuinely explanatory way.

4 Number and kind modification

In Chapter 2, I argued in favor of an account of the (non)intersective ambiguity that distinguished between direct modification and indirect modification as syntactic concepts: adjectives which are sufficiently local to the noun modify it directly (resulting in composition via Function Application), and adjectives which are non-local modify the noun indirectly (resulting in composition via Predicate Modification). In this chapter, I aim to further specify some currently vague components of this analysis. First, what determines locality of an adjective; how close is close enough to be a direct modifier? And second, why do adjectives within that local domain compose differently from adjectives outside it?

The answer to the first question will end up being grammatical number. Adjectives that combine with the noun prior to the introduction of whatever element performs the atomizing function of number - which will vary across languages, between singular/plural morphology on the noun, on the determiner, or independently as a classifier - count as direct modifiers, and adjectives that combine with an already-atomized noun with number count as indirect modifiers.

The answer to the second question is that the addition of number fundamentally changes the type of semantic object that the noun denotes. Prior to composing with number, nouns denote kinds. Composition with number triggers an individuating/atomizing process that converts the kind denotation into an extensional set of objects which instantiate the kind, the set of the kind's members in a particular world. These ambiguous adjectives, then, take as their first argument something that is a kind, and can only combine with an already-individuated object type via Predicate Modification once their initial argument has been saturated. It is this kind argument which can optionally be provided by the context and takes the syntactically present but phonologically null slot that Chapter 3 just argued for.

This chapter will be organized as follows. First, I will give a brief introduction to the semantic notion of a kind and its syntactic correspondences. Then, I will introduce the theory of modification of Leffel (2014), who initially suggested the idea of treating ambiguity in

adjectival modification as an ambiguity between kind- and object-level modification. Leffel, I will argue, was fundamentally on the right track but lacked sufficient empirical basis for this proposal. As a result, I will then introduce a novel source of evidence that the kind/object treatment of the (non)intersective ambiguity is, in fact, correct, from Turkish. Finally, I will return to Bangla, and explain how the number-based account not only covers the data we have already seen but is further bolstered by it, with some necessary revisions to Leffel's original proposal.

4.1 Modifying kinds and objects

This section introduces the semantic distinction between kinds and objects, and the attempt from Leffel (2014) to apply it to the problem of ambiguous modification.

4.1.1 Reference to kinds

Generally, a plural noun tends to refer to some particular set with multiple members.

(147) Dogs ate the meat.

A sentence like this is true if, roughly speaking, there was some meat and there was some particular group (consisting of more than one dog) that ate it. But this is not the only option for the reference of a bare (determinerless) plural noun. Consider:

(148) Dogs eat meat.

To utter (148), you need not have in mind a particular group of dogs; in fact, you really shouldn't. Instead, (148) is making a claim about dogs *in general*: they have some broad tendency, as members of their species, to eat meat. Even though it is a claim about the whole collective of dogs, however, it isn't the same as a universal claim about *every* dog. Trying to give (148) the semantics of a universal statement doesn't quite work. A sentence like that is still true even if some dogs don't eat meat. So maybe it is more akin to:

(149) Most dogs eat meat.

But the truth conditions of (148) and (149) are not at all identical. Imagine, for example, if some highly contagious illness for dogs swept across the planet, and nine out of every ten dogs got the disease. Then we would certainly be willing to say (150a), but much less likely to say (150b).

- (150) a. Most dogs have the dog flu.
 b. Dogs have the dog flu.

The fact that the vast majority of dogs have the dog flu isn't enough to justify the statement that *dogs* have it, or at least not every interpretation of (150b). Perhaps if the situation persisted for generations, and achieved some kind of permanence in the dog population, then we might be willing to say that, yes, having the dog flu is one of the things that dogs do. But not simply because of frequency in the population, so something different is going on compared to (150a). In fact, a sentence like that can be true even if there is really quite a small percentage of the group that the statement applies to:

- (151) Mosquitoes carry malaria.

Only female mosquitoes can carry malaria, and only some small percentage of females do at any given time. Nowhere close to *most* mosquitoes carry malaria, and yet we are generally comfortable assenting to (151). So these kinds of sentences with bare plurals can refer to very few of their members. But not every property that's true of few members, of course:

- (152) Dogs have three legs.

While some dogs have three legs, no one is likely to agree that (152) is true, even if all dogs are disposed to the possibility of at some point having three legs in a way that not all mosquitoes are disposed to the possibility of at some point carrying malaria. So it is not enough that some members have the property, even though it can be enough for very few members to. Perhaps most interestingly, these bare plural nouns can also refer to situations where *no* members of the species have the stated property:

- (153) Dogs are common.

A sentence like (153) is true even though it's not true that all dogs, most dogs, some dogs, or even any individual dog is common; it doesn't really even make sense to speak of whether an individual can be common or not. (153) has to be read as making some claim about a property of the *species as a whole*, in a way that neither universality or genericity can capture.

It is problems like this that motivated Carlson (1977) to introduce the semantic notion of a *kind*. Bare plural nouns can, in some cases, refer to kinds of things, which denote something over and above simply the set that collects all those things together. The kind *dog* is not entirely reducible to the list of all dogs, though it does correspond to the property of *being a dog*. What exactly is a kind? From a semantic point of view, we can look at a kind as a particular type of individual, but not the traditional individual that has a particular manifestation in space and time; rather they are spatiotemporally 'discontinuous' (Chierchia, 1998, 348). A kind is not solely the set of its instances, but it is identified by its instances: 'the dog-kind in our world can be identified with the totality of dogs, the scattered entity that comprises all dogs, or the fusion of all dogs around' (Chierchia 1998:349). When we say that a kind is identified by or with this totality of its instances, we mean that a kind is a function from worlds into the sum of instances of the kind in a world. Give a kind a world, and it will give you the entity that consists of all of its members in that world - but, crucially, the kind is not definable *as* any particular one of those sets of members, or the superset of them across worlds. It is more like the concept that allows us to identify which things count as part of it.

I will return to a longer discussion of how we might think of kinds as a particular type of concept in Chapter 5; for now, we can proceed with the philosophical ambivalence of Chierchia (1998:350): 'Kinds are whatever your favorite worldview says that they are.' What matters presently is their function in a semantic system. Thinking compositionally, kinds are a certain counterpart of properties: they give us the collective individual comprised of those entities that bear the property. Along with the philosophical ambivalence I will also be

adopting from Chierchia some notation for describing this relationship: ‘If DOG (or, more precisely, $\lambda w \text{DOG}_w$) is the property of being a dog, then let $\cap \text{DOG}$ be the corresponding kind. Conversely, if d is the dog-kind, let $\cup d$ be the property DOG of being a dog. ‘ \cap ’ and ‘ \cup ’ are maps that allow us to get a kind from the corresponding property and vice versa.’ (Chierchia 1998:348). We can define these mapping functions as follows:

(154) For any property P and world/situation s ,

$$\cap P = \begin{cases} \lambda s \iota P_s, & \text{if } \lambda s \iota P_s \text{ is in } K \\ \text{undefined,} & \text{otherwise} \end{cases}$$

where P_s is the extension of P in s .

(155) Let d be a kind. Then, for any world/situation s ,

$$\cup d = \begin{cases} \lambda x [x \leq d_s], & \text{if } d_s \text{ is defined} \\ \lambda x [\text{FALSE}], & \text{otherwise} \end{cases}$$

where d_s is the plural individual that comprises all of the atomic members of the kind.

The availability of these functions allows for significant inherent flexibility in modulating the types of nominal elements; constraining their application in principled ways is one of the goals of a fully fleshed-out theory of kind reference. A core insight from Chierchia’s account is that, while null elements instantiating these type shifters do exist, their application is constrained by a blocking principle which prevents the use of a type shifting operation if the language in question has an overt determiner that would create the same meaning. It is important to keep this principle in mind going forward as the default answer to a question that would otherwise pop up frequently, namely, ‘why couldn’t we simply type-shift the noun at this point to achieve the unavailable reading?’ Unrestricted type-shifting would clearly preclude any analysis that tried to attribute restricted distributions of certain ambiguous meanings to the kind/property status of the noun at a given point in the derivation, as we are trying to develop here.

4.1.2 Kinds, nouns, and number

As that last sentence implies, core to the present proposal is that kind denotations and property denotations of nouns are associated with different regions of the extended nominal phrase in the syntactic derivation. While the particular implementation varies, the literature on kinds (Vergnaud and Zubizarreta, 1992; Longobardi, 1994; Chierchia, 1998; Larson, 1999; Zamparelli, 2000; Carlson, 2003) has been for some time circling around an intuitive generalization that, roughly stated, reference to kinds (or types) occurs lower in the structure of the (extended) noun phrase, and reference to properties (or objects or tokens) occurs higher. For example, Vergnaud & Zubizarreta suggest locating this contrast as a contrast between NP and DP:

- (156) Type/Token Correspondence Law (Vergnaud & Zubizarreta 1992)
- a. When a DP denotes, it denotes a token.
(e.g. the *water* refers to a specific, spatio-temporally located quantity of water)
 - b. When an NP denotes, it denotes a type.
(e.g. *water* refers to a kind of liquid [or solid or gas] substance)

As written, this law is merely descriptive - it doesn't formalize any explanation for why these properties are associated with these particular syntactic layers - but versions of it have been argued for repeatedly in subsequent syntactic work (see Carlson 2003 for an overview). On many of these perspectives including that of Vergnaud & Zubizarreta themselves, the law is attributable to the semantic function that the determiner performs: it converts the type denotation of the noun into a tokenized denotation. This is not a language-universal proposal, however - for example, this characterizes the English determiner, but not French, where the fact that the determiner performs an agreement function allows it to appear in a semantically inert expletive form (an option not available to English, because of the determiner's morphological inactivity, and therefore 'in English the definite determiner must be absent in a type-denoting expression' (Vergnaud & Zubizarreta 1992:635). I'll return to this particular contrast between English and Romance languages later in this chapter, but for now we can helpfully restrict our view to English without losing too much information. Leffel

(2014:35) summarizes the hypothesized contribution of the determiner in these accounts:

- (157) a. Nouns lexically denote kinds.
b. Determiners quantify instantiations of kinds.

This summary highlights a missing piece in these kind of analysis. Specifically, quantifying *over instantiations* is not at all the same semantic function as that of *instantiating*. For a determiner to do its regular quantificational work over instances - which we certainly want it to do - it must already be receiving as input an appropriately object-level or property denotation. This suggests that somewhere in between the base generation of the noun and the merging of the determiner, the denotation of the noun has been appropriately shifted from a kind to its instances in the world, a property of objects. We would prefer this to not be a function of conveniently invoked null type-shifting operators, if it is to be a regular and consistent aspect of the ontology of noun phrases that occurs cross-linguistically.

Following Leffel (2014), I am going to here adopt the position of Déprez (2001, 2005) that this conversion function is performed by grammatical number. On this account, the base denotation of a noun is always that of a kind. When grammatical number combines with the noun, it converts that kind denotation into the instances of the kind, thus making it appropriate input for the quantificational function of determiners. On this system cross-linguistic variation in noun meaning, and the ability of nouns to appear as bare arguments, is encoded in the differences in number morphology. Languages differ in whether or not they obligatorily have a functional projection for number in their syntax, as encoded by Déprez' 'Plural Parameter':

- (158) The Plural Parameter (Déprez, 2005): NumP must project and contain a semantic counter.
NP \rightarrow $\langle e \rangle$ = a kind
NumP \rightarrow $\langle e, t \rangle$ = the realization of a kind (object-level individuals or sub-kinds)

The notion of a 'semantic counter' is not going to feature much further in the discussion here, since the work it does will be picked up by other replacement bits of the machinery, but in short it represents a measure function that ensures the countability of the kind real-

ization (whether it's counting objects or sub-kinds). This parameter divides languages into +PLURAL and -PLURAL: languages with overt obligatory plural morphology and those that lack it. If a language is -PLURAL, such that the lack of plural morphology on a noun does not necessarily entail its singular interpretation, NumP neither needs to project nor contain such a counter, though it may, and that may in some cases be phonologically null. As a result, bare nouns in -PLURAL languages have a wider variety of options for their interpretation than bare nouns in +PLURAL languages. The general ban on singular bare arguments in +PLURAL languages like English and Italian follows from this.

This number-parameterized approach contrasts with the noun-parameterized approach of Chierchia (1998), where the relevant parameter is the basic denotation of nouns. While Déprez assumes that nouns cross-linguistically denote kinds and the difference arises in how and when that kind denotation is converted into an object denotation, Chierchia assumes that languages vary in whether their nouns basically denote kinds or object properties. As Déprez puts it, Chierchia assumes a 'flexible mapping between syntactic categories and semantic types' (Déprez 2005:865), while her proposal uses an inflexible syntax-semantics mapping and places the flexibility on the presence or absence of a bit of syntactic structure. Unsurprisingly, evaluating the full scope of this debate - which bears not only on the questions of modification of present interest but really on the fundamental ontology of language and linguistic reference - is well outside the capacity of this dissertation. Déprez argues for her position on the grounds that it makes meaningful reference to visible morphological differences between the varieties of languages (\pm PLURAL in her system; Class I/II in Chierchia's) while carving the joints in the same way, capturing the same divide in empirical behavior and therefore not losing out on coverage while gaining in explanatory power with the appeal to morphology. This is, in my view, a reasonable argument though one the data in this dissertation is not well-suited to adjudicate.

My perspective on the problem is: assuming that nouns have the basic function of kinds and that their conversion into objects always takes place in the syntax, such that there are

distinct domains for kind and object reference, offers enormous utility in explaining the sort of adjectival modification contrasts we are concerned with here. It allows us to not only give a structural account of these modification ambiguities but to do so in a way that reduces the contrast to a preexisting and richly independently motivated contrast (between kinds and properties). Both classes of language appear to show this contrast between lower and higher modifiers, which is significantly more natural to explain via appeal to the kind/object contrast on Déprez' worldview than Chierchia's. Obviously, one might not want to do that. You might prefer to pursue an explanation of this modificational contrast that makes no reference to kinds, and then preserve the bare noun parameter worldview for reasons (maybe quite good ones) external to anything about adjectives. But starting with the adjectival problem points towards the number-focused approach - as will be substantiated more in this chapter - and as a result of this argument, offering an alternate account of modificational ambiguity is potentially added to the list of responsibilities that a proponent of a fully-fledged bare noun-parameter worldview has to resolve. Put more simply: the ability to offer a rather uncomplicated and empirically motivated account of adjectival ambiguities in terms of number domains constitutes only one argument in favor of Déprez' view of kinds, but there is sufficient independent value in that worldview (from the morphological arguments) to make it worth adopting here in order to get the adjectival facts right.

4.1.3 Leffel (2014) on kind modification

It remains now to show exactly how this kind of assumption about nouns and kind conversion actually does capture the modification facts. For that, we will need a particular implementation of Déprez' theory, and here I will be laying out that of Leffel (2014), which is an account developed to explain the kind of adjectival ambiguity contrasts that we are concerned with here. The key difference is that Leffel's account is exemplified on other interpretive ambiguities, like the stage-/individual-level contrast, rather than the intersective/nonintersective contrast, but with some minor adjustments I want to say that the same approach will func-

tion in the latter case just as well. Leffel’s account is composed of a semantics for bare nouns, an analysis of grammatical number, and an analysis of adjectives, which will be introduced in this section.

Nouns. Uninflected nominal roots for Leffel are predicates of subkinds: they take in kinds and evaluate them for whether they are a subkind of the superordinate kind of the noun. So, a root like *dog* is true of those kinds which are subkinds of the DOG kind:

$$(159) \quad \llbracket \mathbf{dog} \rrbracket = \lambda k_{\langle s,e \rangle} : \mathbf{kind}(k) \cdot [k \leq (\lambda w [\iota x [* \mathbf{dog}_w(x)])]]$$

A few things in this denotation require elaboration. First, the symbol \leq denotes the taxonomic subkind relation (Krifka et al., 1995):

$$(160) \quad k_i \leq k_j \text{ iff } k_i \text{ is a more specific kind of } k_j$$

We can formalize this notion of ‘more specific’ as the subkind containing all the individuals of the superkind. Second, the symbol $*$ represents the pluralization operator from Link (1983), which gives the set of all sums that can be created out of the atoms of its input set (where \oplus is the sum operator):


$$(161) \quad \begin{array}{l} \text{a. } \llbracket P \rrbracket = \{a, b, c, \dots\} \\ \text{b. } \llbracket *P \rrbracket = \{a, b, c, a \oplus b, b \oplus c, a \oplus c, a \oplus b \oplus c, \dots\} \end{array}$$

Since kinds are functions from worlds into individuals, the dog kind itself is represented with this term, which appears internal to (159), which in Leffel’s notation we can call **DOG**:

$$(162) \quad \mathbf{DOG} = \lambda w [\iota x [* \mathbf{dog}_w(x)]]$$

This is a function that takes a world as input, and returns the unique individual that is comprised of the pluralized set of all sums of those individuals that are dogs in that world. Therefore, the entire denotation of *dog* in (159) is a function that takes in a kind k and evaluates true iff $k \leq \mathbf{DOG}$, or k is a subkind of the dog kind. This is the basic denotation of a nominal root prior to inflection.

Number. As Déprez argued, inflection with number is what converts a noun from this kind-level denotation to a property of objects. Leffel departs from Déprez in assuming that grammatical number is hosted not in a distinct NumP, but rather in the specifier position of NP. Nothing in particular rides on this choice for us here - this seems like a decision that should be morphosyntactically motivated, though Leffel doesn't offer such a motivation - but for consistency I will for now continue to represent Leffel's choice, a phi-feature bundle in the Spec,NP position.



This phi projection hosts the singular or plural feature, which in those languages that are +PLURAL for Déprez do the work of transforming the kind of denotation in (159) to an individual property. Both singular and plural morphemes do this equally, but the singular performs a double duty: after generating a set of individuals, it performs an additional atomizing function, filtering out non-atomic individuals from the set (Sauerland, 2003). That atomizing operation is represented by \circ :

(164) $\circ P = \{x \in P \mid x \text{ is atomic}\}$

Singular and plural are then defined as follows¹⁴:

(165)

- a. $[\mathbf{SG}] = \lambda P_{\langle\langle s,e \rangle, t \rangle} [\lambda w [\circ [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z) \wedge x \sqsubseteq z_w]]]]]$
- b. $[\mathbf{PL}] = \lambda P_{\langle\langle s,e \rangle, t \rangle} [\lambda w [* [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z) \wedge x \sqsubseteq z_w]]]]]$
(where \sqsubseteq is the part-of relation on \oplus -sum formation)

Denotations for a singular and plural noun, then, will be as:

(166) a. $[\mathbf{dog}] = \lambda w [\circ [\lambda x_e [\exists z_{\langle s,e \rangle} [z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]]]$

¹⁴While Leffel defines these as independent and fully-formed morphemes, it is worth noting that this still captures the insight being picked up from Sauerland (2003) that the plural is the default, unmarked element and the singular is built up on top of that, because the \circ function is a stronger version of the $*$ function and therefore SG could be thought of as PL plus an additional restriction. Those restrictions could be broken apart in distinct morphemes to equivalent semantic effect as this implementation.

- = e.g., {fido, rover, spot,...}
- b. $\llbracket \mathbf{dogs} \rrbracket = \lambda w[*[\lambda x_e[\exists z_{\langle s,e \rangle}[z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]]$
 = e.g., {fido, rover, spot, fido \oplus rover, rover \oplus spot,
 fido \oplus spot, fido \oplus rover \oplus spot...}

In both cases what we have is an intensionalized property that applies to individuals that instantiate any subkind of the dog kind, differing on whether or not that property applies to singular and plural or only singular individuals.

The assumption that nominal roots always denote kinds, and that grammatical number always converts that kind denotation to a property, does run into (at least) one obvious empirical obstacle: sometimes, nouns do actually denote kinds, even by the time that they have combined with a determiner. If number always eliminates the kind semantics, actually having surface kind readings should not be possible. And yet we do see them, of course, in both bare plurals and singular definites.

- (167) a. Dogs have four legs.
 b. The dog and the wolf evolved from a common ancestor.

Leffel addresses these two cases independently. For kind readings of bare plurals as in (167a), all that we need is Chierchia's \cap operator, which can recreate a kind out of the property.

- (168) a. $\llbracket \mathbf{PL(dog)} \rrbracket = \lambda w[*[\lambda x_e[\exists z_{\langle s,e \rangle}[z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]]$
 b. $\llbracket \cap(\mathbf{PL(dog)}) \rrbracket = \lambda s[\iota y[*[\lambda x_e[\exists z_{\langle s,e \rangle}[z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]]]$

Specifically, the application of \cap to a plural property selects the unique maximal member of that property to form the kind from, which is going to be the maximal \oplus -sum over its members, the sum of all members. The fact that we do not see bare singular arguments with kind readings (or at all) in English follows from the way that \cap is defined in (154): since it is defined only over contexts where there is a unique member of the extension of the property in the world, and singular nouns contain only atoms, there is no such unique maximal member to satisfy the presupposition of ι , and therefore the attempt to recreate a kind out of a singular property will cause a presupposition failure.

On this view, then, the distinction between plural and taxonomic kinds (Dayal, 2004) is captured derivationally: the basic denotation of nouns is always that of a taxonomic kind, picking out the set of taxonomic subkinds of the named kind, and reference to plural kinds as in (167a) is always achieved via type-shifting from the taxonomic denotation after the application of semantic plurality.

Accounting for singular definite kind readings is, as a result, slightly more difficult. What Leffel argues is that we can avoid this problem by assuming that ‘singular’ kind terms are, in fact, not semantically singular but rather lack semantic number altogether, like proper names (supported by Carlson 1977’s suggestion that these singular kind-terms are acting as ‘proper names of kinds’). As a result, the basic subkind predicate denotation is passed up to the level of the definite determiner directly. Of course, as identified earlier, this will cause an immediate problem with determiners defined only as quantifiers over individuals. To make this approach work, Leffel is forced to generalize the type of the determiner:

$$(169) \quad \textbf{Generalized polymorphic semantics for } \textit{the} \\ \llbracket \textbf{the} \rrbracket = \lambda P_{\langle \sigma, t \rangle} [\iota x_{\sigma} [P(x)]]$$

With this in place, the derivation of the singular kind reading is straightforward. Interpreting the definite as a quantifier over kinds ($\sigma = \langle s, e \rangle$), it will pick out the unique maximal kind from the noun in this world:

$$(170) \quad \llbracket \textbf{the} \rrbracket (\llbracket \textbf{dog} \rrbracket) = \iota k [k \leq \textbf{DOG}]$$


Interpreting this also requires generalizing the ι operator, so that it can select the maximal sum from the extension of the subkind predicate, which will end up being the kind **DOG** itself. Therefore, we get the singular kind reading desired for sentences like (167b).

The moves that Leffel is required to make to get surface kind readings are undeniably marks against the account, in different ways. Assuming a generalized polymorphic *the* is no small thing and has potentially wide consequences for DP semantics if unconstrained. The move to require reconversion of the kind-derived property back into a property-derived kind for bare plurals is less severe in its potential spillover effects, but is certainly less elegant


than the noun-parameter account that allows the basic kind reading to surface in those cases without redundant compositional steps. While these kinds of downsides seem at present to be reasonable tradeoffs for the mileage that we will gain in tackling modification problems, it is worth flagging here that these are currently the theoretical door charges we're paying to access that upside of the account, which I'll turn to now.

Adjectives. The goal of this account is to capture the difference between direct and indirect modification readings of adjectives as a result of the adjective alternately modifying either the kind-level denotation of the noun or the object-level denotation. With the architecture described above, where kind denotations are converted into object denotations at the introduction of the number morpheme, this amounts to saying that the two domains for adjectival modification are demarcated by number. An adjective which merges below number will be a direct modifier, substantively modifying the noun-as-kind, while an adjective that merges above number will be an indirect modifier, modifying noun-as-property or noun-as-set-of-objects. Syntactically, Leffel implements this by proposing that adjectives may enter the derivation either via head-adjunction (following Sadler and Arnold 1994) or phrasal adjunction:

(171) a. Head Adjunction:



b. Phrasal Adjunction:



The directionality of these adjunction operations is a parameter of cross-linguistic variation; for English the assumption is that head-adjunction is only leftward (as above) but phrasal adjunction can be leftward (as above) or rightward.¹⁵

If adjectives can modify either the kind-level denotation of a noun or its object-level property denotation, then it is necessary that we have versions of the adjective that are of appropriate type to handle both of these cases. Direct (or inner) modifiers are therefore predicates of kinds, while indirect (or outer) modifiers are intensional predicates of individuals:

- (172) a. $[[\mathbf{adjective}]_{\text{INNER}}] = \lambda k_{\langle s, e \rangle} [k \leq \text{ADJECTIVE}]$
 b. $[[\mathbf{adjective}]_{\text{OUTER}}] = \lambda w [\lambda x [\text{adjective}_w(x)]]$

This account of kind modification takes significantly inspiration from McNally and Boleda (2004), who argue that the class of relational adjectives should be treated as predicates of kinds which intersect with nouns that also denote predicates of kinds. While plenty of interpretive contrasts in adjectives have been tied to this syntactic distinction between inner and outer modification (Svenonius, 2008; Alexiadou et al., 2007), Leffel’s exposition focuses on the contrast between stage-level and individual-level modification, as exemplified by the classic Capella sentences:

- (173) The visible stars include Capella.
 a. ‘The stars that are currently visible include Capella’ (stage-level)

¹⁵For present purposes, I will be representing the different attachment heights with this choice between head and phrasal adjunction as well, though no part of my particular proposal is going to turn on this - as long as the adjective is Merged above or below number, the syntactic choice of adjunction type vs. something else shouldn’t matter.

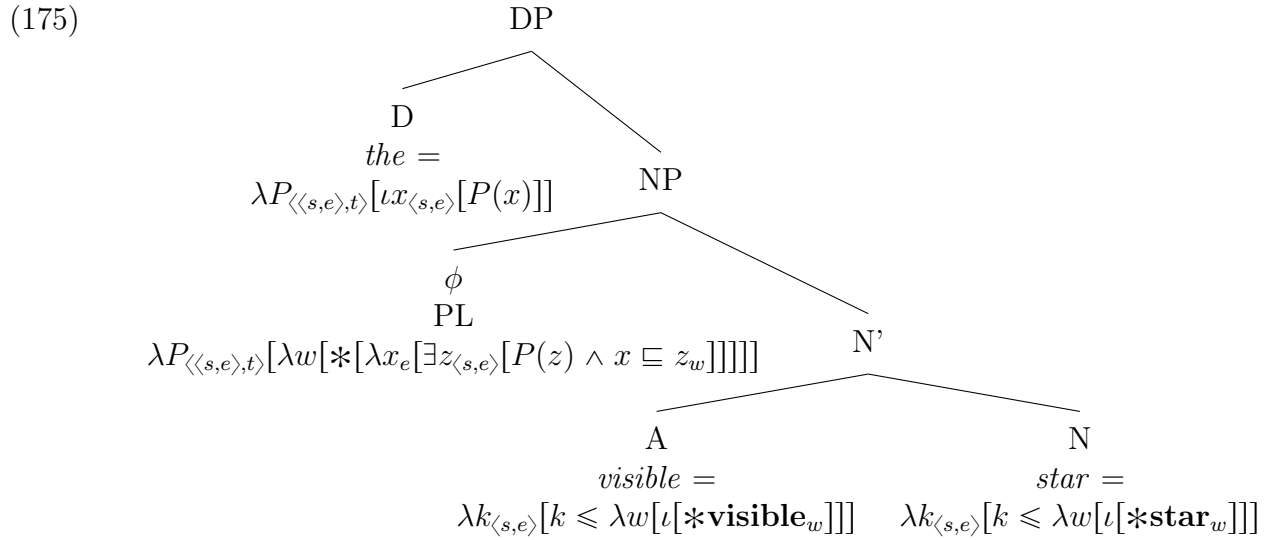
- b. ‘The stars which are characteristically visible include Capella’ (individual-level)

This ambiguity suggests that *visible* can be either an inner/direct or outer/indirect modifier, and so has both potential denotations for composing with kinds or properties:

- (174) a. $[[\mathbf{visible}]_{\text{INNER}}] = \lambda k_{\langle s,e \rangle} [k \leq \mathbf{VISIBLE}]$
 $= \lambda k_{\langle s,e \rangle} [k \leq \lambda w [\iota [* \mathbf{visible}_w]]]$
 b. $[[\mathbf{visible}]_{\text{OUTER}}] = \lambda w [\lambda x [\text{visible}_w(x)]]$

The inner-modifying visible is a predicate over kinds and holds true of any kind which is a taxonomic subkind of the **VISIBLE** kind. The outer-modifying visible is an intensional predicate over individuals, and holds true of any individual that is visible in that world.

We’re now equipped to show the derivation of both readings of (173) in this system. The individual-level reading, again, is the result of applying the inner version of *visible* in (174a) to the subkind-predicate denotation of the noun, *stars*, prior to its composition with number. The derivation of the DP is thus:



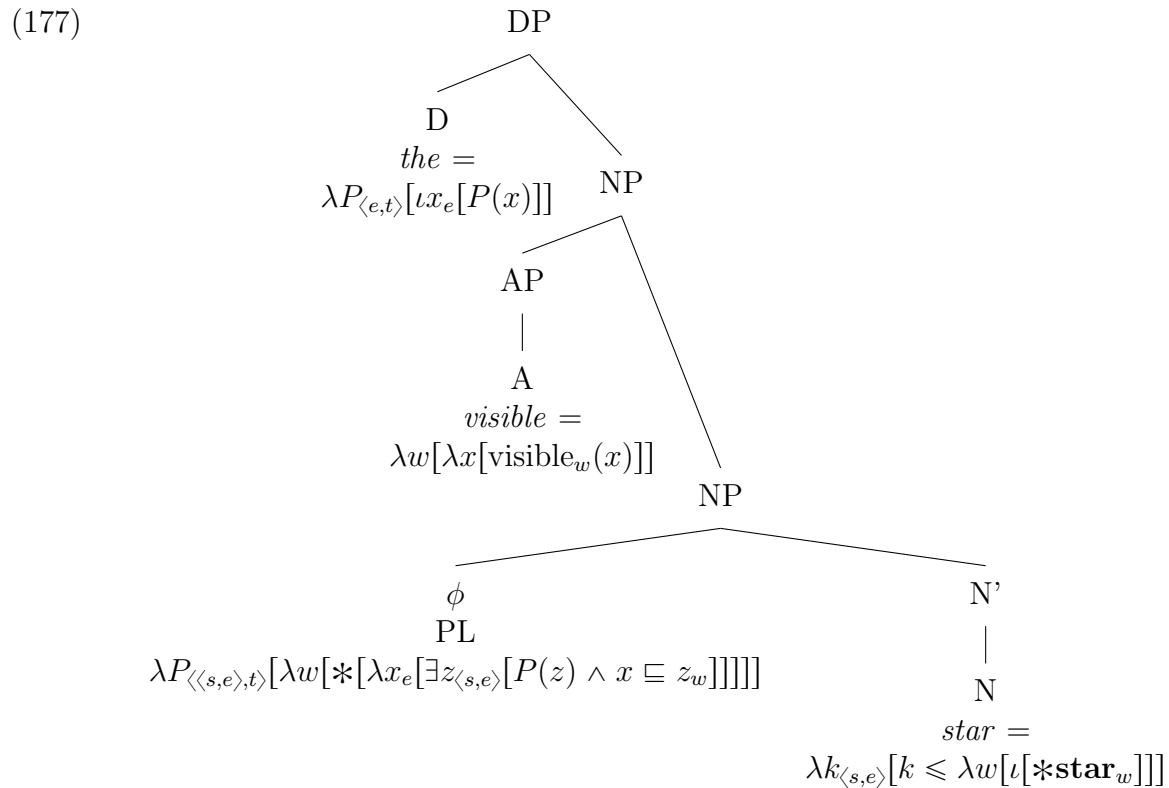
Intersection of the A and N denotations, and then application up the rest of the tree, gives as the eventual result:

- (176) Individual-level reading of *the visible stars*
 $\iota x_e [\exists k_{\langle s,e \rangle} [k \leq (\lambda w [\iota [* \mathbf{visible}_w]]) \wedge k \leq (\lambda w [\iota [* \mathbf{star}_w]]) \wedge x \sqsubseteq_{w'} k]]$

This denotes the unique maximal individual (in this case, a plural one) which is a member

of the kind k , where k is a subkind of **VISIBLE** and a subkind of **STAR**. The individual-level reading here is specifically attributable to the fact that the individual must belong to the visible *kind*, with the understanding that kind membership is not a temporary fact evaluated at a particular instance but rather describes some core characteristic of the individual.

The stage-level reading is derived by instead using the outer modifying version of the adjective, and merging it through phrasal adjunction above ϕ :



In this derivation, PL applies to the kind-level denotation of N and converts it to an intensional object-level property which can compose with the AP via intersection. The whole NP then combines with a suitably typed determiner, yielding:

(178) Stage-level reading of *the visible stars*
 $\iota x_e [\mathbf{visible}_{w'}(x) \wedge \exists k_{\langle s,e \rangle} [k \leq (\lambda w [\iota [* \mathbf{star}_w]]) \wedge x \sqsubseteq_{w'} k]]$

This gives the unique maximal (again, plural) individual who is both visible and a member of a kind that is a subkind of **STAR**. Because the actual application of the ‘visible’ term is to an individual, not to the kind, this will result in the reading of a currently visible individual

and impose no requirement as to the characteristic, essential visibility of the star (though it must be characteristically and essentially a star).

Evaluating the account. Leffel’s story for the individual/stage-level contrast is attractive because of its reliance on preexisting and independently motivated semantic assumptions about the semantics of number, rather than innovating a novel syntactic distinction specifically to explain modification ambiguity, in the way that Cinque (2010) did with small dP. While the assumptions required about number and kind reference to make the baseline theory work are not without some stipulative consequence, the technical derivation of ambiguity is a functional and intuitive use of preexisting compositional machinery. We have already covered the objections from the perspective of a theory of kind reference, namely how the assumption of universal basic kind denotations runs into mechanical issues in deriving surface kind readings; what I want to focus on here is potential objections from the point of view of a theory of modification and ambiguity. I do want to argue in favor of adopting basically a version of this approach for the (non)intersective ambiguity, but there are some issues to overcome first.

From a theory-building standpoint, Leffel’s account highlights one of the central issues with appealing to structural or domain distinctions to explain ambiguity. At least as I see it, one of the primary appeals of a structural account is the hypothetical ability to avoid complicating the lexicon. While polysemy or doublet theories require duplicate lexical entries for adjectives, a structural account that distinguishes meaning based on where the adjective is placed in the syntax relative to the noun should be able to capture that meaning difference with a single semantics for the adjective, because the distinguishing burden is shifted to syntax. That is, arguably, the whole theoretical point of making an appeal to domains. But avoiding lexical duplication is not so easy. By claiming that the core difference in meaning is about the semantic value of the noun below or above Num, we are also forced into making the adjective be able to take both versions of the noun as its argument. This is why Leffel needs

to define two versions of the adjective as in (172). Consequently, we end up with redundant theoretical machinery: a complicated syntax alongside a complicated lexicon, where the goal of complicating the syntax was ostensibly to enable us to simplify the lexicon. Leffel does explicitly acknowledge that it would be nice to have a single version and derive one from the other within the AP, but without any obvious way to do so, leaves that as merely an open possibility.

From an empirical standpoint, what this account lacks is independent motivation *for* treating the individual/stage-level contrast as a kind/object-level contrast. Doing so results in a theoretically attractive simplification of the problem, which of course can count as motivation in itself, but not empirical motivation. The fact that individual-level modification can be reasonably described as kind-level modification (and stage- for object-) isn't itself a reason to do so. Frankly, Leffel offers no syntactic evidence for differences in adjective attachment height being separated by number (the directionality of the adjunction rules needing to be stipulated, after all, to get the word order facts at all), nor semantic evidence for assimilating the two contrasts besides that it is elegant to do so.

This is not to say that either of these problems are unsolvable, however. And that is precisely what the second half of this chapter will be devoted to: providing novel empirical motivation in favor of reducing both the (non)intersectivity ambiguity and the individual/stage ambiguity to the kind/object contrast, revising Leffel's theory of adjectives to simplify the lexicon in the desired way, and applying this revised theory to additional languages. Some of these motivations will come from further application of data already introduced in earlier chapters, and some from novel paradigms. I'll start with some evidence of the latter kind.

4.2 Kind- and object-modification in Turkish

One source of evidence for reducing the adjectival ambiguity to the number-sensitive kind/object distinction comes from the interaction of modification and number in Turkish. This interaction provides a potential diagnostic test for whether a particular instance of modifi-

cation is kind- or object-level. This section introduces the relevant facts about Turkish, their analysis from Sağ (2021), and then applies this diagnostic to the individual/stage and (non)intersective contrasts.

4.2.1 Number neutrality and modification

Turkish is an optional classifier language which, like English, has nouns that can be either marked morphologically with an overt plural morpheme, or unmarked (that is, there is no specific singular morpheme). Data in this section comes from Sağ (2021), unless otherwise stated.

- (179) a. *kitap* ‘book’
b. *kitap-lar* ‘books’

However, Turkish differs from English in the interpretive possibilities for its singular (unmarked) nouns. While English unmarked nouns are unambiguously singular, Turkish nouns which lack plural marking appear more flexible in their interpretation, sometimes allowing a plural interpretation despite the absence of plural morphology. This is not universally true, but plural interpretations of unmarked nouns arise in three specific cases. The first is non-case-marked direct object positions:

- (180) Ali kitap oku-du
Ali book read-PST
‘Ali read one or more books’

This interpretation is not strictly plural or singular, but compatible with both number options. Call this the ‘number neutral’ reading. In direct object position, number neutrality depends on not being marked for case. Case-marked direct objects without overt plurality are instead obligatorily read with a singular interpretation:

- (181) Ali kitab-ı oku-du
Ali book-ACC read-PST
‘Ali read **the** book’

In general, any subject position or case-marked object-position non-plural noun in an

argument position is obligatorily singular:

- (182) a. Çocuk ev-e koş-tu
child home-DAT run-PST
'**The** child ran home'
b. Ali çocuğ-a kitab-ı ver-di
Ali child-DAT book-ACC give-PST
'Ali gave the book to **the** child'

The second position that displays number neutrality is the existential copular construction. Turkish has a dedicated existential copula, *var*, and unmarked nouns which appear immediately preceding this copula can be number neutral.

- (183) Oda-da fare var
room-LOC mouse exist
'There are one or more mice in the room'

The third position is the predicate position with a null copula, which allows both singular and plural subjects with an unmarked predicative noun:

- (184) a. Ali çocuk
Ali child
'Ali is a child'
b. Ali ve Merve çocuk
Ali and Merve child
'Ali and Merve are children'

Recall from the previous discussion that a standard semantic treatment of number in English and English-like languages considers the plural morpheme to be the semantically unmarked one, with a pluralized set including both atomic and non-atomic (summed) individuals, while the singular morpheme is semantically marked and stricter, filtering out non-atomic individuals and leaving only singular atoms. From this point of view, the number neutrality of 'singular' nouns in Turkish is extremely curious. Applying the same perspective to the Turkish number morphemes would predict that overtly marked plurals might allow this kind of flexibility under certain pragmatic circumstances, but that the possibility of plural reference should be precluded by the semantics of the singular morpheme. This problem has led a number of theorists (Bliss, 2004; Bale et al., 2010; Görgülü, 2012) to argue

that Turkish in fact presents the inverse situation as English: singulars are the semantically unmarked, number-neutral sets, and the plural morpheme enforces a strict anti-atomic filter, allowing only sums:

- (185) a. $\llbracket \text{kitab} \rrbracket = \{a, b, c, a \oplus b, b \oplus c, a \oplus c, a \oplus b \oplus c\}$
 b. $\llbracket \text{kitab+PL} \rrbracket = \{a \oplus b, b \oplus c, a \oplus c, a \oplus b \oplus c\}$

This account aligns morphological and semantic markedness: the singular, which is morphologically unmarked in both English and Turkish, would be additionally semantically unmarked in Turkish, and the morphologically marked plural is also semantically marked. This could be considered an advantage of this approach, though it appears to be typologically rare, with a larger majority of languages exhibiting the English-like markedness asymmetry Sauerland et al. (2005)

Sağ (2019, 2021) instead argues for an account that brings Turkish in line with the typological majority, with number-neutral bare plurals (semantically unmarked, morphologically marked) and strictly singular bare singulars (semantically marked, morphologically unmarked). The starting point for Sağ’s argument is that the exclusively plural reading of bare plurals in Turkish, like in English, appears to be an implicature rather than a necessary condition of their semantics. While the exclusive plural reading is the default interpretation of bare plurals (for both Turkish and the English translation):

- (186) Çocuk-lar sokak-ta top oynu-yor
 child-PL street-LOC ball play-PROG
 ‘Children are playing ball on the street’ → ‘More than one child’

...this effect disappears in downward-entailing contexts and questions:

- (187) a. Çocuk-lar sokak-ta top oynu-mı-yor
 child-PL street-LOC ball play-NEG-PROG
 ‘Children aren’t playing ball on the street’ → ‘Not even one child’
 b. Orman-da ayı-lar-la karşılaş-tı-nız mı?
 forest-LOC bear-PL-COM come.across-PST-2PL Q
 ‘Did you come across bears in the forest?’ → ‘Even one bear’

Assuming that plurals are semantically unmarked naturally predicts these readings in

(187); the challenge is to explain why upward-entailing contexts have exclusively plural readings, which Sağ follows Spector (2007) and Zweig (2009) in adopting a scalar implicature account for (in short, the singular alternative is stronger in positive but not negative contexts, and so not using it suggests plurality in the former but not the latter).

The other obstacle that Sağ (2021) argues a strict-plural, neutral-singular account faces concerns certain patterns with modification interacting with number neutrality in the bare singular forms. Specifically, certain modifiers when applied to bare singulars preserve the option of number neutrality, while others either force a strict singular interpretation or result in ungrammaticality without the overt plural marker. Beginning with the non-case-marked direct object position:

- (188) a. Ali, ev-e geldikten sonra, dini / bilimsel kitap oku-du
 Ali home-DAT having.come after religious / scientific book read-PST
 ‘After he came home, Ali read one or more religious/scientific books’
 b. *Ali, ev-e geldikten sonra, eski / küçük kitap oku-du
 Ali home-DAT having.come after old / small book read-PST
 Intended: ‘After he came home, Ali read one or more old/small books’

Adjectives like *dini* ‘religious’ or *bilimsel* ‘scientific’ can modify the bare singular *kitap* ‘book’ while preserving its number-neutral reading, as in (188a). On the other hand, adjectives like *eski* ‘old’ or *kucuk* ‘small’ cannot - trying to use the bare singular with these adjectives in this position results in ungrammaticality. Sağ summarizes the relevant contrast as follows: ‘What seems to be common to the adjectives that are acceptable is that they establish a sub-type of the noun that they modify; they denote classificatory properties in a sense. While *religious/scientific* books can be considered sub-types of books, it is harder to establish this relation with *worn-out/small* books. The second set of adjectives simply define physical properties of the book/books that were read’ (Sağ 2021:5). In that sense, then, the adjectives that are allowed in this context are those that can be reasonably construed as kind-level modifiers, like McNally & Boleda’s (2004) relational adjectives; those that aren’t permitted appear to be object-level modifiers.

The same contrast between kind- and object-level modifiers appears in the other bare

singular positions, though to different effect. In the existential copular position, kind-level modifiers again are grammatical and preserve the number-neutral reading, and object-level modifiers - while grammatical, in contrast to the direct object position - restrict the possible readings of the noun to strictly singular, removing the number-neutrality:

- (189) a. Kutu-da dini / bilimsel kitap var
 box-LOC religious / scientific book exist
 ‘This box has the religious/scientific book’
 ‘There are one or more religious/scientific books in this box’
 b. Kutu-da eski / küçük kitap var
 box-LOC old / small book exist
 ‘This box has the old/small book’
 #‘There are one or more old/small books in this box’

The same effect holds in predicate position. Kind-level modifiers which define a relevant subtype of the bare singular noun are compatible with both singular and plural subjects; object-level modifiers are only compatible with singular subjects, not plural.

- (190) a. Ali (ve Mehmet) pratisyen doktor
 Ali and Mehmet practitioner doctor
 ‘Ali is a practitioner doctor’
 ‘Ali and Mehmet are practitioner doctors’
 b. Ali (*ve Mehmet) yakışikli doktor
 Ali and Mehmet handsome doctor
 ‘Ali is a handsome doctor’
 *‘Ali and Mehmet are handsome doctors’

The resulting generalization is that only kind-level modifiers are compatible with the plural reading of unmarked bare nouns in Turkish. Object-level modifiers eliminate the possibility of the plural reading, either through rendering the whole construction ungrammatical (as with the non-case-marked direct object position) or through solely disallowing that reading (as with existential and predicate positions).

4.2.2 Singular kind terms

The interaction of number neutrality with kind/object-level modification leads Sağ to propose that bare singulars in Turkish denote singular kinds, or at least have the option to.

This picks up on the notion of singular kind terms introduced by Dayal (1992, 2004), who contrasts singular from plural kinds in addition to the contrast between singular and plural sets of objects. Singular kind terms are conceptually plural, representing an individual that is comprised of more than one member, but grammatically atomic. They are not a basic notion, but rather compositionally derived. A singular kind term results from the application of a definite determiner (whether overt or covert, depending on the language) with a noun that denotes a property of taxonomic individuals (= subkinds). So, to derive the singular kind reading necessary for a sentence like:

(191) The African lion is extinct.

We can model *lion* as a predicate of taxonomic individuals, or subkinds, which intersects with *African* to yield the singleton set containing only the African lion subkind of lion, {**AFRICAN LION**}. The definite determiner then picks out the only member of that set, trivially satisfying uniqueness and maximality. For a sentence like:

(192) The lion is extinct.

We have two options, given to us by either Saĝ or Leffel's implementations of Dayal's proposal. On the Saĝ-Dayal account, context shifts the domain of quantification to the level of kind meanings that contains **LION** along with other kinds of similar taxonomic specificity, e.g., **WHALE**, **DOG**, etc. Then the denotation of *lion*, at this level, is the singleton set {**LION**}. On Leffel's revision, there is no role for this shift in the domain of quantification, and *lion* always denotes a predicate of subkinds of lions. The difference is that the determiner is redefined to have its polymorphic semantics in (169) and a generalized notion of maximality allows for **LION** to be retrieved as the maximal subkind of itself, analogous to how the complete sum of all members in a plural set is the maximal element of that set. The choice shouldn't end up consequential. In either case, the basic meaning of a bare singular noun like *lion* ends up as a predicate over subkinds, and the singular kind reading in definite contexts is attained compositionally.

Turkish, in contrast to English, allows bare singulars to carry the definite singular kind interpretation without an overt definite article:

- (193) a. Dinozor 250 milyon yıl önce evrimleş-miş-tir
dinosaur 250 million years ago evolve-PERF-GEN
‘The dinosaur evolved 250 million years ago.’
b. *Dinosaur evolved 250 million years ago.

Sağ attributes this to the lexical absence of a definite determiner in Turkish. Recall Chierchia’s (1998) blocking principle, which prevents the application of a covert type-shifter if the language has an overt determiner that could result in the same meaning. English *the* therefore blocks the application of a covert iota operator that would be required to get singular kind readings with bare singulars, but that covert operator is perfectly fine in Turkish due to the lack of an overt competitor.

Turkish bare singulars are not exclusively read as kinds, however. They also display standard definite singular interpretations over objects, made unambiguous in episodic contexts where the kind-reading is unavailable:

- (194) Köpek bugün havla-mı-yor
dog today bark-NEG-PROG
‘The dog isn’t barking today’
#‘Dogs aren’t barking today’

Sağ argues that this ambiguity is a fundamental one: ‘Turkish bare singulars are ambiguous in denoting an atomic property of object-level entities and an atomic property of taxonomic/kind entities’ (Sağ 2021:21). In both cases, a covert iota type-shifter is applied, deriving either the object-level singular definite interpretation or the singular definite kind interpretation.¹⁶ And it is this ambiguity that gives us the explanation for the perceived

¹⁶One could certainly try to pursue a further unification of Leffel and Sağ’s accounts by rejecting the basic status of this ambiguity and attempting to always derive the object-level property from a basic taxonomic meaning, like Leffel argues for English and Romance. I’m sympathetic to this project, and suspect that it’s both possible and desirable; the likely line would be similar to Leffel’s analysis of singular definite kind readings in English, with singular kind readings in Turkish being the result of an absence of grammatical number entirely and singular object readings from a null singular morpheme, whose denotation would need to be adjusted to account for the impure atomic nature of its singular kind input. However, for present purposes it will not change anything about the analysis of modification whether we do this or not, so I will leave it for future consideration.

number-neutrality of Turkish bare singulars: what looks to be on the surface an ambiguity between singular and plural interpretations is actually an ambiguity between object and kind interpretations. While both are grammatically singular, the conceptual plurality of the singular kind interpretation is what leads to the possible plural truth conditions of bare singulars.

This ambiguity also explains the observed sensitivity of the number neutrality to modification. Recall the pattern: only kind-level modifiers, which serve a classificatory function and describe an established subtype of the noun they modify, are allowed with the plural reading of bare singulars. Attempting to modify a bare singular with an object-level modifier results in either ungrammaticality or the loss of the plural reading. On the hypothesis that the plural reading is in fact the result of the noun being interpreted as a kind, this pattern follows naturally. Object-level modifiers would be simply incompatible with the taxonomic predicate, triggering a type-clash, and so the presence of an object-level modifier necessarily indicates that the bare singular is in its object-level form.

Sağ develops a detailed analysis of how the surface plural readings are derived compositionally from the taxonomic form of the bare singular, and why it is these three particular contexts - non-case-marked direct objects, existential copular arguments, and predicates - that allow such readings. What is relevant here, however, is the conclusion regarding modification. Given this analysis of the modification contrast, Turkish number neutrality can end up functioning as a diagnostic for distinguishing kind-level from object-level modification. If an adjective preserves the plural readings of bare singulars in these positions, it must be of the appropriate type to take a kind-level argument; if an adjective precludes the plural reading, we can conclude that it is an object-level modifier. This allows Turkish to serve as a testing ground for Leffel's hypothesis that individual-level adjectives are kind-modifiers, and for the current proposed extension that nonintersective adjectives are, as well.

4.2.3 Applying the diagnostic

Leffel claims that the individual-level interpretations of ambiguous adjectives like *visible* are, in fact, the result of modifying a kind-level denotation of the noun, while stage-level interpretations are the result of modifying an object-level denotation. If this is true¹⁷, we would expect the meaning of such an ambiguous adjective to change depending on whether a Turkish bare singular noun is interpreted as singular or plural; or, conversely, we should expect the intended meaning of an ambiguous adjective to determine whether a Turkish bare singular noun can be given a plural interpretation.¹⁸

Beginning with the predicate position, we can see that Turkish *görünür* ‘visible’ displays the same ambiguity as its English counterpart:

- (195) Capella görünür yıldız
Capella visible star
‘Capella is a visible star’
- a. ‘Capella is one of the stars that are characteristically visible’ (individual-level)
 - b. ‘Capella is a star which is currently visible’ (stage-level)

Predicating the bare singular *görünür yıldız* ‘visible star’ of a plural subject, however, changes this:

- (196) Capella ve Vega görünür yıldız
Capella and Vega visible star
‘Capella and Vega are visible stars’
- a. ‘Capella and Vega are stars that are characteristically visible’ (individual-level)
 - b. ?*‘Capella and Vega are stars which are currently visible’ (stage-level)

With the plural subject, the stage-level reading becomes significantly degraded. This parallels the way in which plural subjects become unavailable for unambiguously object-

¹⁷Of course, this only means ‘if Leffel’s claim is true cross-linguistically’. It’s entirely possible that the surface ambiguity that appears across languages is attributable to distinct compositional (or noncompositional) processes in different languages, and the failure of this diagnostic in Turkish would not conclusively disprove Leffel’s claim about English or Romance. I assume in saying this that pursuing a unified cross-linguistic account of the phenomenon is, however, the goal absent strong specific evidence to the contrary. And the success of the diagnostic would, again, not be conclusive proof of the kind/object theory of English and Romance, but it would be a strong push towards such a theory.

¹⁸The previous note about all data coming from Sağ (2021) ends here; all the data in this section is novel. Thanks to Yağmur Sağ and Hande Sevgi for providing judgments and helpful commentary.

level modifiers. Further illustrating the effect, placing the sentence in an episodic context with an explicit temporal anchor forces the stage-level reading, and that temporal anchor is incompatible or significantly degraded alongside a plural subject:

- (197) Bu gece Capella görünür yıldız
 tonight Capella visible star
 ‘Tonight, Capella is a visible star’
- a. #‘Tonight, Capella is one of the stars that are characteristically visible’ (individual-level)
 b. ‘Tonight, Capella is a star which is currently visible’ (stage-level)
- (198) ?*Bu gece Capella ve Vega görünür yıldız
 tonight Capella and Vega visible star
 Intended: ‘Tonight, Capella and Vega are visible stars’

Expressing the plural interpretation of these kinds of sentences requires the addition of an overt plural marker. This parallels the obligatory interpretation of bare singulars as object-level singular definites in episodic contexts, as we saw in (194), which disallowed the kind reading.

The same effect occurs in the existential copular position:

- (199) gökyüzün-de görünür yıldız var
 sky-LOC visible star exists
 ‘There is visible star(s) in the sky’
- a. ‘There is a characteristically visible star in the sky’ (singular + individual-level)
 b. ‘There are characteristically visible stars in the sky’ (plural + individual-level)
 c. ‘There is a currently visible star in the sky’ (singular + stage-level)
 d. #‘There are currently visible stars in the sky’ (plural + stage-level)

The individual-level interpretation of the adjective preserves number neutrality and is compatible with both singular and plural interpretations of the noun, while the stage-level interpretation of the adjective is compatible with only a singular interpretation of the noun. This again mirrors the pattern in (189), where the kind-level modifier preserves number neutrality and the object-level modifier forces singularity.

Finally, the non-case-marked direct object position:

- (200) Ali görünür yıldız gördü
 Ali visible star see.PST
 ‘Ali saw visible star(s)’
- a. ‘Ali saw a characteristically visible star’ (singular + individual-level)
 - b. ‘Ali saw characteristically visible stars’ (plural + individual-level)
 - c. #‘Ali saw a currently visible star’ (singular + stage-level)
 - d. #‘Ali saw currently visible stars’ (plural + stage-level)

Here, the stage-level interpretation is entirely unavailable regardless of the number interpretation, much like the the object-level modifiers make (188b) ungrammatical. The individual-level interpretation of the adjective, like the kind-level modifiers in (188a), retain number-neutrality.

The consistent parallels between the clear kind-modifying adjectives and individual-level readings of ambiguous adjectives on one hand, and clear object-modifying adjectives and stage-level readings of ambiguous adjectives on the other hand, are striking and lend substantial support to the idea that this kind of ambiguity can be reduced to a contrast in kind/object-modification. In just the same way, it appears that the stage-level interpretations of *visible* are incompatible with the singular kind reading (= the surface plural reading) of bare singulars, suggesting an underlying type-clash such that stage-level *visible* can only modify predicates of objects, while individual-level *visible* is of the appropriate type to modify subkind predicates.

The next question is whether the same parallels extend to the (non)intersective ambiguity. We can apply the same diagnostics in just the same way, looking at the adjectives *iyi* ‘good’ and *yetenekli* ‘skilled’. Both of these adjectives do display the (non)intersective ambiguity, though much like English, the intersective reading is by default somewhat dispreferred:

- (201) Ali iyi hırsız
 Ali good thief
 ‘Ali is a good thief’
- a. ‘Ali is good at thievery’ (nonintersective)
 - b. ?‘Ali is a good person’ (intersective)

- (202) Oya yetenekli dansçı
 Oya skilled dancer
 ‘Oya is a skilled dancer’
- a. ‘Oya is skilled at dancing’ (nonintersective)
 - b. ?‘Oya is generally skillful’ (intersective)

With plural subjects, however, there is a clear contrast and speakers report that the intersective reading is even more difficult to obtain than in the singular case:

- (203) Ali ve Mehmet iyi hırsız
 Ali and Mehmet good thief
 ‘Ali and Mehmet are good thieves’
- a. ‘Ali and Mehmet are good at thievery’ (nonintersective)
 - b. ?*‘Ali and Mehmet are good people’ (intersective)

- (204) Oya ve Elif yetenekli dansçı
 Oya and Elif skilled dancer
 ‘Oya and Elif are a skilled dancers’
- a. ‘Oya and Elif are skilled at dancing’ (nonintersective)
 - b. ?*‘Oya and Elif are generally skillful’ (intersective)

Placing an overt plural marker on the noun rescues the intersective interpretation (to the extent that it’s available normally), e.g.:

- (205) Ali ve Mehmet iyi hırsız-lar
 Ali and Mehmet good thief-PL
 ‘Ali and Mehmet are good thieves’
- a. ‘Ali and Mehmet are good at thievery’ (nonintersective)
 - b. ?‘Ali and Mehmet are good people’ (intersective)

Further, the existential copular position with bare nouns also behaves nearly identically to the individual/stage-level ambiguity:

- (206) Bu mahalle-de iyi hırsız var
 dem neighborhood-LOC good thief exists
 ‘There is good thief in the neighborhood’
- a. ‘There is a skillful thief in the neighborhood’ (singular + nonintersective)
 - b. ‘There are skillful thieves in the neighborhood’ (plural + nonintersective)
 - c. ‘There is a moral thief in the neighborhood’ (singular + intersective)
 - d. *‘There are moral thieves in the neighborhood’ (plural + intersective)

- (207) Ofis-te yetenekli dansçı var
 office-LOC skilled dancer exists
 ‘There is skilled dancer in the office’
- a. ‘There is a skillful dancer in the office’ (singular + nonintersective)
 - b. ‘There are skillful dancers in the office’ (plural + nonintersective)
 - c. ?*‘There is a generally skillful person who is a dancer in the neighborhood’ (singular + intersective)
 - d. *‘There are generally skillful people who are dancers in the neighborhood’ (plural + intersective)

The only difference between the nonintersective examples and the individual/stage examples is that, for some reason, the intersective interpretation of *yetenekli* ‘skilled’ in particular is degraded in the existential copular construction regardless of which number interpretation the noun receives (though the plural interpretation is still perceived as worse). Overall, however, this shows the same pattern as we have seen before: the intersective interpretation of the adjective precludes a plural interpretation of the noun.

Concluding again with the non-case-marked direct object position:

- (208) Ali iyi hırsız gördü
 Ali good thief see.PST
 ‘Ali saw good thief/thieves’
- a. ‘Ali saw a skillful thief’ (singular + nonintersective)
 - b. ‘Ali saw skillful thieves’ (plural + nonintersective)
 - c. *‘Ali saw a moral thief’ (singular + intersective)
 - d. *‘Ali saw moral thieves’ (plural + intersective)
- (209) Ali yetenekli dansçı gördü
 Ali skilled dancer see.PST
 ‘Ali saw skilled dancer(s)’
- a. ‘Ali saw a skillful dancer’ (singular + nonintersective)
 - b. ‘Ali saw skillful dancers’ (plural + nonintersective)
 - c. *‘Ali saw a generally skillful person who are dancers’ (singular + intersective)
 - d. *‘Ali saw generally skillful people who are dancers’ (plural + intersective)

Here again, the (non)intersective ambiguity tracks precisely with both the individual/stage-level ambiguity and the kind/object-level contrast. This provides empirical motivation for assimilating the three patterns, and reducing both kinds of ambiguity to the more funda-

mental contrast between kind- and object-level motivation, in favor of both Leffel’s proposal and the extension of that proposal to (non)intersectivity. With the empirical grounding for this move established, I’ll now move to implementing it: making the necessary revisions to Leffel’s proposal to extend it to (non)intersectivity, and applying it to the other languages under consideration.

4.3 Nonintersectivity as kind modification

4.3.1 Revising the semantic analysis

The Turkish number diagnostic bolsters the empirical grounding for Leffel’s account, but that still leaves the objection on the basis of theoretical complexity laid out in Section 4.1.3 - the need to lexically encode both a kind-level and object-level version of an ambiguous adjective, which in a sense loses out on some of the theoretical ground we otherwise gain from appealing to structural domains to explain the ambiguity. As noted in that section, Leffel does acknowledge that it would be preferable to have adjectives lexically denoting subkind predicates and be converted to individual predicates at the AP level - which would establish a ‘striking’ parallel with the dual structure of nouns - but doesn’t offer a way in which that might be done.

Achieving the same kind of conversion as we have with nouns would, indeed, be difficult. While number is an obvious candidate for the kind of atomizing conversion that takes kind-level meanings of nouns to sets of objects, there is no such morphological equivalent for adjectives. However, the theory of the (non)intersective alternation developed here in Chapters 2 and 3 offers a different perspective on how we might derive the object-level modifier from the kind-level one, one on which intersective readings are derived from two-place adjective semantics that first takes a nonintersective argument. Chapter 2 demonstrates that this two-place approach allows for a clear explanation of the word order and privativity facts in English and Italian (though there, we were agnostic with respect to the semantics of the initial composition, and here we can clarify it with reference to kinds), and Chapter 3

provided independent morphosyntactic evidence for the presence of such a null argument in intersective readings - a null argument that we can now identify as contributing a kind-level denotation. Directly implementing a single adjective denotation and a kind-level null argument, however, is going to run into some significant technical problems. My hope is that we can take advantage of our evidence for null structure by leveraging that structure in a particularly useful way, but let's first see the problems with the naive approach.

Recall Leffel's two schematic denotations for adjectives:

- (210) a. $\llbracket \text{adjective} \rrbracket_{\text{INNER}} = \lambda k_{\langle s, e \rangle} [k \leq \text{ADJECTIVE}]$
 b. $\llbracket \text{adjective} \rrbracket_{\text{OUTER}} = \lambda w [\lambda x [\text{adjective}_w(x)]]$

The simplest way to combine those denotations into a two-place (or three, counting the intensional argument) predicate would be something like:

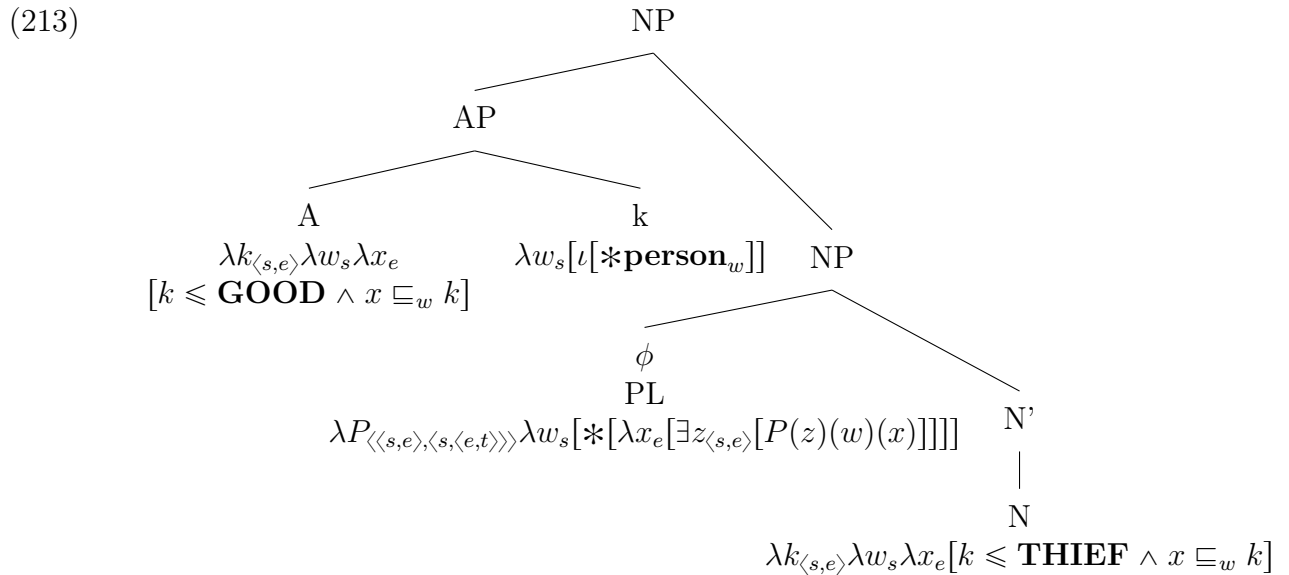
- (211) Revised adjective semantics, Version 1
 $\llbracket \text{adjective} \rrbracket = \lambda k_{\langle s, e \rangle} \lambda w_s \lambda x_e [k \leq \text{ADJECTIVE} \wedge x \sqsubseteq_w k]$

This semantics would first take a kind argument k and then an individual argument x , and evaluate to true if k is a taxonomic subkind of the adjective kind and x is a member of the kind's instantiation in the given world. The idea would be that, in direct modification positions, this composes with the overt noun in the same way that (210a) would, and in indirect modification positions it composes with a covert noun that way, yields an object-level property, and then combines with the already-atomized number-bearing noun in the same way that (210b) would. There are several reasons why this won't do what we want, though. The first is that it won't compose with nouns at all in the pre-number direct modification configuration. If we want to have the adjective output an object property after composing with a kind, then it needs to have an output of type $\langle e, t \rangle$ (or $\langle s, \langle e, t \rangle \rangle$) - but the actual subkind predicate denotation of the nominal root pre-number as defined in (159) only has output type t , and so even though they would both be predicates of kinds, they can't combine via Predicate Modification. One solution to this obstacle would be to adjust the type of the noun roots to match (211) and allow Predicate Modification; this

would also require an adjustment to the semantics of grammatical number, but an easy and inconsequential one since the output would be truth-conditionally equivalent.

- (212) a. Revised noun semantics, Version 1
 $[[\mathbf{noun}]] = \lambda k_{\langle s,e \rangle} \lambda w_s \lambda x_e [k \leq \mathbf{NOUN} \wedge x \sqsubseteq_w k]$
 b. Revised number semantics, Version 1
 $[[\mathbf{SG}]] = \lambda P_{\langle \langle s,e \rangle, \langle s, \langle e,t \rangle \rangle} \lambda w_s [\circ [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z)(w)(x)]]]]$
 $[[\mathbf{PL}]] = \lambda P_{\langle \langle s,e \rangle, \langle s, \langle e,t \rangle \rangle} \lambda w_s [* [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z)(w)(x)]]]]$

These adjustments together would allow composition to proceed through the DP without any type-clash. The adjective and noun could combine via intersection now, and after composition with number the output would be of identical type as in the previous theory, appropriate for composition with the polymorphic determiner. And in the indirect modification configuration for the adjective, these denotations will compose just fine, though we'll have to assume that the null argument is not a predicate of subkinds like overt nouns but rather a kind itself. In this indirect modification configuration, syntactically phrasal adjunction, we would have a structure like (213) for the intersective reading of *good thieves*:



Computing through application of PL to N, application of A to k, and intersection of AP with NP, the resulting denotation is¹⁹:

¹⁹This also involves a necessary simplification step where the pluralization operation $*$ is taken to apply over the entire intersected lambda abstract binding x ; this doesn't change the truth-conditions, since what is effectively happening is the intersection of a set without the plurality requirement with a set with the

$$(214) \quad \lambda w_s[*[\lambda x_e[\mathbf{PERSON} \leq \mathbf{GOOD} \wedge x \sqsubseteq_w \mathbf{PERSON} \wedge \exists z_{\langle s,e \rangle}[z \leq \mathbf{THIEF} \wedge x \sqsubseteq_w z]]]]]$$

This is an intensional plural property which holds of individuals who are (i) members of the **PERSON** kind that is itself a subkind of the **GOOD** kind and (ii) members of a subkind of the **THIEF** kind. While this does compose, an obvious interpretive problem with the result jumps out immediately: this doesn't hold if the individual is a member of some **GOOD** subkind of **PERSON**, but rather if the individual is a **PERSON** and if **PERSON** is a subkind of **GOOD**. That means that someone could only be truthfully described as a *good thief* if they were a thief and all people were good; not at all the result that we want.

This problem shows up in different form for the nonintersective composition, where the denotation of the NP would work out to:

$$(215) \quad \lambda w_s[*[\lambda x_e \exists z_{\langle s,e \rangle}[z \leq \mathbf{GOOD} \wedge z \leq \mathbf{THIEF} \wedge x \sqsubseteq_w z]]]]]$$

Trying to derive nonintersective composition via intersection of subkind predicates leads to a situation where all we can state is that the predicate is true of individuals that belong to one single kind, and that kind is both a subkind of the adjective kind and the noun kind. What might such a kind z as defined above actually look like? Certainly, one way of interpreting it is the kind of people who are good at thieving, which would be the nonintersective reading that we want. But this isn't at all the only way to interpret it, and that uncertainty comes from uncertainty surrounding what this independent kind **GOOD** even is. Is that the type of concept that can correspond to a kind on its own, not interpreted relative to any particular set or other kind? If the kind **GOOD** needs to capture both the 'skill at X' readings and the 'moral' readings that it receives, then a perfectly valid option for z would be the kind comprised of those who are both morally good and thieves - that is, the intersective reading. If the kind **GOOD** that appears in both (214) and (215) is the same thing, then there is no way to prevent the nonintersective composition from being satisfied

plurality requirement, and that intersection will exclude any members of the first set that don't meet the number requirements imposed on the second. This is more consequential with SG than PL, but either way, the simplification here is logically consistent.

by the intersective reading.

4.3.2 Visibility problems

I think this is a consequence of the data that Leffel’s account focuses on: for cases like *visible star*, it may be reasonable to assume that there is some independent kind **VISIBLE**, and the only distinction between the readings is whether the individual is required to belong to that kind as well as the **STAR** kind, or just to the **STAR** kind while bearing the *visible* property. Of course, we could also disagree with this, and say that what it means to be visible as a star is something quite different from what it means to be visible as, say, a table. After all, you would probably be inclined to utter both of the sentences in (216) when the amount that you could see the table and the star were the same - if you could only see a table to the same extent that we can normally see stars in the sky (at a massive distance, making out almost none of its features, only just able to tell that it exists) then you would probably not call it *clearly visible*.

- (216) a. The table is barely visible.
 b. The star is clearly visible.

Taking this kind of contrast seriously suggests perhaps that Leffel’s denotation for the individual-level reading of *visible star*, repeated below as (217), is inadequate.

- (217) Individual-level reading of *the visible stars*
 $\iota x_e [\exists k_{\langle s,e \rangle} [k \leq (\lambda w [\iota [*visible_w]])] \wedge k \leq (\lambda w [\iota [*star_w]])] \wedge x \sqsubseteq_{w'} k]$

Recall that the approach Leffel takes, of intersecting predicates of subkinds that gives us these $k \leq A \wedge k \leq B$ denotations, was inspired by McNally & Boleda’s (2004) analysis of relational adjectives like *technical*:

- (218) a. John is a technical architect.
 b. El Marti es arquitecte tecnic (Catalan)
 the Marti is architect technical
 ‘Marti is a technical architect’

McNally & Boleda (2004) argue that these relational adjectives are predicates of kinds,

and compose via the noun - which is also a predicate of kinds - via simple intersection. The result is denotations of (218b) like (219), where some kind k is connected to the subject via a relation R , is an **ARCHITECT** kind, and is a **TECHNICAL** kind.

$$(219) \quad R(\mathbf{marti})(k) \wedge \mathbf{architect}(k) \wedge \mathbf{technical}(k)$$

Leffel extended this kind of analysis to adjectives like *visible*, resulting in denotations like (217) with the same intersective $k \leq A \wedge k \leq B$ structures. But it's not clear if this extension is desirable given the syntactic patterns that motivate McNally & Boleda's analysis. Specifically, they argue that relational adjectives should be analyzed as intersective (just intersective over kinds) because they appear obligatorily postnominal in Romance languages like French and, shown here, Catalan, just like intersective adjectives:

- (220) a. un escriptor jove
 a writer young
 ‘a young writer’
 b. una malaltia pulmonar
 a disease pulmonary
 ‘a pulmonary disease’

But as was covered extensively in Chapter 2, this is not at all the behavior of either individual/stage adjectives like *visible* or (non)intersective adjectives like *good* or *fake*. The individual or nonintersective readings, which Leffel models with kind intersection like relational adjectives, are available in both prenominal and postnominal positions in Romance, in fact being the only readings available in prenominal positions where relational adjectives are entirely unavailable. This strongly suggests that it is undesirable to assimilate the inner readings of these adjectives to the same compositional process as relational adjectives.

Additionally, extending this to nonintersective adjectives runs into not only the interpretive problems described above, but it would only get worse if we tried to model privative adjectives in the same way. An output for something like *fake thief* would end up as:

$$(221) \quad \lambda w_s[*[\lambda x_e \exists z_{\langle s,e \rangle} [z \leq \mathbf{FAKE} \wedge z \leq \mathbf{THIEF} \wedge x \sqsubseteq_w z]]]]$$

This is obviously problematic - the inclusion of the term $z \leq \mathbf{THIEF}$ guarantees that the

output set contains only thieves, specifically the opposite of what we want for the nonintersective and nonsubsective reading of an adjective like *fake*. What we need is for the adjective to be able to directly modify the kind under consideration, rather than contributing a kind of its own. This means that it cannot be composing with the noun via intersection in direct modification positions; instead, we want to raise the type of the adjective further so that it can take a subkind predicate as argument.

4.3.3 A genuinely nonintersective analysis

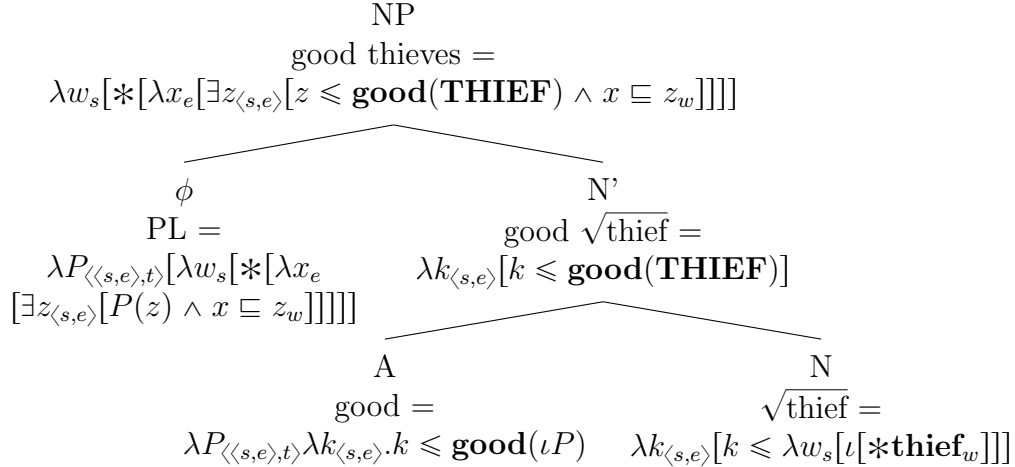
What kind of output do we want? In order to keep the same semantics for number, the adjective+noun phrase should be the same type as the nominal root, a predicate from subkinds to truth values. This is also important for potential stacking of nonintersective adjectives: having a nonintersective adjective directly convert its argument into a object-level property would prevent sequential application of nonintersective adjectives, only allowing mass intersection of the adjectives before application, which will not derive the correct ‘nested’ readings for these constructions. And we want the content of the adjective to directly modify a kind. The most natural candidate for that kind is the maximal kind of the noun’s subkind predicate, e.g., the **DOG** kind for the noun *dog*, which we can retrieve with an *iota* operator just like we saw for singular kind readings. So what we end up with is:

(222) **Schematic denotation for a (non)intersective adjective:**

$$\llbracket \mathbf{A} \rrbracket = \lambda P_{\langle \langle s,e \rangle, t \rangle} \lambda k_{\langle s,e \rangle} . k \leq A(\iota P)$$

Let’s see how this composes in a direct modification context with the *good thieves* example, and be a little more meticulous about it now that we are demonstrating the final proposed analysis. Here **THIEF** refers to the maximal kind of the subkind predicate denoted by *thief*; when a kind is being modified directly by a nonintersective adjective, I’ll generally use the capital shorthand for ease of readability.

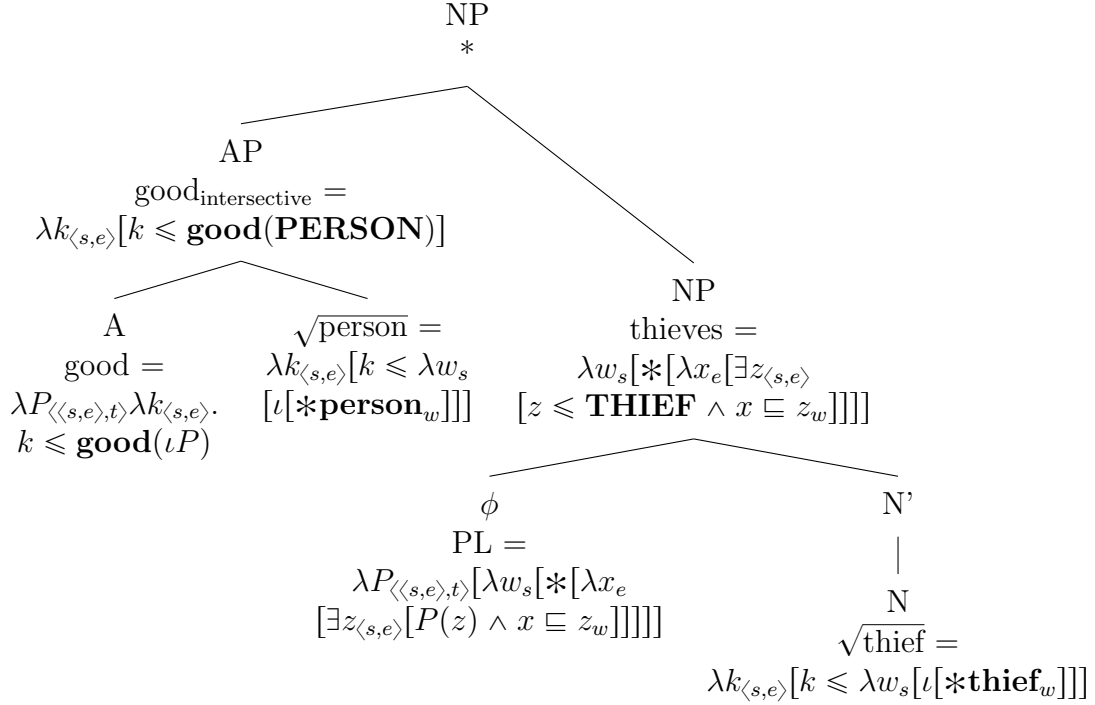
(223) **Derivation of a nonintersective meaning:**



The obvious question about this derivation is the interpretation of **good(THIEF)**, or generally any adjective directly modifying a kind. This will be in large part the focus of Chapter 5. For now, it should be thought of as a process that creates a new kind by adding a particular dimension of *good* to the preexisting *thief* kind, in contrast to the semantics offered by Leffel, which establishes a subkind relation between the two kinds. Importantly, it is just as easy to state a privative adjective's denotation in this form as a subsective one's, because no direct subkind relationship between the existentially quantified kind and the **THIEF** kind itself is stated, allowing for potential nonsubsectivity.

There is still work to do in deriving intersective meanings with this approach, however. The original idea floated at the beginning of this section was to allow adjectives to take in kind meanings and convert them to object-level meanings in much the same way that number does, such that saturating the kind argument with the null argument proposed in Chapters 2 & 3 would result in a property meaning appropriate to intersect with the singular or plural noun. The denotation proposed in (222) can no longer do this, because doing so would result in a type-clash with either additional direct modifiers or with number. But this means that giving a null subkind predicate argument in the indirect modification configuration will lead to a type-clash at the phrasal adjunction level instead:

(224)



Preserving a single denotation for the ambiguous adjective alongside a compositional derivation of ambiguity and the idea that number converts kind-denotations to object-denotations ends up being no trivial enterprise. There are a number of ugly semantic maneuvers that might work, by adjusting the basic denotations of nouns and number, but all of those run into obstacles in their inconsistency with established semantic analyses of those areas, and one of the primary appeals of this reduction-to-kinds analysis was the preexisting nature of these pieces of semantic machinery. Because of this, the line that I want to pursue here is one that preserves as much as possible about nouns and number and instead takes advantage of the novel evidence offered here for null structure inside the adjective phrase. The cost we have to pay for avoiding polysemy and preserving a consistent semantics of the rest of the structure is introducing additional complexity into the null structure, which is only somewhat syntactically motivated - Chapter 3 provides evidence for the syntactic reality of *some* element in that position, but not what it is or how much structure is there. Here, I'm going to make a concession on that front by proposing that it's more complicated than simply a null version of the overt noun.

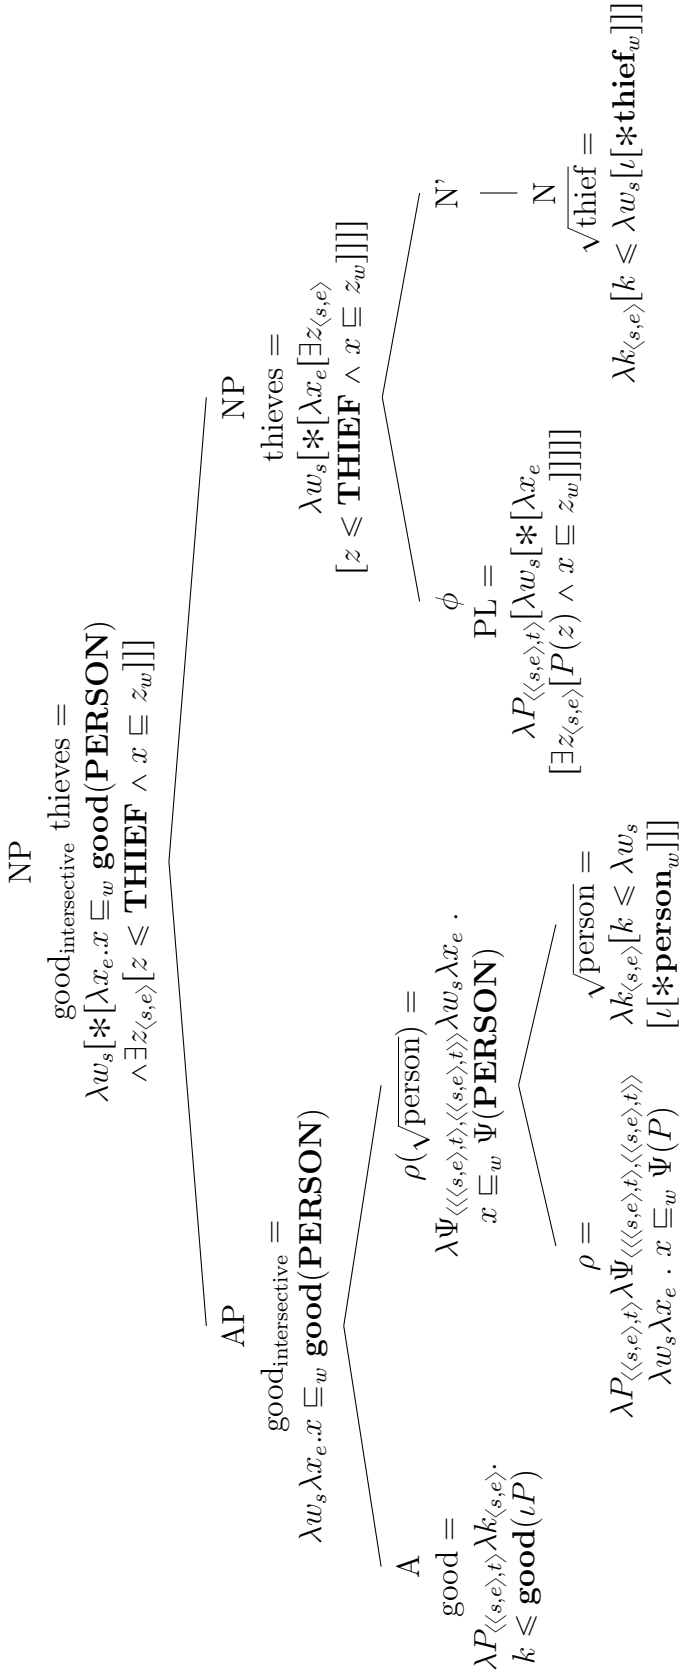
Specifically, on top of a null subkind predicate that is semantically equivalent to an overt

noun, I suggest that there is an operator that mediates between the adjectival root and that null nominal, creating an object-level property out of the adjective that is interpreted relative to the null nominal. Call this operator ρ , for ‘relative’, and use Ψ for the type of function from subkind predicate to subkind predicate which we’ve now defined nonintersective adjectives as.

$$(225) \quad \llbracket \rho \rrbracket = \lambda P_{\langle\langle s,e \rangle, t \rangle} \lambda \Psi_{\langle\langle\langle s,e \rangle, t \rangle, \langle\langle s,e \rangle, t \rangle\rangle} \lambda w_s \lambda x_e . x \sqsubseteq_w \Psi(P)$$

The relativizing operator ρ takes a subkind predicate - which will be the null argument discussed throughout, N_C from Chapter 2 - and a nonintersective adjective, applies the adjective to the subkind predicate, and then returns an intensional predicate over individuals that checks if the individual is a member of the new modified kind. This output is of appropriate type for combining via Predicate Modification with the number-bearing noun. Assuming that the null syntax that Chapter 3 provided evidence for consists, in fact, of a complex of both ρ and the null nominal, e.g., PERSON, we can derive the intersective meaning of *good thieves* as follows:

(226) Derivation of an intersective meaning:



This correctly derives an adequate intersective semantics for *good thieves*: an intensional predicate of individuals which checks if an individual is a member of the ‘good person’ kind, and a member of some subkind of the ‘thief’ kind. The asymmetry here - where it checks directly the ‘good person’ kind but checks ‘thief’ via existential quantification over subkinds - is a potentially odd, but inconsequential, relic of the fact that the number denotation comes from Leffel’s formulation and the ρ conversion is my own; we could easily change the ρ denotation to also include this existential step without affecting the truth conditions, since a member of a subkind of a kind will always trivially be a member of that superkind.

One open question is whether it is desirable to revise Leffel’s semantics for *visible* to match the proposed semantics here for (non)intersective adjectives. This partially depends on the discussion above regarding (216), whether you’re satisfied with kind intersection or think that it’s not appropriate to treat *visible* as a kind in itself. In general, the parallels in terms of syntactic distribution of individual/nonintersective and stage/intersective meanings might speak in favor of this. However, given that the core empirical motivation for this particular approach to nonintersective adjectives is from privative adjectives, the same evidence can’t be used about the individual/stage-level contrast, which doesn’t have a privative equivalent in its class. If such assimilation is desired, though, it should be straightforward enough to redefine the relevant adjectives.

4.3.4 Grammatical number versus classifiers: the view from Bangla

All of the discussion in this chapter thus far has centered around the way that grammatical number - namely, the singular and plural morphemes - encodes the conversion process from kind-level denotations to object-level denotations, with the result that full NPs denote object-level properties (assuming the location of SG and PL in [Spec,NP]). But this is not a cross-linguistically resilient assumption: one of the basic parameters that distinguish languages in terms of their kind reference and bare argument options is whether countability is enforced via morphological number or via a classifier. In classifier languages, the mass-count

distinction familiar to speakers of English is grammatically absent, and all nouns require the addition of a classifier word for counting nouns. Bangla is one of those languages.

- (227) a. kal ek *(-t̥a) / du *(-t̥o) tʃʰatro eʃe tʃʰilo (Dayal 2012:196)
 yesterday one CL two CL student came
 ‘Yesterday a student/two students came.’
 b. anu ek *(-t̥a) / du *(-t̥o) boi kinetʃʰilo
 Anu one CL two CL book bought
 ‘Anu bought a book/two books.’

Attempting to use a numeral to count the nouns without the inclusion of the classifier results in ungrammaticality. Bangla has a diverse lexicon of classifier options which are restricted to different noun class; here we’ll stick to examples using the most general classifier option *-ta/to/te* (with its allophonic variants). Recall from Chapter 2 that (non)intersective modification appears to be sensitive to the position of the adjective and noun relative to the classifier: when the adjective and noun are separated by the classifier, intersective readings result, and when the adjective and the noun are together (on either side of the classifier), the result is nonintersective readings (with the strengthening effect of the preference for privative adjectives):

Adjective Type: Interpretation:	Subsective		Privative	
	Nonintersective	Intersective	Nonintersective	Intersective
DEM NUM-CL [ADJ NOUN]	✓	?	✓	??
DEM [ADJ NOUN] NUM-CL	✓	?	✓	??
DEM [ADJ] NUM-CL [NOUN]	?	✓	*	✓
[ADJ] DEM NUM-CL [NOUN]	?	✓	*	✓

(228)

Given the analysis of (non)intersective readings developed in this chapter, this pattern suggests that kind-level modification is only possible when the adjective is local to the noun relative to the classifier, and when the classifier intervenes, only object-level modification can take place. In that way, classifiers in Bangla would be performing the same role defining modification domains as morphological number marking does in English and Turkish. This conclusion is natural if we assume that the necessity of classifiers for counting is due to clas-

sifiers performing the same kind-to-object conversion in Bangla that morphological number marking does in English and Turkish.

And in fact, Dayal (2012) argues that that is precisely what Bangla classifiers do: mediate between kind- and object-level denotations of nouns. Bare NPs, sans classifier, in Bangla denote kinds, as is the standard account of non-number marking languages (Krifka et al. 1995, Chierchia 1998), as indicated by their ability to be arguments of kind-level predicates:

- (229) a. $\text{\textcircled{t}atayat-er}$ $\text{sahod\textcircled{m}ad}^{\text{h}}\text{yom holo}$ gari (Dayal 2012:209)
 transport-GEN convenient means be car
 ‘Cars are a convenient mode of transport.’
 b. $\text{pat}^{\text{h}}\text{odoffotabdi-te boi}$ $\text{t}^{\text{h}}\text{apa juru hoi}$ $\text{t}^{\text{h}}\text{ilo}$
 fifteenth century-LOC book printed start happened
 ‘Books started to be printed in the 15th century.’

The classifier is a function from kinds to sets of objects:

- (230) $\llbracket \text{-ta} \rrbracket = \lambda k_{\langle s, e \rangle} \lambda x_e [\cup k(x) \wedge AT(x)]$
 where AT is the atomic function, denoting the set of those individuals that do not have proper parts

The AT function, though of a different type, is equivalent in function to the \circ operator used by Leffel (2014), making this the classifier equivalent of the singular morpheme; specifically plural classifiers (e.g., *-gulo*) also exist with a non-atomic requirement. The classifier thus creates sets of object-level individuals which are of appropriate type for combination with numerals, which pick out sets of a particular cardinality. (230) is defined with the assumption that bare nouns denote kinds directly, rather than subkind predicates as Leffel (2014) assumes (and Dayal also assumes in other work for other languages); again, conversion between the two theories is easy to accomplish, and so that we are able to retain a consistent schematic denotation for the adjectives cross-linguistically, I’ll assume instead a subkind predicate ontology for nominal roots. Adjusting the semantics of the adjective and ρ would be simple enough if one wanted to maintain a kind denotation for nouns, though:

- (231) a. $\llbracket \mathbf{A} \rrbracket_{\text{over kinds}} = \lambda k_{\langle s, e \rangle} . A(k)$
 $= \lambda k_{\langle s, e \rangle} \lambda w_s [\iota [*A(\mathbf{thief})_w]]$
 b. $\llbracket \rho \rrbracket_{\text{over kinds}} = \lambda k_{\langle s, e \rangle} \lambda \Psi_{\langle \langle s, e \rangle, \langle s, e \rangle \rangle} \lambda w_s \lambda x_e . x \sqsubseteq_w \Psi(k)$

Retaining a subkind predicate semantics for nouns, however, would just involve redefining (230) to be like Leffel’s SG morpheme (165a).

Either way, Dayal’s analysis predicts that kind-level modification should only be able to take place if the adjective accesses the noun prior to its composition with the classifier. On the view that nonintersective modification is kind-level modification, the pattern from (228) validates precisely that. However, we can make it even more clear by testing unambiguous examples of kind and object modification, much like the data that we used to establish the diagnostic in Turkish. For all of the following data, keep in mind that the adjective movement alone automatically induces some degradation in acceptability, so the paradigm below with ? on the moved instances is the maximally acceptable situation:

- (232)
- a. oi du-to bhari boi
DEM 2-CL heavy book
‘these two heavy books’
 - b. ?oi bhari du-to boi
DEM heavy 2-CL book
‘these two heavy books’
 - c. oi bhari boi du-to
DEM heavy book 2-CL
‘these two heavy books’
 - d. ?bhari oi du-to boi
heavy DEM 2-CL book
‘these two heavy books’

Here, *b^hari* ‘heavy’ is unambiguously an object-level modifier; ‘heavy books’ is not a subkind of books in the way that ‘scientific books’ in Turkish was. And, with the standard caveat that the lone-adjectival movement is slightly degraded by default, this kind of modification is acceptable in all positions. With a subkind modifier, however:

- (233)
- a. oi du-to golper boi
DEM 2-CL story book
‘these two story books’
 - b. *oi golper du-to boi
DEM story 2-CL book
‘these two story books’
 - c. oi golper boi du-to
DEM story book 2-CL

- ‘these two story books’
 d. *golper oi du-to boi
 story DEM 2-CL book
 ‘these two story books’

Here, the modifier *golper* ‘story’ - which picks out a subkind of books - is only grammatical when local to the noun, and ungrammatical when split across the classifier. This is exactly the pattern that we see with nonintersective interpretations of the privative adjectives, in the same way that nonintersective interpretations pattern alongside kind modifiers with respect to number neutrality in Turkish. Bangla therefore provides additional evidence for assimilating the (non)intersective contrast to the kind/object-level distinction, and specifically that the syntactic domain that demarcates the two interpretations will change cross-linguistically depending on how that language encodes countability. In languages where morphological number marking converts from uncountable kinds to countable objects, that number marking splits the domains of nonintersective from intersective modification; in languages where classifiers do the same conversion, they instead split the domains of modification.

4.4 Interim summary

This chapter laid out a proposal for reducing the (non)intersective ambiguity to an ambiguity between kind-level (nonintersective) modification on the one hand and object-level (intersective) modification on the other. Basic principles behind reference to kinds in natural language were established, and a particular implementation of those principles by Leffel (2014) was introduced. On this implementation, bare nominal roots denote predicates of subkinds, and number marking (singular and plural morphemes) convert these kind-level denotations to properties of objects at the NP level. Leffel specifically argued that adjectives which show a contrast in individual vs stage-level modification, like *visible*, come in two forms: one that is a subkind predicate and adjoins to and intersects with nouns for individual-level modification, and one that is an individual predicate and adjoins to and intersects with noun phrases for stage-level modification. On this account, then, both meanings are in fact the result of

intersection.

Three problems with Leffel’s account were identified. First, it lacks empirical motivation for the assimilation of direct/indirect modification to kind/object-level modification - while the proposal technically functions, the motivation for treating the contrasts as instances of each other is merely intuitive. Second, it still requires a polysemous semantics for adjectives, despite complicating the syntactic structure and proposing different attachment heights for the two meanings. Third, the Predicate Modification approach to kind-level modification fails to derive the correct readings for individual-level or nonintersective adjectives.

Further empirical support for Leffel’s hypothesis is found in Turkish and Bangla. In Turkish, it has been argued by Sağ (2021) that bare singular nouns are in fact ambiguous between denoting (grammatically singular, but conceptually plural) kinds and definites, resulting in what on the surface appears to be a number-neutral reading where they may be read as either plural or singular in certain syntactic configurations. Adjectival modification interacts with this number-neutrality, however, with kind-level modifiers preserving the plural interpretations and object-level modifiers eliminating them. This constitutes a potential diagnostic for kind-level modification, one which we applied to both the individual/stage contrast and the nonintersective/intersective contrast. The results supported Leffel’s hypothesis and the (non)intersective extension: individual and nonintersective interpretations of adjectives behaved like kind modifiers, preserving plural interpretations, while stage and intersective interpretations behaved like object modifiers, disallowing plural interpretations of bare singular nouns. Extending the pattern to Bangla, a classifier language, shows that nonintersective modifiers again behave like kind modifiers (here, only being able to receive nonintersective interpretations when they compose with the noun prior to composing with the classifier) and intersective modifiers behave like object modifiers.

A revised version of Leffel’s account was developed, with the following properties:

- The basic denotation of nominal roots is a subkind predicate.
- Number (either morphological or classifiers) converts subkind denotations to object

properties.

- (Non)intersective adjectives have a uniform semantics as functions from subkind predicates into subkind predicates, as a result of modifying the kind itself.
- Intersective readings in indirect modification configurations are derived through a relativizing operator which creates an object property out of the adjective by saturating its kind argument with a null subkind predicate.

This account is advantageous in its ability to dispense with duplicate lexical entries for adjectives, instead deriving ambiguous meanings compositionally; the true nonintersective meanings derived through direct modification of kinds (rather than intersection of two subkind predicates); and the reduction of the (non)intersective ambiguity to the independently motivated and necessary kind/object distinction. The downside of the account is the need to posit additional covert structure in the relativizing operator and its null subkind predicate argument; however, this maneuver is at least interface-accountable - in Bouchard (2002)'s sense - because of the morphophonological evidence from Chapter 3 for additional null syntactic structure in the adjectival phrase. The need to make that structure a complex with a novel operator semantics is a significant downside of the proposal, compared to requiring only a null nominal argument as suggested in earlier chapters, but I believe the advantages we get in preserving a compositional account without polysemy make it a line worth pursuing.

As I note earlier in this chapter, the most significant open question for this account regards the nature of the direct nonintersective kind modification - when the denotation for a phrase like *good thief* involves creating the composite kind **good(THIEF)**, the result of modifying the **THIEF** kind with **good**, what exactly is occurring? What does it mean to modify a kind with an adjective and create a novel kind? One, probably common, perspective would be that this kind of question is no longer the responsibility of the compositional semantics. Kinds, on this view, may be regarded as monistic from the perspective of semantics; not internally compositional with parts that can be modified, but simply a conceptual individual - **DOG**

encapsulates whatever it means to be a dog, which is the purview of some world knowledge component and not narrow semantics, and the same for **good**, so determining the output of **good(DOG)** is equally not a compositional question but something that semantics simply receives from external modules of cognition. Even if it is not frequently spelled out explicitly, this view is most likely common among formal semanticists. One who shares this view may choose to take the analysis developed up til this point as complete, in that sense, then, and either stop here or read on with the understanding that the remaining chapter is functionally a work of cognitive psychology, not formal semantics. Personally, I suspect that we have good reason to consider the inner workings of nonintersective kind modification to be the domain of compositional semanticists, and so it is to that question that I will now turn.

5 Inside nonintersective modification

In Chapter 4, I argued that both nonintersective and intersective interpretations of ambiguous adjectives result from the same kind of compositional operation: modification of a subkind predicate. Nonintersective interpretations result from the adjective directly modifying the subkind predicate denoted by their overt modificand noun, while intersective interpretations are the result of the adjective first directly modifying a covert subkind predicate providing a different scale of evaluation (via a linking operator which creates an object-level property out of the result), and then intersecting with the property denoted by the enumerated noun. On this account, both interpretations in fact involve a critical step of what we might call nonintersective modification, in the sense that the adjective's semantic content modifies a noun's semantic content via Function Application rather than set intersection. The goal of this chapter is to investigate further the nature of that nonintersective step.

The evidence reviewed thus far supports a treatment of the 'semantic content' involved in this modification as kinds. We have also so far described kinds as some abstract conceptual object - a function from worlds to individuals, for sure, but the internal component of the kind that determines an object's membership in the category, the bold lowercase aspect of denotations like (234), has been left underspecified, or in Chierchia's words repeated again, 'Kinds are whatever your favorite worldview says that they are.' (Chierchia 1998:350).

$$(234) \quad \mathbf{DOG} = \lambda w[\iota x[*\mathbf{dog}_w(x)]]$$

But understanding what it means for an adjective to nonintersectively modify a kind, creating a novel kind out of the noun's lexical kind, requires that we crystallize our worldviews somewhat and understand what that kind term actually means. Towards that end, this chapter explores and sketches the boundaries of a semantic worldview on which kinds refer to sortal concepts - collections of conceptual knowledge that allow for the categorization of individuals - in a particular way. This is combined with a syntactic worldview on which acategorical word roots identify and index concepts, and the process of syntactic categoriza-

tion (into nouns, verbs, adjectives...) imposes certain informational structures onto those concepts to organize them into compositionally usable representations (of kinds, in the case of nouns). Combined with the description of the enumeration process in Chapter 4 - wherein number marking or a classifier converts kind representations to properties - the result is a sketch of the full compositional pipeline from pre- or non-linguistic conceptual knowledge to the type of $\lambda x.dog(x)$ terms with which undergraduate semantics students are familiar. Treating these stages compositionally enables a theory of modification like the one proposed here, where differences in intersectivity arise from differences in attachment height causing modification of distinct compositional stages of nominal meaning.

The particular motivation and empirical testing ground for this sketch is going to be privative adjectives, which were introduced as the focus of Chapter 2 but have taken a background role in Chapters 3 and 4. But privative adjectives provide the ideal test case for a theory that enriches the conceptual structure of semantic denotations, because I believe the problems for a simple compositional theory of semantics that they have been long recognized to pose in turn constitute linguistic evidence they can provide for such conceptual enrichments. The structure of this chapter will therefore be as follows: I'll begin by reviewing the core problem of privativity, the issues it raises for compositional semantics, the canonical account of privative adjectives from Partee (2007, 2009) and its problems. I'll then describe two experiments aimed at clarifying the behavior of privative adjectives across contexts, which will help motivate a counterproposal for privative semantics from Del Pinal (2015, 2018), with some revisions. Then I will connect this semantic proposal to the discussion of kinds from Chapter 4, and the syntactic half of the worldview described above, concluding with some exploratory suggestions for what certain scattered empirical problems might look like under this worldview.

5.1 Returning to the problem of privativity

5.1.1 Privativity as a compositional problem

Privativity is the linguistic phenomenon in which the extension or reference of an adjectivally modified noun phrase is entirely disjoint from the reference of the head noun. The paradigmatic case of a privative adjective is *fake*, and the canonical example of a privative noun phrase is *fake gun*, where the intuition is that the set of *fake guns* contains none of the members of the set of *guns*, and vice versa. Nor can we derive the extension of *fake gun* through any set-theoretic operations on the *gun* set and the *fake* set, if the latter set exists at all.

Contrast this with other members of the adjective category. The classical typology of adjectives following Kamp (1975) and Montague (1970) splits the category first by subsectivity. The subsective adjectives are those for which the reference of the noun phrase is a subset of the reference of the noun. For a given subsective adjective A_S and a given noun N , then, $A_S N \subseteq N$. These adjectives license thus inferences from the noun phrase to the noun: if X is an $A_S N$, then X is an N . A special case of the subsective adjectives are the intersective adjectives, where the reference of the noun phrase is given by the intersection of the references of the adjective and the noun. Then for a given intersective adjective A_I and a given noun N , $A_I N = A_I \cap N$. On top of the standard subsective inference, intersective adjectives license additionally inferences to the adjective and inferences across nouns: if X is an $A_I N$, then X is A_I and for any other noun N' , X is an $A_I N'$. Intersective adjectives are often considered the simplest case, usually analyzed as one-place predicates of type $\langle e, t \rangle$, semantically equivalent to common nouns.

(235) Subsective inference

- a. $\frac{X \text{ is an } A_S N.}{X \text{ is an } N.}$

(236) Intersective inferences

- a. $\frac{X \text{ is an } A_I N.}{X \text{ is } A_I.}$

- X is A.
- b. X is an A_I N.
- X is an N'
- X is an A_I N'.

Intersective adjectives are assumed to include things like *red*, *circular*, *German*, *carnivorous*. A *red ball* is both *red* and a *ball*, and it is also *red* as any other noun: it is a *red object*, a *red toy*, and so on. Non-intersective subsectives, which license the inference in (235) but not those in (236), include *skillful*, *good*, *beautiful*, *large*. We can infer that a *skillful dancer* is a *dancer*, but not that they are *skillful* in any more general sense, and we can infer that a *short basketball player* who is also a *chef* on the side need not be *short chef* at all. These subsectives are more semantically complex than intersectives, but crucially determining the reference of the modified noun phrase is still possible using only the set denoted by the noun, combined with perhaps some degree of comparison contributed by the adjective: all members in the *skillful dancer* set can be found in the *dancer* set, just like all members of the *red ball* set can be found in the *ball* set.

Non-subsectives are divided into the modals and the privatives. Modal non-subsectives, like *alleged* and *potential*, license no inferences of any kind, suggesting only the possibility of set membership, but validating none. Privative adjectives, by definition, license the opposite of the subsective inference: for a given privative adjective A_P , if X is an $A_P N$, X is not an N . Canonical examples of privative adjectives include *fake*, *mock*, *counterfeit*, *false*, *former* and perhaps also *imaginary*, *pretend*, *virtual* (see Nayak et al. 2014 for an attempt at comprehensive categorization). Again, contrasting with subsectives, the set denoted by a privatively modified noun phrase is disjoint from the set denoted by the unmodified noun, and no set-theoretic operation over the unmodified noun set can derive the noun phrase set.

- (237) Privative inference
- a. X is an A_P N.
 - X is not an N.

Hence, the problem of privativity. On a standard, strong interpretation of compositional-

ity, the meaning of a complex expression like a modified noun phrase is to be fully determined by the meanings of its constituent parts (Szabó, 2012). In the case of intersective adjectives, this is satisfied straightforwardly: intersect the meanings. For other subsectives, the process may be slightly more complex, but it is still easily computable, since we are selecting a subset of the noun's meaning according to a standard provided by the adjective's meaning. For privatives, however, it appears that we must pull from beyond the extensional denotations of both the noun and the adjective to generate the denotation of the noun phrase. Of course, this is not much of a problem for the cognitive system to accomplish. We all readily compute the set of *fake guns* one way or another; its membership is no mystery to a competent language user. But how such computation is achieved by the compositional semantic system is significantly more opaque, given that no meaning contributed by *gun* seems to be preserved in the end result, at least on a standard semantic conception of extensional, referential meaning.

The inference pattern alone creates problems when faced with data like (238), where the truth of (238a) predicts, incorrectly, that a question like (238b) should not be well-formed.

- (238) a. A fake gun is not a gun. (Partee, 2009, 14)
 b. Is that gun real or fake?

Privative adjectives have received substantial attention in the frameworks of cognitive or conceptual semantics (Franks, 1995; Coulson and Fauconnier, 1999), because these inference facts appear initially to be more easily captured through an account of shifting senses or conceptual blending rather than a strictly compositional one. However, there has also been a mounting defense of compositional accounts of privativity. In the following sections, I will introduce what has become the canonical compositional account of such phenomena, due to Partee, and then raise some potential problems left unresolved in this account.

5.1.2 The coercion account

An analysis with its roots in Kamp and Partee (1995) and further developed in Partee (2007, 2009, 2010) argues that we can derive so-called privative behavior by assuming that such adjectives are, in fact, subsective rather than privative, and that their apparently non-subsective behavior arises via interaction with natural principles of semantic interpretation. Motivating the first assumption, that there are no ‘true’ privatives, is primarily syntactic data from Polish, which I will not review here besides to say that certain syntactic tests like the ability to split the adjective from the rest of the NP show that nonsubsective privatives pattern with subsectives rather than with nonsubsective modals, which appear to form their own class.

The basic structure of the account is the following. Both *real*-type adjectives and *fake*-type adjectives are subsective. Our default interpretation of any bare noun is literal, and only includes real instances of that noun category. Thus, when *real* combines with a noun, e.g., *gun*, the composition is vacuous; all guns are real guns, by default, and so $\llbracket \mathbf{real\ gun} \rrbracket = \llbracket \mathbf{gun} \rrbracket$. This violates the following principle of interpretation:

(239) **The Non-Vacuity Principle (NVP)**

In any given context, try to interpret any predicate so that both its positive and negative extension are non-empty. (Kamp and Partee, 1995, 161)

In the context of *gun*, here, the predicate *real* would have an empty negative extension, as it excludes no member of $\llbracket \mathbf{gun} \rrbracket$ from $\llbracket \mathbf{real\ gun} \rrbracket$, and so incurs a violation of NVP²⁰. A similar, but inverse, problem arises in composition with *fake*. As $\llbracket \mathbf{gun} \rrbracket$ is taken to include only real guns, $\llbracket \mathbf{fake\ gun} \rrbracket = \emptyset$ and *fake* has no positive extension in its role as a predicate here. This violates NVP as well. So, neither composition can proceed as currently stated.

²⁰There is some confusion about the way that this process is stated in Partee’s work. The way I have chosen to describe the process here, where it is *real* (and *fake*) that are the problematic predicates incurring violations, is the most natural way to follow the argument I see. However, to be precise, (Partee, 2009, 18) explicitly clarifies that the NVP ‘applies not only to simple predicates but to predicates formed by combination of an adjective and a noun: these should be interpreted in such a way that the ADJ + N combination is a non-vacuous predicate.’ This implies that Partee takes *real gun* (pre-coercion) to be what incurs the NVP violation. This version of the logic, I do not follow - even if *real* ends up being redundant, *real gun* would have the same extension as *gun*, which is not itself a vacuous predicate.

This kind of NVP violation is not limited to *real* and *fake*, of course; otherwise it would be a very stipulative principle. It also occurs in contexts like (240). In (240a), *tall* behaves much like *real* - on the natural assumption from world knowledge that all basketball players (or, at least, professional ones) are *tall* according to a more general standard, the application of *tall* here would be vacuous. Similarly, *tall* in (240b) behaves like *fake* - again, on the assumption that all relevant jockeys are short by general standards, *tall* would lack positive extension.

- (240) a. He is a tall basketball player.
b. He is a tall jockey.

The difference between the cases in (240) and those in (238) is the repair mechanisms available. As we can intuitively observe, the actual interpretations of the sentences in (240) involve some recalibration of the adjective, relativizing *tall* to a narrower standard of comparison than in its more general use. Our understanding of the nouns does not change. This is due to a competing principle of interpretation:

(241) **The Head Primacy Principle (HPP)**

In a modifier-head structure, the head is interpreted relative to the context of the whole constituent, and the modifier is interpreted relative to the local context created from the former context by the interpretation of the head. (Kamp and Partee, 1995, 161)

The HPP acts as a kind of faithfulness constraint (in Optimality Theory terminology) on the syntactic head of the compound, and says that you should adjust your interpretation of the modifier before you adjust your interpretation of the modified head; that is, the head outranks the modifier in preserving its default semantics. This is why (240a) does not result in an interpretation where we allow a different class of individuals, with more varying heights, to count as *basketball players*.

However, it seems that such an approach, which would violate HPP, is precisely what obtains in the context of *fake*. We understand *fake guns* to be literally fake, and it is rather the interpretation of the head noun *gun* which has undergone some adjustment. Specifically, **[[gun]]** has been coerced into expanding to include both real and fake guns (in order to

avoid violating NVP), and then *fake* acts as subsective adjective over this expanded set. An exactly parallel process occurs for *real*.

The fact that such a process can obtain, in violation of HPP, is explained by the assumption that NVP outranks HPP. If a coercion of the modifier's meaning which would prevent an NVP violation is not available, it is acceptable to incur a violation of the lower-ranked HPP by coercing the head's meaning in order to rescue the compound from violating NVP.

Thus, on Partee's account we are able to treat privative behavior in a compositional system by eliminating the class of privative adjectives and instead assuming that they obtain their meaning via standard subsective interpretation, after an instance of a generally-available pragmatic modulation, induced by natural principles of interpretation. From the point of view of an adherent to compositionality, this is rather attractive, and I do not intend to dispute the basic structure of the argument. However, I believe there are good reasons to suggest that this process is, at best, underspecified and in need of significant enrichment.

5.1.3 The limits of coercion

As an instruction manual for composing privative noun phrases, Partee's account is missing a few pages, even if the steps described are broadly accurate. As I see it, the main downside is that this story tells us nothing about what something like *fake gun* actually means - we have a rough map of how to get there, and some motivation for why it does not mean nothing at all, but no specifics for what it does end up meaning. We can buy into the general idea that it ends up being a subset of some expanded denotation of *gun*, but we have no idea what that expanded denotation is. How, for example, are we to distinguish *fake gun* from *fake knife*, or any other noun? In theory, both suffer from the same initial problem - they would be empty if the noun was taken literally - and they could both be resolved by expanding the noun's denotation to include *the same extra members*. If we expanded *gun* to include real guns and also fake watches, and expanded *knife* to include real knives and also fake watches, then **[[fake gun]] = [[fake knife]]**. This is obviously not what we obtain, but nothing in

Partee's account rules it out, or tells us what we are supposed to do instead to obtain a more reasonable result, since NVP only instructs us to include some positive extension of *fake* in the expanded denotation of the noun. Such nonsensical expansions would be just as acceptable as the actual ones, if we considered only violations of NVP and HPP. The natural response is to say that there is some additional principle of minimal modulation, which is violated if we expand the meaning of a term in an unnatural way. However, such a principle would be radically underdetermined - what information can we use to evaluate if a particular expansion is more or less natural than another option?

Another way of stating this limitation is quite simply that Partee provides no semantics for *fake*, only stating that it does not include any of the items in $\llbracket \mathbf{gun} \rrbracket = \llbracket \mathbf{real\ gun} \rrbracket$. Without a precise semantics for *fake*, it becomes impossible to predict the particular outcome of any given instance of modification with it.

This objection motivates the incorporation of enriched conceptual information about a noun into the compositional process, in order to have a predictive theory about what features a *fake gun* will have distinguishing it from a *fake knife*. Such an approach is offered by Del Pinal (2015), to which we will return in Section 4. First, I want to raise two other concerns for Partee's account, in the form of novel experimental data. Her account makes two additional assumptions:

1. The default interpretation of a noun is literal and includes only real instances.
2. A privative adjective will always uniformly incur NVP violations and induce coercion.

In the next sections, we will see that both of these can be seriously questioned on empirical grounds, which will require revision to both Partee and Del Pinal's approaches. However, I will argue, that rather than abandoning either approach, it is a unification of the two that derives the correct results.

5.2 Experiment 1: shifting the default

Experiment 1 tests the hypothesis that, in the absence of a modifier like *real* or *fake*, the default interpretation of a bare noun includes only the real instances of the noun. On the strictest interpretation of Partee’s proposal, which we might term the *narrow hypothesis*, this should be uniformly the case: only such modifiers trigger expansion and coerce the noun to include non-real instances. Fully in the other direction, the *wide hypothesis* predicts that nouns across the board refer to both real and fake instances by default. Any intermediate prediction would fall under the *context-sensitive hypothesis*, in which features of the linguistic or extra-linguistic context other than adjectival modification modulate the denotation of an unmodified noun. Here, we evaluate these hypotheses against two specific contextual variables which might effect the noun’s domain: the presence of real vs. fake instances of the same noun in the context, and the presence of comparison classes of other nouns in the context. Both are predicted to have an effect.

This experiment narrows in on the specific privative adjective *pretend*, for two reasons. First, *fake* potentially involves an element of deliberate deception, which poses a challenge for the visual display of images: any object which is unambiguously and immediately identifiable as not real, which is necessary for evaluating participant judgments of non-real images, may fail to count as *fake*. This confound is not present with *pretend*. Second, to lay the groundwork for follow-ups in acquisition work, *pretend* is preferred, since it is acquired significantly earlier than *fake* by children and is more appropriate to describe the kinds of objects with which they are familiar. Preempting acquisition work also influenced the choice of nouns and images, as described below. In either case, since the critical condition in this experiment is actually the one where no adjective is present, the influence of the particular adjective in other questions may be diminished.

5.2.1 Methods

Participants One hundred and twenty volunteers participated in the experiment over the Internet. Participants were recruited from Amazon’s Mechanical Turk system for human intelligence tasks via CloudResearch’s MTurk Toolkit. All reported that they spoke English as their first language, were born in the United States, and graduated from high school in the United States.

Stimuli Twelve nouns were chosen, half animate (*bear, cat, dog, duck, fish, pig*) and half inanimate (*airplane, car, guitar, house, phone, rocket*), and for each noun, five-to-ten images each of unambiguously real instances of the noun and unambiguously pretend instances (generally, toys) were selected. Images were photographs in naturalistic settings, taken from various online databases licensed for free use. These images were arranged into 3x3 grids mimicking the familiar Google CAPTCHA system; each question consisted of a single grid and a question about it, and grids were constructed according to the following criteria.

Grids varied along two parameters: within-noun alternatives and between-noun alternatives. For within-noun alternatives, a grid may have contained either only images of pretend objects (*no real comparison*), or also images of real objects (*real comparison*). For between-noun alternatives, a grid may have contained only images of the target noun (*zero alternatives*), both the target noun and one other noun (*one alternative*), or the target noun and four other nouns (*many alternatives*). Crossing these parameters gives us six grid types. Whenever we reference ‘half’ of an odd-numbered set of images, the division is random.

1. *Zero alternatives, real comparison*: all images of a single noun, half pretend and half real
2. *One alternative, all real*: images of two nouns, all real
3. *One alternative, real comparison*: images of two nouns, half pretend and half real
4. *One alternative, no real comparison*: images of two nouns, all pretend


5. *Many alternatives, real comparison*: images of five nouns, half pretend and half real
6. *Many alternatives, no real comparison*: images of five nouns, all pretend

Then, for each grid type, there were four possible question wordings. All questions asked the participant to select all of the images which met a specific condition. Instructions made clear that selecting any number of images was allowed, from zero to nine. For a target noun N , the questions for each grid would be:

1. Select all the images that contain a N .
2. Select all the images that contain a pretend N .
3. Select all the images that contain a [filler subjective adjective] N .
4. Select all the images that contain something pretend.

This results in 24 total question types, crossing the 4 wordings with the 6 grid types.

Select all images that contain a bear.



Select all the images that contain a pig.




Figure 1: Example question grids from Experiment 1. Displayed are conditions with zero alternatives (only one noun) and real comparison on the left, and many alternatives and real comparison on the right.


Procedure Questions were arranged into 24-question surveys, such that each survey contained one of each question type, counterbalanced via Latin Square. Since there were 12 nouns, each survey contained two questions where each noun was the target. Each participant received a random one of these surveys, and the order of questions within the survey was then additionally randomized. Participants were informed that they were participating in a linguistics study, with no additional specifics about the design or intent, and received the following instructions prior to the questions:

For each question, you will read a sentence and see nine images. Please select all the images that best match the condition provided in the sentence. You may select as many (possibly all of them) or as few images (possibly none of them) as you think match the conditions asked for.

The order of image presentation within any given grid was also randomized. No participant's data was excluded from the analysis, leading to an eventual count of 5 responses per unique question (question wording + grid).

5.2.2 Results

In analyzing the results, we will focus on the target question condition ‘Select all the images that contain a N’, since the effect we are interested in is the categorization of items as category members or not in the absence of a coercion-inducing adjective. Recall that the target condition of interest is whether participants are willing to apply the bare noun label to non-real instances of that noun, e.g. labeling an image containing a stuffed dog as containing a dog. Figure 2 presents the mean responses of participants on this metric; specifically, it presents the percentage of *pretend* nouns which were selected as instances of the bare noun, across grid types. If a given grid contained five pretend dogs, for example, and a participant applied the label *dog* to four of them, such a result would appear as a value of 80%.



We can see from Figure 2 that participants were, on average, willing to classify the majority of *pretend* items as members of the bare *noun* category. Recall that the factors hypothesized to affect this categorization were REALCOMPARISON - the presence of contrasting, real instances of the same noun in the visual context - and ALTERNATIVES - the presence of contrasting non-category members. The predicted effect for REALCOMPARISON is to make participants less likely to label pretend items with the bare noun, on the theory that the availability of a contrasting, more literally true target for the category label would cause participants to sharpen their criteria for category membership. The predicted effect for ALTERNATIVES is that a larger number of alternative noun categories would make participants more likely to label pretend items with the bare noun, on the theory that the e.g. pretend dogs are significantly better candidates for the *dog* label than any obviously non-dog items.

A 3 x 2 linear mixed effects model with the percentage of *pretend* items in a grid labeled with the bare noun as the dependent variable, and REALCOMPARISON (yes real competitor, no real competitor) and ALTERNATIVES (zero, one, many) as independent variables, with a

random effect for participant, was performed. A significant main effect of REALCOMPARISON was found ($\beta = -0.08, p < 0.05$), but no significant main effect of ALTERNATIVES ($\beta = 0.01, p = 0.35$) and no significant interaction ($\beta = 0.0006, p = 0.97$). Analyses of the simple main effect showed that significantly higher percentages (i.e., more pretend nouns labeled as nouns) for grids without real competitors than grids with ($t(570) = -2.94, p < 0.01$).

5.2.3 Discussion

The results of Experiment 1 suggest that people’s criteria for membership in the category denoted by a noun are variable with respect to certain features of the context, but not others. Specifically, there is variation in whether items which meet some criteria but not others of a literal interpretation of a noun - such as a stuffed dog, which has the perceptual features of a dog but not its biological essence - are sufficiently noun-like to be correct applications of the label. When the only available potential referents for a noun are of this type, participants were more willing to extend the noun label to such non-prototypical cases. When both non-prototypical and literal, prototypical category candidates are present in the context, it can push participants to choose a more narrow interpretation of the noun, excluding the pretend cases. This real competitor effect behaved as hypothesized.

By contrast, we saw no effect of the presence or number of alternative noun categories entirely, items which do not fall under the target noun label under any widened or narrowed interpretation. The absence of such an effect goes against the experimental hypothesis, which predicted that unambiguously non-noun items would raise the level to which maybe-noun items (the pretend cases) were grouped under the noun category.

It is important to keep in mind, in interpreting the main effect of REALCOMPARISON, that even in the presence of real competitors, participants were broadly very willing to count pretend items as category members despite the depressive effect of the competitor. The mean acceptance rate for pretend items in the presence of real competitors was 0.80, significantly but not dramatically lower than the acceptance rate without real competitors of

0.88. This suggests that participants defaulted to a more generously inclusive interpretation for nouns, and only sometimes were pushed to revise and narrow it to highlight a contrast in the contextually provided options.

Crucially, this process is occurring even in the absence of an adjective to induce coercion. On Partee's account, our default interpretation of nouns is literal, including only their real instances, and it is composition with a privative adjective which violates NVP and triggers a coercive expansion. These results suggest one of three different stories. First, it could be the case that our default interpretation of nouns is broad. This would bolster further Partee's argument that privatives are truly subsective, without requiring any of the NVP violation or coercion interactions. However, this interpretation raises serious questions about what it means to have a 'default' semantics for a noun that is non-literal. It seems much more natural to take the constrained, extensional meaning to be literal and expand it based on the extending of certain associated perceptual features, etc. The competing story, of starting with a broad denotation, struggles to identify exactly what pragmatic mechanisms are involved in the narrowing. How does one determine what literal, real dogs are from the broad set?

It could also be the case that our default interpretation is narrow, but that non-linguistic context is sufficient to expand it. This seems to be the most natural, and perhaps null hypothesis. There is no principled reason why constraints like NVP should only care about the application of predicates to other predicates; if the predicate *dog* has no positive extension in a context (because there are only pretend dogs) and there is a natural expansion to fix that, context should be sufficient to trigger such coercion. However, we run into the same problem as we did with *fake guns* - how does the semantic system know what to expand to? It's not the case that, if the context contains only real and stuffed cats, and you are asked to select all the dogs, you will incur a NVP violation and expand *dog* to include any cats; you will simply select none of the options. We will return to addressing this problem in the eventual discussion.

Before that, it is worth raising the third possibility from this experiment. Given that participants answered many questions with a bare noun target and also many questions that involved the adjective *pretend*, it is entirely conceivable that the *pretend noun* questions caused an adjectivally-induced coercive expansion of *noun*, and that such expansion lingered, such that participants were still primed to interpret it broadly even when *pretend* did not appear in a particular question. This seems plausible, although it might also predict question ordering effects which we do not find, and it would be equally plausible to suggest a pragmatic explanation in the other direction, where the presence of *pretend* in some but not all questions heightens the expectation of contrast for participants and primes them to interpret the noun differently in the absence of *pretend*.

5.3 Experiment 2: privative variation

Experiment 2 tests the hypothesis that so-called privative adjectives exhibit uniform behavior with respect to inference licensing. Partee’s account of privatives suggests that an adjective like *fake* will trigger the same coercive expansion regardless of the noun it combines with, and also that any privative adjective composes via the same process.

In reality, however, substantial variation is observed. Within a single adjective-noun combination, we might observe variation across contexts: a fake gun might be considered a gun in some respects and not a gun in others (Taylor, 2003, 96), and depending on the respects which the contextual demands draw attention to, our categorization may differ, as Experiment 1 also illustrated. With a single adjective, but across nouns, we find even more substantial and widespread variation: a *fake gun* may not be a *gun*, but *fake news* might be news, a *fake ID* is often an *ID*, and a *fake person* is usually a *person*. Cappelle et al. (2018) bring to bear the methods of computational distributional semantics to this question, and find that *fake* in nearly half its uses on Wikipedia behaves non-privatively, or at least not unambiguously privatively. The observation that ‘privative adjective’ may be too strong a term, and that we should instead speak of ‘privative uses’ of an adjective (Cappelle et al.,

2018, 7) is not a new one (Kamp, 1975; Boleda et al., 2013), nor is it anywhere near fatal for the Partee-style account: we simply need to restrict the coercion-expansion story to said privative uses. However, the challenge for the theory then becomes to determine when and how the composition distinguishes privative from non-privative uses.

Classically ‘privative’ adjectives are not the only ones which show variation. It is well known, for example, that constitutive material adjectives like *stone* and *plastic* alternate between intersective uses (*stone table*) and privative uses (*stone lion*) (Kamp and Partee, 1995; Coulson and Fauconnier, 1999; Oliver, 2014). Other ‘basic’ intersectives like color have less obvious examples, but phrases like *red meat* (not red in color) and *red pen* (not necessarily red on the outside) fail to validate the intersective inference from $AN(x)$ to $A(x)$.

There has been relatively little experimental investigation of these patterns of variation. Distributional semantic approaches like Cappelle et al. (2018) and Boleda et al. (2013) are more common, but behavioral investigations are limited. One of the only to focus on this question is Pavlick and Callison-Burch (2016), who had participants perform an entailment judgment task across sentences where a privative adjective has been inserted or deleted. So, participants were given a sentence containing a phrase like *fictitious company* and the same sentence without *fictitious*, and asked to rate how likely it is that the former sentence entails the latter, or vice versa. While rigid lexical categories would predict that entailment was consistent for substantive adjectives and never present for privatives, there was in fact significant variation, where inserting a privative adjective only produced contradiction 40% of the time, and deleting a privative adjective in fact produced entailment nearly 50% of the time. This demonstrates variation between adjective-noun pairings, which they take to suggest ‘how this $AN \rightarrow N$ inference, in the case of privative adjectives, often depends less on the adjective itself, and more on properties of the modified noun that are at issue in the given context.’ (Pavlick and Callison-Burch, 2016, 115) Crucially, however, this study asked about the entailment relationship between statements about the noun phrases, not about taxonomic categorizations of the privatively modified noun phrases as instances of the noun:

‘Rather than asking “Is this instance of AN an instance of N?” we ask “Is this statement that is true of AN also true of N?”’ (Pavlick and Callison-Burch, 2016, 116). Experiment 2 here is designed to explicitly test the categorization question.

5.3.1 Methods

Participants Forty volunteers participated in the experiment over the Internet. Participants were recruited from Amazon’s Mechanical Turk system for human intelligence tasks via CloudResearch’s MTurk Toolkit. All reported that they spoke English as their first language, were born in the United States, and graduated from high school in the United States.

Stimuli In order to evaluate the inferences from AN to N categorization made by participants, we asked them to answer questions of the form ‘Is an AN an N?’, varying the adjective and noun. We selected 11 adjectives: 5 canonically described as privative (*artificial, counterfeit, fake, false, mock*), 5 more traditionally intersective (*blue, red, round, square, plastic*), and 1 nonsubsective modal (*alleged*), and for each adjective selected 8 nouns with which combinations were attested. An obvious confound in any experiment of this design is the selection of nouns. For reasons elaborated in the discussion, we chose not to use a neutral selection mechanism like sampling of random or high-frequency bigrams from corpora. Here, we manually selected nouns which were likely to show variance with each adjective due to differences in their semantic features. In addition, because adjectives in general and privatives especially show restricted ranges of nouns which they can modify, it is not possible to evaluate all adjectives with respect to the same list of nouns, and so a unique list had to be selected for each adjective. With this understood, it will be appropriate to evaluate this particular experiment as (i) proof of concept for this method of identifying variation and (ii) illustration that at least some variation does exist, which will be expanded in the discussion. Less stipulative and significantly larger selections of both adjectives and nouns

will be a desirable extension of this work going forward.

For the present study, regardless, this procedure resulted in 88 unique A-N pairs. For each A-N pair, we created four questions:

- NOMINAL: ...an N? (e.g. Is a fake gun a gun?)
- SUPERSET: ...a [superset of N]? (e.g. Is an alleged thief a criminal?)
- ACCIDENTAL: ...[some property that Ns might have]? (e.g. Is a plastic flower yellow?)
- CONTRADICTION: ...[some label that cannot apply to Ns]? (e.g. Is a blue pencil a cat?)

These additional question categories function as a baseline for comparison to the target questions, which are in the NOMINAL category. Example questions in the NOMINAL condition, as participants saw, are shown in Figure 3 below.

Is an artificial tree a **tree**?

Definitely not	Probably not	Unsure	Probably yes	Definitely yes
----------------	--------------	--------	--------------	----------------

Is a mock election an **election**?

Definitely not	Probably not	Unsure	Probably yes	Definitely yes
----------------	--------------	--------	--------------	----------------

Figure 3: Sample questions for Experiment 2

Procedure The resulting 352 questions were divided into 8 surveys of 44 questions each. Surveys were counterbalanced via Latin square such that each survey contained one of each four question types per eleven Ns. 40 participants took one survey each, resulting in 5 responses/survey = 5 responses/question.

Participants were instructed to respond on a five-point Likert scale, from "Definitely not" to "Definitely yes". Participants were excluded from analysis if they endorsed (answered 'Probably yes' or 'Definitely yes' to) a CONTRADICTION in more than 10% of their total answers. Two participants were excluded according to this criterion.

5.3.2 Results

Between conditions, participants' ratings conformed to the expected patterns: contradictions were rarely endorsed, and all other conditions received moderate levels of endorsement. Specifically, ratings in the target NOMINAL condition ($M = 3.80, SD = 1.4$) were significantly higher ($t(565) = -35.4, p < 0.001$) than the CONTRADICTION condition ($M = 1.2, SD = 0.57$), and showed significantly more variation. Ratings in the NOMINAL condition were also significantly higher ($t(710) = -8.1, p < 0.001$) than the ACCIDENTAL condition ($M = 3.2, SD = 0.86$), though both means were squarely in the middle of the scale. By contrast, ratings in the NOMINAL condition were not significantly different ($t(835) = -0.54, p = 0.6$) than in the SUPERSET condition ($M = 3.75, SD = 1.3$), and the variance of ratings in the two conditions was also not significantly different ($F = 0.89, p = 0.22$). This lack of difference is to be expected on the assumption that determining membership in the nominal category also determines membership in a superset of the nominal category ('Is an artificial tree a plant?' = 'Is an artificial tree a tree?' + 'Is a tree a plant?').

We will focus on reporting results from the target condition NOMINAL. Recall that the hypothesis we are interested in evaluating is whether the inferential behavior of an AN pair (namely, whether it is counted as an N) is determined by the class of the adjective (privative vs. intersective). On first pass, it might appear to confirm the usefulness of the categories: A-N pairs in the NOMINAL condition are rated significantly higher ($t(359) = -7.44, p < 0.001$) when the adjective is intersective ($M = 4.34, SD = 1.15$) than when the adjective is privative ($M = 3.31, SD = 1.54$). However, the results from within the privative class raise suspicion. Based on the canonical theory of adjective classes, we would expect the ratings in

the privative conditions to be near floor, as what it means to be a privative adjective is to license the inference AN \rightarrow !N. Instead, we see the mean ratings at 3.31, just above the center, and significantly higher than chance ($t(194) = 30.1, p < 0.001$). We also see that the privative class has significantly higher variance than the intersective class ($F = 0.56, p < 0.001$). This suggests that, while in broad strokes the privatives license the AN \rightarrow N inference at a lower rate than the intersectives, the picture is more complex at a more granular level.

To examine the finer-grained results, we turn to a by-adjective analysis. Figure 4 plots the data across all adjectives in the NOMINAL condition, where the large circle represents the mean.




Figure 4: Ratings in the NOMINAL condition, between-adjective

While, again, we see a general tendency for privative adjectives (in blue) to have lower means than intersectives (in red), such tendencies are not categorical. For the given set of nouns, the intersective adjective *plastic* here behaves as if it were privative, and the

privative *artificial* patterns more closely to the intersectives. By-mean analysis hides a significant amount of variation - for every privative adjective, there are responses at the top and bottom and throughout the scale.

To examine the usefulness of categorizing individual adjectives according to their inference-licensing properties, a natural next question to ask is, to what degree can the inferential property of a given AN phrase (i.e., does $AN \rightarrow N$?) be determined from the A alone? Recall from the introduction that on the simple classical story, the nature of the noun is not a factor in determining inferential behavior, which is licensed entirely by the adjective. To answer this question, we can first binarize the response data, by eliminating response means of exactly 3 and coding all values above 3 as licensing the intersective inference, and all values below 3 as licensing the privative inference. We then can construct a logistic model with the binarized response as the dependent predicted variable and the individual adjective as the independent predictor variable.


On this model, only three adjectives reach significance as predictors of the inferential behavior of AN phrases containing them: *fake* ($\beta = -1.31, SD = 0.51, p = 0.01$), *red* ($\beta = 2.45, SD = 1.08, p = 0.02$), and *round* ($\beta = 2.45, SD = 1.08, p = 0.02$), with *blue* approaching marginal significance ($\beta = 1.35, SD = 0.71, p = 0.06$). Even *fake*, however, as the best predictor of all adjectives, only receives responses below 3 on 59.0% of trials. That is, if you tried to use the adjective alone as a cue to guess the inferential behavior of the full AN phrase, your best guess will be incorrect four out of ten times. Given the variance observable from Figure 4, this should be unsurprising; from the perspective of a category-based theory of inferences, this should be catastrophic.

The high within-adjective variance suggests the need to dig down further. The following figures break down the variance within each adjective, showing response ratings for each noun. At this level of granularity, sample sizes are too small for any meaningful statistical comparison. The takeaway, however, should be clear to the eye test - adjective alone is a woefully insufficient determinant of inferential behavior, and nearly all adjectives show

significant variation across nouns.


Responses to 'Is a fake N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a Counterfeit N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a mock N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a false N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is an artificial N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is an alleged N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a plastic N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a square N an N?'

Where 1=Definitely Not, to 5=Definitely Yes




Responses to 'Is a round N an N?'


Where 1=Definitely Not, to 5=Definitely Yes



Responses to 'Is a blue N an N?'

Where 1=Definitely Not, to 5=Definitely Yes





5.3.3 Discussion

The emergent hypothesis The results of Experiment 2 appear to validate the predictions of Pavlick and Callison-Burch (2016), in that neither privative nor intersective adjectives are entirely well-behaved according to their category.²¹ Privative adjectives show significantly wider variance than the intersectives, though we can likely predict this is due to the selected intersectives - given the behavior of *plastic*, it should be easy to assume that *wooden* and *metal* would show a similar pattern, while choosing *green* and *yellow* would reinforce the intersective patterning. We'll return to an argument for why colors and materials pattern this way in the next section, but for now the relevant observation is that neither class of adjective allows comfortable prediction of inferences. To determine if a modified noun phrase is going to license inferences to the unmodified noun, it is necessary to know both the semantic content of the adjective and the noun. In this way, we can start to consider privativity as

²¹In much of what follows, I'll be using the terms *privative adjectives* and *intersective adjectives* as shorthand for *those adjectives which, in canonical theory, have been labeled privative/intersective*, which would be even more laborious to read than write. We'll eventually return to a way to perhaps use these terms to usefully label a type of semantic property, but even when we do so, the intention won't be to reinstate the distinction, but again as a shorthand.

an emergent property of specific instances of composition, rather than a lexical property of specific adjectives. Call this the emergent hypothesis.

The task of the emergent hypothesis, then, is to explain why certain combinations are privative, and in doing so, explain what the actual semantic contribution of a privative adjective is. Informally, the data above suggests the following proposal: nouns have a variety of semantic features, some of which are more important than others. By important, I mean something like ‘core to determining category membership’. Let’s consider as a first example the case of *counterfeit*. One thing that we might know about a noun, either as part of the noun’s actual essential, definitional character or as part of our associated world knowledge, is the means by which it comes into existence. For something like a *dollar*, this origination process is completely core to what it means to be a dollar - fiat currency has value only according to the authority that minted it, it has a very distinct official process of creation, and anything that didn’t come into being in that particular way cannot perform the functions which a dollar is supposed to perform. For something like a *watch*, we certainly may have our default assumptions about how a watch comes into being, the process of crafting and design, and the various people who might make one. But it doesn’t seem to matter *for being a watch* that the object takes a particular one of those paths, so long as it ends up being something that functions to tell time and is appropriate to sit on a wrist, perhaps. Whether a wrist-mounted time-telling device is analog or digital, made by a master craftsman at Rolex or 3D-printed in a basement, or anything else, the outcome is the same - so long as it comes out as a watch, it’s a watch.

We can model *counterfeit*, then, as contributing a deprivation of such an ‘origination’ feature. For any noun *N*, the set of *counterfeit Ns* will contain those objects which are exactly like *Ns* in all respects other than their origin. And for something to be a *dollar*, it is required that that thing has the distinct origin of a dollar. So the set of *counterfeit dollars* will be normally disjoint from the set of *dollars*, and we obtain privative behavior for the adjective. But for something to be a *watch*, there can already be significant variation

in the origin process, which would not be uniform across the set of *watches* even before modification. So, the set of *counterfeit watches* will be those objects which satisfy all other watch qualities but were not made in the normal way that you might expect a watch to be made in the context. Such objects, by nature of the semantics of *watch*, are already in the set **[[watch]]**. Thus, *counterfeit* is subsective in this instance.

What about our intermediately rated cases, like *drug* or *contract*? I posit that these effects arise due to variation in exactly which part of the origination process is implicated by the adjective. All that we are saying *counterfeit* contributes is, in a sense, a negative requirement - don't have the prototypical genesis. Perhaps the signatures were forged - in this case, a counterfeit contract likely isn't a contract. But it could also be the case that one party lied about their identity or engaged in some deception in the process of contract creation, but the actual text is valid and even though you didn't enter into an agreement with the party you intended to, you did enter into some agreement. Or, perhaps, people are alternating between understanding *contract* as denoting the physical object (which would cause intersective interpretation) or the intellectual content (which would cause privative interpretation), as is well-known for such nouns (Asher, 2011). Similarly, a *drug* could be counterfeited in being made by the wrong company and so it's a cheaper, less effective version, but still meets the chemical qualification of being a drug. Or, it could be oregano rather than marijuana, in which case it isn't a drug at all. Alternations in the noun's meaning, then, also induce alternations in the inferential behavior with the adjective - another wrinkle in the categorial hypothesis.

What happens if a noun doesn't have any kind of origin encoded at all? There are obviously objects which don't have as any part of their meaning anything about who made them; something like a *boulder* certainly has normal ways in which it might come into being, but it's unlikely that any speaker of the language is considering those as part of their use of the term. We would expect the phrase *counterfeit boulder* to be infelicitous, and it is, absent some substantial supporting context which adds origination information not normally

present to the word's meaning. We can attribute this to the Non-Vacuity Principle: if the contribution of *counterfeit* is to remove a semantic feature that was never present in the first place, it will be ruled out for the same reason that *real* applied to a literal denotation would.

We see that the same type of story can be constructed for the intersective *plastic* - in short, natural kinds like flowers and diamonds need very much to be made of flower and diamond, while artifacts like chairs and spoons denote only their function as a tool and are amenable to any material (though they will still have some default material interpretation, given by world knowledge).

How might the same data be accounted for on a theory where adjectives are lexically labeled as privative in some way? A basic, stipulative solution would be to posit two lexical entries for privative adjectives, homophonous but allosemous: one with whatever semantic piece triggers privative inference, and one without. However, this type of analysis is rather obviously undesirable. First, it would require stipulating selectional restrictions, such that nouns like *diamond* could only be modified by lexically privative *counterfeit*, nouns like *watch* could only be modified by lexically intersective *counterfeit*, and nouns like *contract* (which show ambiguity with respect to judgments of privation) can take either. Such a stipulation would seem to entirely miss the generalization, one which appears to crucially involve elements of the compositional process in a way that would be undesirably duplicated if we simply encoded a list of possible adjective-noun pairings. In addition, in order for such an approach to be extended to the classically intersective adjectives, we would need to similarly posit underlying lexical ambiguity, such that material adjectives and potentially significantly more classes all have intersective and nonintersective variants. Finally, such a solution would be unable to account for the context-sensitivity illustrated by Experiment 1, unless we also posited an underlying ambiguity in the meaning of nouns like *dog* (with a narrow version and a wide version) and our list of allowed adjective-noun pairs was specified for such nominal variations as well. This quickly becomes untenable.²²

²²In theory, one could also preserve the principle of compositionality and such a lexical-categories account by appealing to the other determinant of compositional meaning: the way the constituents are combined.

Cherry-picking arguments Now, caveating any of this discussion, Experiment 2 highlights a particular difficulty of studying adjective variation in a rigorous, quantitative manner. Namely, adjectives display not only variation within the nouns with which they can combine, but also have distinct selectional restrictions (to sloppily borrow a syntactic term). It is not possible to test every adjective in a study on the same set of nouns, given that they will frequently be infelicitous together. Privative adjectives are particularly pernicious examples of this, with highly restricted behavior. The range of nouns which can be modified by *mock*, for example, appears to have almost no overlap with the range modifiable by *counterfeit*. So, there is inevitably going to be some degree of inconsistency. How much of a problem is this?

From the perspective of one defending a lexical perspective on canonical adjective categories, the most obvious objection to raise is, ‘You’ve hand-selected nouns to maximize the likelihood you get these high-variance graphs’. And that is absolutely true. Rather than use any experimenter-independent, quantitative method for selecting the target nouns, these were cherry-picked precisely to create this effect. However, it is unclear how this normally suspect design choice actually impacts any of the theoretical conclusions here. There are two reasons for this.

First, the burden of proof for the strict, lexicalized nature of adjective inferences is, I think, on the theorist defending them. Our null hypothesis should not involve the existence of distinct categories with semantic reality. In that sense, we could view Experiment 2 as merely the pilot demonstrating that if one would like to argue for such a categorization, they have some empirical work to do themselves. Another, perhaps more forceful way of putting this point is that the bar for *falsifying* such a theory is extremely low. On its strictest interpretation, the existence of *any* variance whatsoever is something which needs explanation as a problematic exception. Whether such counterexamples were happened upon

Such an approach would say that adjectives are always privative or intersective, but some invisible structural variation in their syntactic relationship to the noun triggers the variance we see here. I won’t try to sketch or argue against such an account here, as it seems to require significantly more empirical work and theoretical argument than could be easily summarised or imagined, but it is in principle possible, and even desirable.

in corpus scraping or a norming study, or selected manually and then verified experimentally, seems irrelevant for the theoretical point they make: variance exists. And it exists for more than one or two lexical items, and the phrases it exists with don't appear to be idiomatic or metaphorical or easily captured with some special dimension of meaning. Variance exists in abundantly normal compositional contexts. It is difficult to see how the charge of cherry-picking those contexts blunts the force of this observation.

Second, no obvious alternative is distinctly more attractive. We've seen that we cannot select a random wordlist and use the same list across adjectives, given their restricted range of application. What about by-adjective selection from a corpus? Let's imagine that, to keep our hands off the data for as long as possible, we instead take the top n most frequent bigrams containing a given adjective from a corpus, or a random subset of those bigrams, and perform the same experiment using them. What are we likely to see? Let's set aside the outcome where we reproduce the pattern observed here, which would obviously not be useful as an objection to our manipulations. Say instead that we find significantly reduced variation, with *fake* behaving strictly privatively in some near-ceiling number of cases, in a way that tracks with a categorial hypothesis. Could we take this as support for the categorial hypothesis, over the position sketched above? I'm not so sure.

If we take the semantic contribution of a privative adjective to be deprivation of a certain semantic feature, as I argue, we should expect to find precisely such a broadly categorial pattern in corpus data, for the same reason that we should not expect many naturally occurring examples of *colorless green ideas* outside of linguistics texts. If intersective interpretations of privative adjectives arise when the semantic feature targeted by the adjective is weakly specified or unimportant to the noun, then such compositions run the risk of being vacuous. The less important that such a feature is for the noun's meaning, the more likely it is that an operation that deprives it of that feature will be contributing close to nothing to the semantics of the resulting noun phrase. By contrast, the more that a noun 'cares' about a given semantic feature, the more likely the composition is to be interpreted privatively, and

the more semantic work the adjective is doing in distinguishing the modified noun phrase from the bare noun. Thus, assuming that we should be able to see reflexes of the Non-Vacuity Principle in action across frequency of constructions, we would expect that privative adjectives, from a bird's eye view, appear to be so because there is pragmatic pressure for them to preferentially compose with nouns with which they do behave privatively.

So, even in the case where some large percentage of a wider, random sample of nouns patterns according to the categorial hypothesis, we have good reason to still conclude against it. Any amount of variance at all, beyond idioms, would remain a problem, and both the categorial and emergent hypotheses predict a large-scale preference for privative behavior. From these frequency arguments, we can account for both the existence of variation and the presence of significant trends towards privativity simultaneously. In the next section, we can spell out the semantic machinery necessary to implement these intuitions.

5.4 Dual Content Semantics

This section introduces the framework of Dual Content Semantics (Del Pinal 2015, 2018), a theory of semantics that allows compositional operations to access and manipulate certain elements of the conceptual structure of nouns, enriching the information comprising lexical entries while also maintaining a standard extensional semantic component of denotations. Much of the basic structure of this theory

5.4.1 Motivations for lexical enrichment

Before diving into the analysis, let's quickly review the goals for a compositional system. We need a system which allows for modulation of a noun's meaning to be induced by both linguistic and non-linguistic context, shrinking or widening the boundaries to allow for different standards of category membership, as per the Experiment 1 results. We need it to be able to allow for all of the relevant kinds of variation, within-adjective and across-adjective, that we observe in Experiment 2. It should preserve the basic structure of Partee's argument,

in terms of reference to simple primitive principles of interpretation, rather than stipulating more complex motivating principles. The overarching goal is to have a system which provides us with a compositional blueprint for the output of these modifications: we want to be able to predict the membership of sets like $\llbracket \text{fake gun} \rrbracket$, and whether such membership will be intersective with $\llbracket \text{gun} \rrbracket$, from the individual components.

This problem of being unable to predict the outcome is even more pernicious, in fact, than just quibbling for more specification. Because expansion is only triggered by a violation of Non-Vacuity, namely seeing that the result set would be empty absent some expansion, extensionally this requires checking the set of guns and determining that there are no fake guns in it. But how does this checking process proceed? It must involve evaluating the members of the gun set with respect to the fake gun set, and seeing that none of them are members of the latter. But this creates an obvious problem: to perform such an operation as is required to trigger a Non-Vacuity violation, we already need to have access to the *fake gun* set, to determine that none of its members are in the *gun* set! We rapidly enter circularity: to get the *fake gun* set we need, we must compose *fake* and *gun*, which will either result in the empty set or, by hypothesis, trigger Non-Vacuity, and so on.

To resolve this problem, what is needed is some sort of predictive mechanism which can look only at the meanings of the adjective and the noun in isolation and be able to determine that the outcome will be contradictory. Thus, to preserve anything like the Partee story, we need a systematic way of analyzing the meanings of nouns.

Dual Content Semantics (Del Pinal, 2015) preserves a system of composition nearly identical to classical function application, with minimal modifications, by adopting the assumption that the default lexical entries for nouns are notably more complex than in prior systems. On this view, common nouns have a binary semantic structure consisting of their extensional meaning (E-structure) and their conceptual meaning (C-structure). E-structure is the atomic extension-determining component, of the form $\lambda x.\text{STONE}(x)$. C-structure does not determine the extension of a noun, but instead consists of ‘representations of perceptual

features, functional features and genealogical features related to [the noun]’ (Del Pinal, 2015, 4). These take the form of a qualia structure in the style of Pustejovsky (1991). While only E-structure determines the extension of the noun’s denotation, C-structure, it is argued, is a necessary component of speaker’s linguistic competence in their ability to correctly identify members of a kind and use the term dynamically and productively in different contexts; it might be considered an instruction manual for correct and useful application of the linguistic term.

This non-atomistic theory of linguistic meaning may seem like a radical departure from the denotations formal semanticists are used to working with, but objections to purely definitional denotations have a rich history in the philosophy of language and mind, beginning with Putnam (1970), whose criticisms formed the basis for a number of holistic theories of linguistic meaning.²³ Del Pinal’s adaptation of the framework takes Pustejovsky’s insights about the structure of the lexicon and endeavors to apply them to privative adjectives as a prime example of the kinds of tricky compositional meaning and pragmatic influence on composition that motivated such theories. Motivating this shift is a desire to unify the mechanical power of formal semantic frameworks with a ‘psychologically realistic account of lexical semantics’ (Del Pinal, 2015).

Privative adjectives provide the ideal test case for demonstrating not only its plausibility but its necessity. ‘They provide us with a concrete example of a kind of (literal, non-pragmatic) general linguistic competence - the capacity to systematically and productively use privative NPs - that seems to essentially involve the C-structure of terms.’ (Del Pinal 2015:7) As discussed in the last section, the function of privative adjectives is naturally described as the deprivation of certain conceptual features of a noun’s semantic content - because they modify these features in a regular and productive way, rather than the noun’s extension, they appear to provide evidence that compositional interaction with those features

²³Of course, a rich philosophical history also brings with it a deep and heated debate. Holism is not without its critics. Here, we focus on the consequences for privativity, ignoring both Putnam’s arguments against atomism and others’ arguments against dual content-style representations, but an interested reader is pointed to Fodor and Lepore (1992) for these objections, and to Bilgrami (1998) for a holistic response.

is possible. Specifically, the fact that this composition is sensitive to aspects of syntactic structure motivates giving it a compositional rather than purely pragmatic treatment. In addition, the fact that the primary contribution of the adjective is negative, the removal of a feature, suggests that the typology of privative adjectives may constitute evidence for which particular features of concepts are lexically represented. If we conclude that a specific adjective - in particular, a monomorphemic one, not build out of other properties with something like *un* or *non* - targets a specific feature, and adjectives with that meaning recur cross-linguistically with some regularity²⁴ (to the exclusion of other possible meanings that do not), it would be reasonable to conclude that the targeted feature is a regular component of nominal conceptual structures. This contrasts with adjectives that add descriptive features or create deprivation through composite means (by constructing an ad hoc new featural deprivation, again through something like *un-*), since neither of those meanings would imply anything primitive about the features that are already represented on the noun. To see how this logic might be implemented, let's turn to the details of the Dual Content semantic system.

5.4.2 The semantic framework

A sample Dual Content entry for a noun, *gun*, is given below.

- (242) **[[gun]]** =
 E-structure: $\lambda x. \text{GUN}(x)$
 C-structure:
 CONSTITUTIVE: $\lambda x. \text{PARTS-GUN}(x)$
 FORMAL: $\lambda x. \text{PERCEPTUAL-GUN}(x)$
 TELIC: $\lambda x. \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$
 AGENTIVE: $\lambda x. \exists e_1[\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)])]$

Not all of the precise semantics of the C-structure elements in (242) will come into play

²⁴Of course, a complete version of this argument depends on a typological survey of the range of privative meanings across languages that has not been conducted here; this is a potentially fascinating avenue for future work.

here; what is necessary is that the C-structure of *gun* encodes that it is composed of gun parts, has the perceptual form of a gun, is generally used in shooting events, and was made with the goal to be used in shooting events. Now that lexical entries are decomposed in this format, we also introduce operations which are able to access specific components of a lexical entry’s meaning.

- (243) *Dimension operators*: partial functions from the meaning of terms into their respective C-structure denotations, namely, constitutive, formal, telic, and agentive. The qualia functions are Q_C , Q_F , Q_T , Q_A . For example, using the denotation for *gun* in (242): $Q_C(\llbracket \mathbf{gun} \rrbracket) = \lambda x. \text{PARTS-GUN}(x)$
- (244) *Core enrichment operators*: partial functions from the meaning of terms into certain combinations of their E-structure and C-structure. The operators are C , T , A , and F . For example, $T(\llbracket \mathbf{gun} \rrbracket) = \lambda x. \text{GUN}(x) \wedge \text{GEN } e[\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)]$

Adjectives can have a Dual Content structure as well. Intersective adjectives could theoretically be represented in a simpler manner, perhaps with only E-structure, such as $\llbracket \mathbf{red} \rrbracket = \lambda D_C. \lambda x. D_C(x) \wedge \text{RED}(x)$, or even more simply as type $\langle e, t \rangle$ and composing with nouns using Predicate Modification. Privative adjectives, then, make use of these qualia functions, by operating over different elements of the noun’s C-structure, preserving some in the resulting NP denotation and rejecting others. I will skip over much of Del Pinal’s exposition and argument for how he arrives at this eventual lexical entry for *fake*, and simply present the final version. Here, D_C is the domain of ‘ordered sets of the E-structure and C-structure of common Ns’ (Del Pinal, 2015, 14).

- (245) $\llbracket \mathbf{fake} \rrbracket =$
 E-structure: $\lambda D_C. [\lambda x. \neg Q_E(D_C)(X) \wedge \neg Q_A(D_C)(x)$
 $\quad \wedge \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]$
 C-structure:
 CONSTITUTIVE: $\lambda D_C. Q_C(D_C)$
 FORMAL: $\lambda D_C. Q_F(D_C)$
 TELIC: $\lambda D_C. \neg Q_T(D_C)$
 AGENTIVE: $\lambda D_C. [\lambda x. \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]$

With this entry for *fake*, the resulting referent will have the constitutive and formal qualia of the noun, will not have the telic or agentive qualia, and will have a new agentive qualia suggesting that the referent was made with the goal of having the same formal qualia as the noun (i.e. being a convincing fake). In this formalism, the negation of a qualia function indicates that the function does not apply to that entity; perhaps a notation like $Q_T(D_C) = 0$ would be more natural, but I will preserve Del Pinal’s notation here to avoid confusion with the original work. To compose this complex modifier with our complex noun in (242), we will need a more complex notion of function application:

- (246) Dual Content Function Application (FA^{DC}):
 If α is a branching node, $\{\beta, \gamma\}$ is the set of α ’s daughters, and $\llbracket \beta \rrbracket_E$ is a function whose domain contains $\llbracket \gamma \rrbracket$, then $\llbracket \alpha \rrbracket_E = \llbracket \beta \rrbracket_E(\llbracket \gamma \rrbracket)$ and $\llbracket \alpha \rrbracket_C = \langle Q_C(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_F(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_T(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket), Q_A(\llbracket \beta \rrbracket)(\llbracket \gamma \rrbracket) \rangle$.

Per (246), the E-structure of a modifier takes in the E-structure of the noun as its argument, as does each C-structure take in its corresponding C-structure argument. Then by FA^{DC} , the result of applying **fake** in (245) to **gun** in (242) is:

- (247) **fake gun** =
 E-structure: $\lambda x. \neg Q_E(D_C)(\llbracket \mathbf{gun} \rrbracket)(x) \wedge \neg Q_A(\llbracket \mathbf{gun} \rrbracket)(x)$
 $\wedge \exists e_2[\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(\llbracket \mathbf{gun} \rrbracket)(x))]$
 C-structure:
 CONSTITUTIVE: $Q_C(\llbracket \mathbf{gun} \rrbracket)$
 FORMAL: $Q_F(\llbracket \mathbf{gun} \rrbracket)$
 TELIC: $\neg Q_T(\llbracket \mathbf{gun} \rrbracket)$
 AGENTIVE: $\lambda x. \exists e_2[\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(\llbracket \mathbf{gun} \rrbracket)(x))]$

Thus, we get a class of entities which are not guns, do not have the origins of guns, and were made to appear as if they were guns, but do not have the purpose of guns (i.e. are generally not used in shooting events).

In addition to the Dual Content structure and the particular semantics needed for privatives, Del Pinal (2018) also introduces a modulation operation following Recanati (2010) and Pagin and Pelletier (2007), constrained relative to the operations available in Dual Content:

- (248) $\llbracket \mathbf{e} \rrbracket_{ME}^c = \text{mod}_l(e, c)(\llbracket \mathbf{e} \rrbracket_M^c)$

The function of the modulator mod_l is to select a lexical modulation function l which ‘takes the full meaning of e and returns an E-structure that is relevant for e in c .’ (Del Pinal, 2018, 179) Unlike in preceding work, e.g. Recanati (2010), the modulation function is rather severely constrained. The only available lexical modulation functions draw from the pool of dimension operators and core enrichment operators. That is, the only way to modulate a lexical item’s meaning is through some manipulation of its C-structure, replacing some elements of the E-structure with C-structural qualia.

This is the basic architecture of Dual Content Semantics: split lexical entries which involve both an extensional component and a structured conceptual component, and compositional and pragmatic operations for manipulating the structured conceptual features and allowing them to interact with extensional semantics. This satisfies the primary requirement of a theory of privative modification: allowing adjectives to modify conceptual features. Both to account for our experimental results in this chapter, however, and to synchronize with the syntactic components of the modification theory developed so far, I want to propose a number of revisions to this theory. The first couple revisions will be somewhat theory-internal, different ways to do what DCS is doing. I’ll then move to a more fundamental adaptation of the theory, which will take its inspiration from DCS at its core but change its basic principles to make them cogent with our syntactic architecture.

5.4.3 Making privativity emergent

The first revision that’s necessary is enabling privativity to be an emergent compositional phenomenon. Taken together, the results of the two experiments suggest that our semantic implementation of privativity should be capable of capturing a range of variations in meaning. In particular, our system should derive privation via the context-sensitive interaction of conceptual features associated with the adjective and with the noun. What does this mean? Even a system like DCS, which allows for variation in the output meaning of a privative adjective’s composition with a noun, based on the semantic features of the noun, does not

allow for variation in whether the output is privative (when the adjective and noun are fixed). This is because privation is hard-coded into the lexical entry for an adjective: part of the meaning of *fake* is the negation of the extensional structure of its nominal argument. That this assumption may be too strong is not an entirely novel observation; Del Pinal considers it in a footnote, but remained relatively agnostic on the choice. Experiment 2 here provides evidence against this decision: hard-coding privation into the meaning of an adjective will cause even a dual-content system to undergenerate and fail to produce the empirically attested range of variation.

Luckily, the basic architecture of DCS provides us the material for a more compositionally parsimonious analysis. We just need to strip away some of the extraneous pieces to begin: namely, as suggested above, the lexical entries for privative adjectives need to be revised to eliminate the extensional-negation component. So, the semantic contribution of a privative adjective like *fake* consists only in the privation of a particular quale or set of qualia. Most simply, that just involves removing the $\neg Q_E(D_C)(X)$ term from (245), leaving:

$$\begin{aligned}
 (249) \quad \llbracket \mathbf{fake} \rrbracket = & \\
 \text{E-structure: } & \lambda D_C. [\lambda x. \neg Q_A(D_C)(x) \\
 & \wedge \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]] \\
 \text{C-structure: } & \\
 \text{CONSTITUTIVE: } & \lambda D_C. Q_C(D_C) \\
 \text{FORMAL: } & \lambda D_C. Q_F(D_C) \\
 \text{TELIC: } & \lambda D_C. \neg Q_T(D_C) \\
 \text{AGENTIVE: } & \lambda D_C. [\lambda x. \exists e_2 [\text{MAKING}(e_2) \wedge \text{GOAL}(e_2, Q_F(D_C)(x))]]
 \end{aligned}$$

Without the extensional-negation component, of course, we immediately run into the opposite undergeneration problem: we now have no way to generate nonsubsective meanings at all. If nonsubsectivity is no longer provided directly by our adjectives (and certainly should not be encoded in our nouns, which would render most nominal meanings absurd), and we wish to adhere to strict compositionality, then the only remaining option is to derive that component of the phrase's meaning from the way the parts are combined. So, we need a principle of the compositional process that tells us when to interpret a particular instance

of composition as privative. The intuition given to us by Experiment 2 is that privative meaning occurs when the semantic feature which the adjective removes from the noun is a particularly central feature to the nominal concept. At a first pass, we can suggest something like (250).

- (250) An A-N compound should be interpreted privatively when any qualia Q negated by A is central to the meaning of N.

This seems intuitively correct, but from a formal standpoint, offers us relatively little, because it shifts the explanatory burden onto the phrase ‘central to the meaning of’, which is a relation that dual content representations do not encode. But what we do have access to is *mod_l* and the qualia-manipulation operations it can provide. The availability of dimension operators (which select only a particular quale of C-structure) and core enrichment operators (which select a quale and the E-structure) give us all of the power we need to capture variation. Modulation can select any combination of the noun’s qualia as extension-determining, induced by either context or linguistic factors. We can assume that some nouns come prepackaged with a preferential or default weighting, as we described before for things like *dollar* - now, we can say that the default interpretation of *dollar* involves core enrichment of its AGENTIVE feature.²⁵ Privativity, then, arises when a quale which has been ‘uploaded’ (to use Del Pinal’s terminology) to E-structure, through either a dimension or core enrichment operator, is negated by an adjective.

What cannot be allowed is free contextual modulation of *which qualia an adjective targets*. Such an attempt is pursued by Guerrini (2018), who argues for a representation of nouns as collections of prototypical features, any of which may be made relevant by the context at any time. *Fake*, then, operates as a privative over whichever features are relevant. This makes two incorrect predictions. First, it predicts that *fake* will always be privative, since no matter which features of the prototype are relevant for determining extension, it will be

²⁵Whether the ‘underlying’ lexical entry for *dollar* has such an enrichment, or it comes with some kind of default modulation that occurs in any standard context, is going to be a completely immaterial choice for this process - either way, before we get to composition with the adjective, it will look the same.

those which are negated. This is obviously an undesirable outcome.

Second, it predicts - or, rather, states up front - that any of the features associated with a noun's prototype may be contextually relevant, and thus that *fake* is free to modify any of those features so long as the context makes it appropriate. This appears to overgenerate: even if a context makes one and only one feature of a noun relevant, it is still not the case that that feature can be targeted by *fake*. Imagine that we are at a metalworking shop and have just recently purchased a collection of guns from a customer. Our only interest in purchasing the guns was to melt them down; the metals that guns are always made of is valuable to us for some other project, so it's irrelevant whether the guns are broken or function. When going through the pile and melting each gun down, we find that one doesn't melt properly, and upon closer inspection is actually made of a futuristic, high-quality plastic. It functions perfectly well, with the same internal mechanisms as a standard gun, but has none of the same materials. The manufacturer didn't want their guns to stand out, for some reason - perhaps the army didn't want to reveal their new technology - and so the plastic is designed and painted deceptively to look precisely like a standard metal gun. In this context, would something like 'They sold us a fake gun!' be an appropriate utterance? My intuition is that no, it would not be. *Fake gun* is not an appropriate label for this item, even though it lacks the most contextually relevant feature of guns and even has a layer of deception to look like other guns in its creation, because it crucially does function as a gun.

What this example illustrates is that, regardless of the ways in which the context may freely manipulate the centrality of different nominal conceptual features, an adjective like *fake* is grammatically restrained in which features it operates over. One can manipulate context to decentralize the telic quale as much as one wants, but it will never result in *fake* targeted anything else - the result will simply be either vacuous or intersective modification.

What about Experiment 1? We also want to capture the variation observed in the absence of linguistic triggers. Recall that participants expanded or contracted their meaning of a noun in virtue of whether real or fake instances of the noun appeared in the context, but were

unaffected by the presence of other nouns entirely. This pattern is, in fact, cleanly derived if we assume that the Non-Vacuity Principle operates as a trigger of mod_l in such cases. Modulation in the form of promotion and demotion of C-structure is relatively powerful but crucially does *not* allow for the manipulation of the content of qualia. The context that contains both real and fake dogs, for example, is easy to handle - speakers are free to choose between a more restrained, literal interpretation of *dog* or a wider one which allows the pretend dogs. Neither of these interpretations will incur an NVP violation, as in both cases there will be a positive extension - at least the real dogs - and a negative extension - any non-dogs. What about when there are only pretend dogs, and no real ones? A strictly literal interpretation now runs into a problem - there is no positive extension of *dog* in such a context under such an interpretation. But mod_l to the rescue, producing a modulation function consisting of the dimension operator F ; the new E-structure for *dog* then only makes reference to the perceptual, formal characteristics of dogs, and the stuffed dogs are happy category members.²⁶

But what about the cases where there are no dogs, real or pretend, in the context, and we ask to point to the dogs? We don't see participants selecting other objects at random, when in fact, this is what we would predict solely from NVP - *dog* has no positive extension in the context! It should be coerced into expanding. Given the absence of any other lexical item in the phrase, no recourse to HPP can save us, since nothing says we can't adjust the meaning of a lone noun. And perhaps NVP does raise the alarm here, but luckily (given that we do not see expansion), mod_l is powerless to assist. There is no combination of dimension and core enrichment operators which can manipulate conceptual structure in order to make a cat count as a dog, since the internal content of qualia cannot be adjusted.

²⁶Note that this process does not lead to the same look-ahead problem we identified for the earlier version of Partee's story. In particular, we do not need to look ahead to the output of a compositional operation to determine that there will be an NVP violation - we can do so simply by evaluating the terminal node *dog* against the context.

5.4.4 The role of extensional structure

At this point, a natural question to ask is to what extent we want to hold on to the original distinction between E- and C-structure, if we are so regularly using dimensions of C-structure to replace the E-structure. Wouldn't it be easier to assume that lexical entries are fundamentally composed of C-structure, and mod_i simply promotes or demotes components of that to be extensional, on their own? Such an approach would absolutely be compatible with the phenomena sketched here. The relevant considerations for making this decision, as I see it, would be the existence or nonexistence of cases where an entity violates all elements of a noun's C-structure but where we would still want to classify it as that noun anyway, perhaps because it still had some core essence that was never manifested in any way. Do such cases exist? Well, that is partially a question about which qualia exist (nothing in this system prevents there from being more, and shortly we'll see an example that suggested one), and about how much information is actually encoded into each quale. What constitutes the telic function of an entity? Which information is allowed to be represented in such a slot? It seems like, with a generous enough answer to these questions, E-structure could easily be eliminated as a distinct level. A more narrow answer might require it to be retained. It appears to me that the revision which collapses E-structure into a privileged status for certain qualia is better suited to capture the variation we're interested in here, and might be modeled like so:

- (251) *An attempt at reducing extensional to conceptual structure:*
- $$\begin{aligned} \llbracket \mathbf{gun} \rrbracket &= \lambda x. \forall C_{\text{GUN}} [C_{\text{GUN}} \in E_{\text{GUN}} \rightarrow C_{\text{GUN}}(x)] \\ C_{\text{GUN}} &= \{ \\ &\quad \text{CONSTITUTIVE: } \lambda x. \text{PARTS-GUN}(x) \\ &\quad \text{FORMAL: } \lambda x. \text{PERCEPTUAL-GUN}(x) \\ &\quad \text{TELIC: } \lambda x. \text{GEN } e [\text{SHOOTING}(e) \wedge \text{INSTRUMENT}(e, x)] \\ &\quad \text{AGENTIVE: } \lambda x. \exists e_1 [\text{MAKING}(e_1) \wedge \text{GOAL}(e_1, \text{GEN } e [\text{SHOOTING}(e) \wedge \\ &\quad \quad \text{INSTRUMENT}(e, x)])] \} \\ E_{\text{GUN}} &= \{ \text{TELIC} \} \end{aligned}$$

Here, our equivalent of E-structure is replaced with a function E over qualia, which

returns true for qualia which are extensional and false otherwise. The actual function $\llbracket \text{gun} \rrbracket$ returns true of an entity to which all qualia in E correctly apply. I assume that the default lexical entry for *gun* has only the telos as extension-determining, but an adapted version of mod_i could easily manipulate that. For example, we see that in Experiment 2, *fake guns* are actually judged to be *guns* in the majority of cases. To achieve such an interpretation, mod_i could apply and give a lexical modulation function l that modulated (251) by uploading FORMAL into E . Then, modification by *fake* would behave similarly, but no privation would arise.

This is not at all to argue that (251) is the *correct* way to reduce E-structure to C-structure, but merely that such an option would be available to derive the same privative variation we wish to capture. I think what we see here is that, in order to capture the rough intuition that extensional structure shouldn't be a distinct independently structured layer so much as a qualification of particular parts of conceptual structure, the basic architecture of DCS might need to be significantly revised. I'll return to that in the next section shortly, but first let's look at another motivation for adjusting our notion of E-structure.

Specifically, the adjustment to privative adjective denotations captures the first half of our variation question, that is, how can we allow privatives to behave subsectively? But the results of Experiment 2 also suggest a second question: how can we allow intersectives to behave privatively? In particular, how can we do so while maintaining a uniform lexical entry across uses of an intersective predicate? It would be simple enough to give the exact same treatment to material adjectives like *plastic* as we did to *fake*, but we would then run into trouble for simple nominal uses. We would like to preserve the idea that such adjectives are one-place predicates like nouns, rather than functions over DC lexical entries. I think we can do so, if we take a look at how Predicate Modification should function in this system. To illustrate, let's consider the classic example of a *stone lion*. We would like to derive that this is not a lion. Let's assume that the lexical entry for *stone* is something quite simple, like the following, contributing only a material constitutive property:

(252) $\llbracket \mathbf{stone} \rrbracket =$
 E-structure: $\lambda x. \text{STONE}(x)$
 C-structure:
 CONSTITUTIVE: $\lambda x. \text{STONE}(x)$

This denotation for *stone* is type $\langle e, t \rangle$ and so will compose with its same-typed head noun *lion* through Predicate Modification. *Lion* is perhaps specified for constitutive, formal, and agentive features, but might not have a telos - the choice won't really affect the argument here.

(253) $\llbracket \mathbf{lion} \rrbracket =$
 E-structure: $\lambda x. \text{LION}(x)$
 C-structure:
 CONSTITUTIVE: $\lambda x. \text{SUBSTANCE-LION}(x)$
 FORMAL: $\lambda x. \text{PERCEPTUAL-LION}(x)$
 AGENTIVE: $\lambda x. \exists e_1 [\text{BIOLOGICAL-BIRTH-LION}(e_1, x)]$

Straightforward Predicate Modification with these denotations would lead to the following denotation for *stone lion*:

(254) $\llbracket \mathbf{stone lion} \rrbracket =$
 E-structure: $\lambda x. \text{LION}(x) \wedge \text{STONE}(x)$
 C-structure:
 CONSTITUTIVE: $\lambda x. \text{SUBSTANCE-LION}(x) \wedge \text{STONE}(x)$
 FORMAL: $\lambda x. \text{PERCEPTUAL-LION}(x)$
 AGENTIVE: $\lambda x. \exists e_1 [\text{BIOLOGICAL-BIRTH-LION}(e_1, x)]$

But this is problematic. A *stone lion* should not be a lion, nor should it be the result of a lion birth, or composed of biological lion parts, and the latter is explicitly contradictory with the constitutive quale of *stone* such that the resulting set is empty, a violation of Non-Vacuity. In fact, all that we want from *lion* is its formal quale, namely having the perceptual features or shape of a lion. An important observation is that even if we adopted a much more minimal semantics for *lion*, where it's only specified for an agentive feature at all, we would still have a Non-Vacuity problem: there are no things that were born from a lion and made of stone. What this tells us is that the problem is not merely when there is competition for

the value of a single quale.

A rather natural way to model the kind of process that occurs in understanding *stone lion* is as a kind of competitive reasoning process where the language user considers different possible featural combinations and selects the one that preserves the most information from both halves of the composition without violating Non-Vacuity; this is similar to an Optimality Theoretic account of ‘Interpretation by Optimization’ proposed by Oliver (2014). On that account, certain features are ranked more highly than others, which won’t always cause one to override another, but can when they are incompatible. Here, the fact that *stone* contributes exclusively a constitutive quale is the reason why the resulting *stone lion* is something made out of stone - any resolution to the feature selection process that left out stone’s single contribution would automatically make the composition vacuous. Anything that competes with that - as determined by world knowledge - then is automatically disqualified. This process isn’t fully deterministic, however, and has to be sensitive to the context of evaluation. For example, Del Pinal observes that the word *literally* can instruct the listener to relax their commitment to Non-Vacuity and accept seemingly empty denotations in the case of constitutive material adjectives, but not in the case of true privatives:

- (255) a. Something unbelievable happened in a laboratory at Harvard. Scientists discovered a way of making, literally, stone lions.
b. #Something amazing happened in a laboratory at Harvard. Some engineer managed to make, literally, a fake gun.

The observation is that we are willing to imagine a hypothetical living lion composed of stone in the case of (255a), but that we cannot do any parallel operation for (255b). In the current analysis, we could explain this by saying that *literally* is a signal for the listener to ignore world knowledge constraints in their interpretation. This would affect the interpretation of constitutive material adjectives, because the reason why being a real stone and a real lion are incompatible has to do with potentially malleable facts about the world, that a sufficiently science fiction scenario could avoid. But it wouldn’t affect privatives, because the incompatibility is due to explicit negation which is the only contribution of the

adjective.

The fact that there is a conflict at all between *stone* and *lion* only occurs, of course, on the most literal interpretation of the noun. We productively and regularly use words like *lion* to refer to their stone versions in the right context, freely expanding to more liberal denotations including objects that include only the formal, representational qualities of the original noun. That process isn't incompatible with what's being described in this intersection here - in fact, it is likely one and the same process, which can either be triggered by the introduction of a literally-incompatible modifier like *stone* or a literally-incompatible discourse context, like pointing to a statue of a lion. Clearly we do have context-triggered expansion processes in language, and it would be natural from this perspective to model those processes as exactly the same kind of qualia loosening or 'upload' of particular qualia into extensional status that takes place in both cases with and without overt modification.

There is a lingering problem, however - even if natural interpretive processes can resolve the conflicts between conceptual structures, the denotation (254) still includes the contradiction in E-structure: we are claiming that the resulting item is both a lion and stone. This kind of consequence of intersection is going to be unavoidable unless we have a way to relate E-structure to C-structure systematically, and so this provides another motivation for the reanalysis of the role of E-structure suggested above. What we need is an indication that holding onto the formal quale of *lion* isn't enough for the resulting entity to be still considered in the extension of *lion*.

At this point, I could go into the weeds and propose a more detailed revision to Dual Content Semantics that implemented a new version of extensional structure. However, I want to make a larger move instead that would invalidate that work entirely. I think that the structured conceptual information of DCS is critical to explain the modification patterns that we have been concerned with, and so any theory I pursue here will want to hold on to that insight, and the choice to make them compositionally active. However, I think that we can effectively restrict the power of Del Pinal's system in a highly desirable way while making

it consonant with the syntactic ideas that the rest of the dissertation has been developing by abandoning the simultaneous dual representation of concept and extension altogether. In the next and final section, I want to start sketching what that combination might look like.

5.5 Theorizing nominal concepts

The aim of this last section will be to explore some possible motivations and options for unifying a conceptually-enriched system like Del Pinal's Dual Content Semantics, which has its advantages for representing the low-level details of modification, with the syntactic-semantic account developed in the rest of the dissertation, which has its advantages for capturing the structure-sensitivity of modification types. It will be a rough sketch at most, and raise significantly more questions and gesture at significantly more areas of research than it will contribute to itself; my goal here is to suggest that bringing these various lines of thinking together might be an enormously valuable endeavor, and that research in these various branches should be taking the other branches quite seriously as sources of evidence to bear on questions each branch is already interested in. Covering the empirical and theoretical ground to unify all these thoughts would be the project of several additional books. So for now, a sketch.

5.5.1 Derivational, not simultaneous, duality

Here is the basic idea: we have seen some evidence that certain instances of composition, namely nonintersective modification, require access to detailed conceptual information about the noun being modified. That motivates an enriched representation of nominal semantics including various conceptual features, i.e., qualia. However, we have also seen evidence that this type of modification is available in only very restricted configurations in the course of the syntactic derivation. Specifically, it's only allowed when the modifier composes with the noun prior to the noun being enumerated - composing with a grammatical operation (realized as morphological number or a classifier) that converts its kind-level meaning into a set of

objects - and after that enumeration, the strategy of modification that involves conceptual features of the noun is blocked.

Against this backdrop, it does appear that nouns have a dual semantic content. Compositional operations make reference both to their conceptual content, e.g., in nonintersective modification, but also to their extensional content, e.g., in intersective modification. I have argued that, despite surface appearances, nonintersective and intersective modification are never taking place in the same syntactic position - it just so happens that sometimes, the intersective mode of composition can create a sort of bootleg version of the nonintersective semantics and lead to the same surface reading. So it follows that this dual content is not, as Del Pinal assumes, simultaneously active and available for concurrent compositional access, but rather split across the syntactic lifespan of a noun. If it were present consistently across the noun's syntactic realizations, we would expect nonintersective modification to be available in all configurations. But, as we have established, it is not. It is restricted to those areas where existing semantic theories would say the noun denotes a kind.

If we take seriously both the idea that nonintersective modification is the modification of kinds and also involves the manipulation of structured conceptual features, the intuitive next move is to identify kinds with those conceptual structures, and that is the move I want to suggest. Del Pinal's C-structure is the grammatical representation of a kind, which exists independently of its extensional level, and it can be converted into extensional structure via the operations of grammatical number. When we take this idea seriously and look at the motivation for both of these ideas - conceptual structure and grammatical kinds - what we find is actually significant overlap in theorizing. C-structure is a way to sort and structure some subset of our conceptual beliefs about a kind of thing, serving as a sort of fallible instruction manual for identifying members of that kind. Del Pinal describes the role of C-structure often in these kinds of language: 'The C-structure of *e* can be thought of as a restricted set of general beliefs associated with *e*'s extension' (Del Pinal, 2015, 12), or 'C-structure of lexical items encodes what, according to our best psychological theories,

are the basic components of the corresponding kinds of concepts' (Del Pinal, 2018, 175). These structures aren't definitions - something may violate them and yet still fall under the extension; it 'traces, but does not fully determine' extension (Del Pinal, 2015, 6). This is quite similar to the way that kinds are often described: their status as 'regularities' (Chierchia 1998) that can enable us to pick out individuals but without strict applicable rules that can accurately determine membership for every possible individual; recall the discussion of *Dogs have four legs* from the beginning of Chapter 4.

A useful perspective on kinds here comes from Mueller-Reichau (2013), who argues that we should view kinds as *sortal concepts* - 'cognitive categories for identifying and classifying objects'. While our object-based system of perception and cognition centers around individuation of objects, we concurrently develop a system for kind-based cognition that centers around categorizing (Xu, 2005, 24), leading to a dual system of reference (Dölling, 1993; Reyes, 1994). The process of developing kinds is illustrative for thinking about them relative to conceptual structure: we observe how an object is used, allowing us to establish a mental concept for that type of object, and over repeated exposure this crystallizes into a new sort of thing, a kind representation for objects that are used that way or share some other core characteristic or property, and this process itself repeats: 'by successively establishing more and more kind representations, a new domain besides the real object domain evolves within our commonsense ontology, viz. the abstract and in a sense 'unreal' kind domain' (Mueller-Reichau 2013:35). In this way kinds are 'reifications' of our conceptual descriptions (Lobner 2002:20), some existent but 'unreal' entity that instantiates the description of a concept, comprised of the individuals that fall under that concept without extensionally picking them out.

Kinds on this view are a way of linking grammatical representations to conceptual representations, or as Mueller-Reichau describes them (2013:53), 'bridgeheads between the conceptual system and the grammatical system'... 'on the one hand, kinds are sortal concepts belonging to the general conceptual system. On the other hand, kinds are ontological prim-

itives, i.e. possible references of linguistic expressions’. Those concepts exist in a language-independent sense (though the role of language in building them up is open for debate), but they also ‘linguistically manifest themselves as kind symbols in the mental lexicon of speakers’ (Mueller-Reichau 2013:55). But it is not the case, I think, that kinds *are* these concepts, diverging from Mueller-Reichau’s perspective. Concepts are richly complex and involve potentially infinite features that our world knowledge associates with a particular type of thing; I know about dogs some general and some specific facts, and my concept of a dog is highly informed by my own personal experience with my dogs, but the kind **DOG** that has the status of a grammatical object in my internal semantic system might lack that experiential information in favor of just the general principles of dogness. Kinds, if we are going to identify them with C-structures, are just that: structured. And they need to be, if they are going to participate in compositional operations. There’s plenty of debate over whether nonlinguistic concept representations like prototypes can compose with each other (Gleitman et al., 2012; Prinz, 2012), but they almost certainly can’t be the arguments of Function Application in the formal semantic sense; I’m claiming here that aspects of kinds can be.

One objection that might come up early on from the traditional way of thinking about kinds is that while ‘properties are unsaturated, something that (at a world) is true or false of individuals. Kinds are saturated, something that at a world has concrete, if possibly spatiotemporal discontinuous, manifestation’ (Chierchia 1998:354). The qualia in C-structure are unsaturated: they are properties, after all. But this treatment of kinds as unsaturated (at a given world) seems correct; we want to retain an $\langle s, e \rangle$ type for them. This requires some precision in what exactly I am proposing when identify kinds with C-structures. The semantic representation for a kind that we have been working with so far is like this:

$$(256) \quad \mathbf{DOG} = \lambda w[\iota x[*\mathbf{dog}_w(x)]]$$

And I intend to retain this idea: kinds are grammatically functions from worlds into an individual, a plural individual comprised of instantiations of it in a world. What I want to

identify C-structures with is the internal function that builds the plurality: the **dog** under \ast , not the **DOG** on the left side of the equality. C-structure can describe the sortal concept of a kind, the criteria by which the kind is built up in a particular world. And given our sense of modulation, those things might change in a given context, just like we may loosen or tighten our understanding of what a dog or any other object needs to be, without changing the overall structure of the kind function. It is this thing that is modified in our nonintersective modification:

$$(257) \quad \lambda w_s[\ast[\lambda x_e[\exists z_{\langle s,e \rangle}[z \leq \mathbf{good}(\mathbf{THIEF}) \wedge x \sqsubseteq z_w]]]]$$

The application $\mathbf{good}(\mathbf{THIEF})$ on this view is the application of the adjective's semantic content to the qualia that make up the conceptual structure of the kind. And that will be different depending on the adjective - **good**, like **fake**, probably modifies the telic function of the noun it modifies, declaring that an individual in the new $\mathbf{good}(\mathbf{THIEF})$ kind does not just perform the function of a thief but does so well. This contrasts with **fake**, which would target the same quale but delete it rather than adjusting something about it or adding additional descriptive information to specify it further. In some sense, then, privative adjectives are the extreme version of subsectives: while **good** here prunes out some thieves by specifying a stricter requirement for the quale, **fake** prunes out them all.

The kind itself, then, is something saturated at a particular world - it is an individual comprised of all its members. But the sortal concept that builds up that kind, the cognitive guide for categorization needed to collect up those members, that thing is unsaturated, a set of functions for checking against individuals to determine membership. And that thing is compositionally active exclusively in the narrow domains where nonintersective kind modification is allowed, sealed off afterwards. It is at this sealing process - the reduction of the kind's complex conceptual structure to the enumerated noun's simpler extensional structure - that pragmatic modulation intervenes as Del Pinal intended to select which of those conceptual features are going to become extension-determining in this particular context.

5.5.2 Roots identify concepts; kinds structure them

I suggested above that kinds are not just concepts, but some structured subset of the information that concepts contain. How does that structuring and filtering process happen? I think that we can point to a significant body of work occurring in contemporary syntactic theory as a way of talking about this: the study of roots and their categorizing process. A particular thrust in research surrounding syntactic theories of words like Distributed Morphology has been stripping semantic content from roots (Borer, 2005; Harley, 2014) with the eventual conclusion that roots - which Leffel argued denoted subkind predicates - are merely indices, contentless themselves absent additional structure. I think that this area of thought can offer some needed specification to the syntactic component of how a proposal like the one above might function.

One branch of thought in particular from Acquaviva (2009, 2014a,b) and Panagiotidis (2011, 2014, 2015) has argued for a vision of roots on which they function as simply pointers to concepts. Roots don't represent conceptual information themselves, but do identify which concept contains the information that can be represented with the inclusion of additional grammatical structure. Roots alone aren't interpretable at the Conceptual-Intensional interface on this perspective, but rather require categorization (via a category-determining head like *n* or *v*) for semantic legibility. The perspective that roots require categorization for syntactic reasons is not an unusual one, following the 'no free roots' rule of Baker (2003) and the Categorization Assumption of Embick and Marantz (2008); Panagiotidis (2011, 2014) follows in the footsteps of Arad (2005) in arguing that this is due to roots being semantically impoverished:

(258) Free (acategorial) roots are not legitimate LF objects, because they are unreadable by the Conceptual-Intentional/SEM systems. (Panagiotidis 2014:290)

The mechanism that syntax uses to convert roots into legitimate LF objects is categorizers: categorizing heads which apply a syntactic category label to roots also impose onto that root an 'interpretive perspective'. Those category heads at least minimally come in nominal

and verbal flavors, which Panagiotidis argues impose distinct modes of interpretation onto the conceptual information identified by the root (identified by, not contained by - the root can tell us what concept to look at, but doesn't itself denote that concept):

- (259) LF-interpretation of categorial features:
A [V] feature imposes an extending-into-time perspective at LF; an [N] feature imposes a sortal perspective at LF. (Panagiotidis 2011:371)

Focusing on the [N] feature for nouns, this sortal perspective is that of a kind: 'the representation is understood to apply to things of a certain kind, but not others. Thus, the sortal DOG allows us to think about dogs, but not tables, trees, wood or any other kind of thing.' (Prasada, 2008, 6) This contrasts with the temporal perspective imposed onto verbs - kinds, as we have discussed, may be both spatially and temporally discontinuous, while verbs are interpreted as events that have particular temporal instantiations.

The unification of this perspective and Del Pinal's would naturally be that the compositional reflex of this sortal perspective that the nominalizing categorizer imposes is C-structure: the root indexes a concept, and the [N] feature imposed by the nominalizing projection is interpreted as pulling certain pieces of information from that concept and organizing them into the grammatical representation of a kind, the sortal concept that determines kind membership, which is composed of some bits and pieces of useful categorizing information. The notion of indices to complex conceptual information isn't only discussed in the context of roots, but sometimes words more broadly - e.g. Pietroski (2018), who views semantic meanings as 'fetch instructions' to retrieve concepts from a particular place in memory - or even concepts themselves, e.g., Quilty-Dunn (2021), who argues that concepts themselves are pointers to areas of memory, based on considerations from exactly the same kind of polysemy problems that motivated this dissertation. The point is that that concept in nonlinguistic cognition has quite a lot of information associated with it - a conceptual package in a long-term memory network (Borer, 2005) or a 'grab-bag' of associated information (Rayo, 2013) - and the root can tell us where to look for that, but grammar can't

manipulate that information until it is organized appropriately by imposing the interpretive structure required of a lexical item (a noun or a verb, and potentially an adjective, though see Mitrović and Panagiotidis 2020 on how adjectives might be some composite of nominal and verbal perspectival features). A nominalizer imposing the structure of a sortal concept makes it CI-interface-legible, but that doesn't necessarily finish the lifespan of a nominal concept - in most cases, it will end up as a property of individuals, but that requires the additional semantic contribution of number.

For our purposes, the important component of this line of thinking is that this process creates multiple distinct domains for modification to occur. Intersective modification can only occur once a concept has been retrieved, rendered legible as a sortal kind, and then enumerated and individuated as a property. Nonintersective modification can only occur after a concept has been converted into a kind, but prior to its conversion to an object-level property. Potentially there is a third domain, too - modification that occurs between the root and its categorizing head, before the conceptual information has been rendered into a structure appropriate for any compositional operations. The most likely candidate for this kind of noncompositional, pre-category modification is idioms (Borer, 2014).

The end goal, I think, of considering these factors is that a complete understanding of the nature of modificational ambiguity is likely to rely on a complete understanding of the target of modification. If we want to move away from theories that rely on lexically listed polysemy for adjectives, our best bet is going to be to develop a sophisticated theory of what kind of object a noun can be, and when it means what in the derivation, paying close attention to both the syntax and semantics and being as explicit as possible about our mapping hypothesis between the two at all points in the derivation. If we do that, then the study of modification and the study of fundamental ontological questions of language (are meanings extensional? what is a kind? are there linguistic concepts? how do words come to be?) can be made more effectively symbiotic - knowledge about the fundamental structural and interpretive properties of different domains can give us valuable theoretical tools to use to explain

patterns in modification, and those same patterns of modification can be used as illustrative diagnostics for what kinds of meanings must be available at different locations in the structure. The result may be a simplification of our concept of meaning (i.e., the elimination of polysemy) or the opposite (i.e., the expansion of compositional noun meanings from solely extension to include rich conceptual structures), but the important part is that either of those moves is both empirically grounded and implemented with interface-accountable moves, the latter of which is probably not possible for either domain of investigation (modification or grammatical ontology) without reference to the other.

6 Conclusion

This dissertation set out to motivate and develop a theory of the (non)intersective ambiguity found in certain natural language adjectives. Rejecting theories that relied on lexical polysemy to derive the ambiguity by positing two distinct denotations, it has instead pursued a compositional analysis where a single denotation for the adjective results in two (or more) possible surface readings depending on its syntactic relationship to the noun it modifies (along with certain pragmatic factors). Synthesizing the arguments of the preceding chapters, the analysis can now be summarized in full.

Adjectives (of the relevant class) are uniformly functions from subkind predicates into subkind predicates. They take as their initial argument a predicate that maps kinds into truth values, evaluating whether the given kind is a subkind of a particular kind, and return a new predicate that adjusts some feature of the superkind of evaluation. The denotation for the adjective *good*, for example, is given below.

$$(260) \quad \llbracket \mathbf{good} \rrbracket = \lambda P_{\langle \langle s, e \rangle, t \rangle} \lambda k_{\langle s, e \rangle} . k \leq \mathbf{good}(\iota P)$$

This semantics for *good* enables it to take in a subkind predicate P , pull out the maximal kind of that predicate ιP , create a new kind by applying the feature **good** to ιP , and return a new predicate that checks if subkinds are in that new kind. This is all that *good* ever does - nonintersective and intersective readings alike are derived from this denotation interacting with different components of nominal semantics.

The basic type of common nouns is that of a subkind predicate. Once a root has been categorized as a noun in the syntax, it denotes this $\langle \langle s, e \rangle, t \rangle$ -type term, like the following denotation for *thief*:

$$(261) \quad \llbracket \mathbf{thief} \rrbracket = \lambda k_{\langle s, e \rangle} [k \leq \lambda w_s [\iota [* \mathbf{thief}_w]]] \quad (\text{unmarked for number})$$

One way that nonintersective interpretations of phrases like *good thief* - the interpretation on which it is the individual's thieving that is good - is through the direct modification of (261) by (260). The adjective Merges with the noun inside the noun phrase and composes

with it via Function Application, resulting in:

$$(262) \quad \llbracket \mathbf{good\ thief} \rrbracket_{\text{NI}} = \lambda k_{\langle s, e \rangle} [k \leq \lambda w_s [\iota [* \mathbf{good}(\mathbf{thief})_w]]] \quad (\text{unmarked for number})$$

But this kind of composition via Function Application is only available up til a certain point in the noun's development in the syntactic derivation - specifically, between its categorization as a noun and its enumeration via composition with a grammatical element like number marking or a classifier. The function of grammatical number is to induce countability by converting the noun's denotation from the kind-level to the object-level, resulting in an intensional property, or a set of the individuals that instantiate that kind in a given world. The result of composing (261) with a grammatical number element, like the singular morpheme SG, is:

$$(263) \quad \llbracket \mathbf{thief+SG} \rrbracket = \lambda w_s [\circ [\lambda x_e [\exists z_{\langle s, e \rangle} [z \leq \mathbf{THIEF} \wedge x \sqsubseteq z_w]]]] \quad (\text{singular})$$

Intersective readings are the result of composing the adjective with this kind of nominal denotation. At this point, however, attempting to compose (260) with (263) would result in a type clash; neither Function Application nor Predicate Modification are appropriate for this pair of types. In order for composition to take place between an enumerated noun and an adjective that adjoins to NP, the adjective needs to be converted into an object-level property as well by combining with a different subkind predicate via a linking operator. The different subkind predicate is a phonologically null but syntactically realized nominal element that is provided by the context and saturates the kind argument of the adjective, specifying, in this case, what kind of 'goodness' is being predicated of the eventual individual argument provided by the linking operator. In the default intersective case, this null element is some semantically minimal kind like *person* or even just the kind of individuals corresponding to D_e , which is the source of the standard intersective reading 'a thief and a good person', but it may also be other more specific contextually-salient kinds. The linking operator needed for composition, which mediates between the adjective and the null argument within the adjective phrase, is defined:

$$(264) \quad \llbracket \rho \rrbracket = \lambda P_{\langle\langle s,e \rangle, t \rangle} \lambda \Psi_{\langle\langle s,e \rangle, t \rangle, \langle\langle s,e \rangle, t \rangle} \lambda w_s \lambda x_e . x \sqsubseteq_w \Psi(P)$$

And the full result of this compositional option, the default intersective reading, is:

$$(265) \quad \llbracket \mathbf{good\ thief+SG} \rrbracket_I = \lambda w_s [\circ [\lambda x_e . x \sqsubseteq_w \mathbf{good}(\mathbf{PERSON}) \\ \wedge \exists z_{\langle s,e \rangle} [z \leq \mathbf{THIEF} \wedge x \sqsubseteq z_w]]]$$

Early in the dissertation it was argued that it is critical to distinguish between nonintersective and intersective as interpretive descriptions of the resulting surface-level readings, and application and intersection as semantic descriptions of the mechanisms of composition. This is because there is not a one-to-one correspondence between application and nonintersectivity on the one hand, and intersection and intersectivity on the other. In fact, there are multiple distinct compositional routes to nonintersectivity. One was described above, but nonintersective readings may also arise via compositional intersection in specific cases. For example, when composing *good thief* via intersection, a contextually salient kind needs to be selected to saturate the kind argument of the adjective. While this is often something like *person*, which is nonspecific enough to be made sufficiently salient in most contexts, another particularly salient kind in this instance is *thief* itself. No principle of composition prevents this redundant choice, and the end result is a denotation with identical truth conditions to the result of the application path: the nonintersective reading.

$$(266) \quad \llbracket \mathbf{good\ thief+SG} \rrbracket_{\text{NI-via-I}} = \lambda w_s [\circ [\lambda x_e . x \sqsubseteq \mathbf{good}(\mathbf{THIEF})_w \\ \wedge \exists z_{\langle s,e \rangle} [z \leq \mathbf{THIEF} \wedge x \sqsubseteq z_w]]]$$

Because the **good thief** kind is a subkind of the **thief** kind, the truth conditions of (266) are the same as those of the ‘true’ nonintersective reading generated via application, (262). This is not true of all adjective-noun combinations, however: certain instances of privative adjectives, like *fake*, are nonsubjective, and therefore the novel kind they create via application is not a subset of the original kind denoted by the noun. Attempting to engineer a nonintersective reading via composition with intersection in these circumstances will lead to a vacuous output and therefore crash.

$$(267) \quad \llbracket \mathbf{fake\ thief+SG} \rrbracket_{\text{NI-via-I}} = \lambda w_s [\circ [\lambda x_e . x \sqsubseteq \mathbf{fake}(\mathbf{THIEF})_w$$

$$\begin{aligned} & \wedge \exists z_{\langle s,e \rangle} [z \leq \mathbf{THIEF} \wedge x \sqsubseteq z_w]] \\ = & \emptyset \end{aligned}$$

It is this unique behavior of privative adjectives that motivates the syntax-semantics mapping hypothesis of this analysis. In the prior structural account, each surface reading is associated with its own underlying syntactic configuration, such that an adjective that is interpretively ambiguous in a specific position must therefore be syntactically ambiguous as well. In contrast, in the proposed account, while interpretively ambiguous surface positions may sometimes correspond to syntactic ambiguity, they do not do so by default; instead, one surface position may correspond to one underlying syntax and further to one compositional semantics, but one compositional process (i.e., intersection) may result in multiple surface readings via this artificial nonintersectivity process. The fact that certain positions across languages, such as the postnominal position in Italian and the preclassifier position in Bangla, display an interpretive ambiguity for subsective but not privative adjectives provides evidence for this approach over the syntactic ambiguity approach.

This account reduces the (non)intersective ambiguity to an ambiguity between kind-level and object-level modification, and identifies the relevant syntactic domains of modification as the pre-number and post-number domains. Empirical motivation for this decision is found in Turkish, where bare nouns unmarked for number display both singular and plural interpretations, but in a way that is sensitive to modification facts: adding kind-level modifiers preserves both interpretations, but adding object-level modifiers eliminates the plural interpretation (or, in some positions, makes the unmarked noun entirely ungrammatical). The nonintersective interpretation of ambiguous adjectives in Turkish pattern exactly like unambiguous kind-level adjectives, and the intersective interpretation like unambiguous object-level adjectives. The same conclusion is supported by Bangla, where the independent structural position of the classifier clearly demarcates the domains of kind and object-level meanings for the noun, and both nonintersective and kind-level modification can only occur below the classifier.

The account also revolves around the existence of phonologically unrealized semantic content within the adjective phrase; namely, both the contextually-provided kind argument and the operator linking it to the adjective. From one perspective, this merely technical solution preserves compositionality and avoids polysemy only from a superficial point of view, by extracting the complexity from the lexicon and artificially inserting it into syntactic structure. These kinds of technical solutions to compositional problems that proliferate the possible maneuvers in syntax can be argued to guarantee the principle of compositionality holds in a trivial sense, as sufficiently unconstrained syntactic operations can always ensure that meanings are compositional. As a result, the account risks not meeting its original criterion for successful theory-building.

However, in contrast to many previous theories of this kind, the covert maneuvers utilized here *are* interface-accountable. The existence of null structure within the adjective in intersective readings specifically is supported by morphophonological evidence in the form of suppletion patterns. Across a moderate cross-linguistic sample of sixteen languages, a generalization arises where intersective interpretations of adjectives appears to block them from having suppletive phonological forms in the comparative or superlative. Given that this kind of comparative suppletion has been extensively argued to depend on locality between the adjective's root and the comparative morpheme, this supports the conclusion that intersective readings involve some additional syntactic structure intervening between the root and its suppletion trigger. This phonological reflex of the proposed additional structure in intersective readings is one way in which this theory is made accountable to the interfaces.

This account also locates the primary burden of deriving nonintersective meanings - not where they appear, not distinguishing them from intersective meanings, but the actual resulting content of nonintersectivity, that is, which aspects of the kind are modified, e.g., in which ways different nouns are modified to be *good* - in a single step of kind modification. Some perspectives on semantics may be comfortable leaving this as is, and arguing that exactly how the concept of *good* gets applied to the concept of *thief* (or, its corresponding kind) in

nonintersective readings is the job for a non-linguistic conceptual system, potentially interfacing with pragmatics. This is particularly complicated for privative adjectives, which seem to require access to the noun's conceptual semantics in order to create their novel meaning, since extensions seem completely uninvolved (though on the present account, this is precisely the same process as any nonintersective reading). In contrast to what might be the standard view on the semantics/conceptual systems divide, the sensitivity of these readings to syntactic variation and the fact that only certain options for interacting with conceptual features - in the form of a limited lexicon of privative adjectives - appear cross-linguistically suggests that, perhaps, this problem too should be given a compositional treatment. This decision has been motivated with the results of two experiments showing that the interpretations of nouns and (privative) adjectives respectively are subject to context-induced modulation, but in a particularly constrained way, only allowing the manipulation of certain features. As a result, the theory is augmented with a compositional system for interacting with conceptual features of nouns. The denotation of a kind on this account is that of a sortal concept, a series of violable categorizing conditions for determining whether an individual counts as a member of the kind, and nonintersective readings are the result of an adjective modifying certain features of this sortal conceptual structure. As a result, the lexical classes of 'privative' vs. 'subsective' adjectives are made obsolete, replaced instead with particular descriptions of how those adjectives interact with conceptual features, and the categorial inferences that have previously described these classes are instead treated as emergent properties of their interaction with specific nouns in different instances of composition.

As with any work of this scope, significant portions of a truly complete analysis are left unresolved and many new questions have been raised. Some of those questions involve how the grammar treats multiple ways to end up at the same meaning - how do principles of economy and interpretation interact with the multiple compositional routes to nonintersectivity? Other questions involve the particular syntactic implementation of the null argument analysis developed here in light of the suppletion data. Still unresolved are the reasons why

certain locality violations in intersective meanings result in ineffability, rather than other options for phonological realization, and how exactly this null material should be placed in the adjective phrase, a question that likely requires more morphological sophistication than developed here. Furthermore, in the theory of kinds utilized here, a number of assumptions are made in order to most effectively facilitate a parsimonious analysis of adjectival modification specifically, but these decisions are inextricable from a much larger debate surrounding how to derive the distribution and various properties of kind reference in natural language independent from modification. And the relationship between nonintersective, direct modification of kinds, the modification of grammatically-structured conceptual features, and the morphosyntax of root meanings and categorization proposed in the final chapter is still radically underspecified, at most a gesture of goodwill between some currently distinct areas of research, identifying that synthesizing these systems which have been developed in large part independently is likely to be necessary for an adequate theory of modificational ambiguity.

Despite all of these open questions and others, I want to conclude by highlighting what I think is the central contributions that this dissertation has made to thinking about ambiguity. And that is that an empirical focus on edge cases can allow us to build theories that, despite the apparent complexity of the edge cases, are fundamentally simpler. The process of building theories of ambiguous adjectives has for some time proceeded by beginning with the simplest, most canonical cases of ambiguity - our beautiful dancers - and constructing an account of them before optionally moving on to incorporate more and more ‘complicated’ cases via the addition of distinct interpretive principles. Often this incremental process ends before adjectives like *fake* are ever reached, either implicitly because they have not been recognized as members of the ambiguity phenomenon of interest or explicitly because their logical properties caused them to be footnoted. And of course, at least half of this dissertation has been dealing entirely with non-privative adjectives, too. Practically no mention of them is made in Chapters 3 and 4 combined. But neither of those chapters would exist without having been able to notice the core distributional pattern in Chapter 2; the proposal that

there is a strict domain of only modification via intersection that nonetheless can achieve nonintersective readings for most adjectives would never have come into being absent the observation of how privative adjectives broke that pattern. In this instance, what has been seen as their unique and compositionally problematic logical property of nonsubsectivity has, with careful attention to their interaction with syntactic factors, come into the light as actually being a fully compositional and context-sensitive, emergent property that has been critical in revealing the underlying behavior of certain syntactic positions. Subsective adjectives, on this view, are too flexible and resilient - they can make do in a number of semantic environments (that is, locations in the syntax where one particular compositional operation is available) and, through adaptive behavior, still achieve a variety of surface interpretations. Privatives are inflexible, like the sensitive canaries from our earlier metaphor, and so directly reflect the compositional nature of their environment in a way that makes them an ideal tool of the semanticist. However, because subsective adjectives are found so much more frequently across a variety of environments, they have become the standard starting place. These subsective birds survive well into the dangerous parts of the cave, and so of course they will be seen as the default case and draw the attention of us miners. But it is the canaries that really tell us what is happening where our eyes are not good enough to see the underlying features of the cave. And so if we as semanticists are in the business of uncovering covert structures and processes, perhaps we should be more often turning our attention to the canaries in the mine.

References

- Acquaviva, Paolo. 2009. Roots, categories, and nominal concepts. *Lingue e linguaggio* VIII(1). 25–51.
- Acquaviva, Paolo. 2014a. Roots, concepts, and word structure: on the atoms of lexical semantics. In Franz Rainer; Francesco Gardani; Hans Christian Luschützky; Wolfgang U. Dressler (ed.), *Morphology and meaning : Selected papers from the 15th international morphology meeting*. Vienna.
- Acquaviva, Paolo. 2014b. The roots of nominality, the nominality of roots. In Artemis Alexiadou; Hagit Borer; Florian Schäfer (ed.), *The syntax of roots and the roots of syntax*, 33–56. Oxford: Oxford University Press.
- Alexiadou, Artemis, Liliane Haegeman and Melita Stavrou. 2007. *Noun phrase in the generative perspective*. Mouton de Gruyter.
- Alexiadou, Artemis and Florian Schäfer. 2010. On the syntax of episodic vs. dispositional -er nominals. In Artemis Alexiadou and Monika Rathert (eds.), *The syntax of nominalizations across languages and frameworks*, vol. 23, Interface Explorations, 9–39. De Gruyter Mouton.
- Anagnostopoulou, Elena and Yota Samioti. 2013. Allosemy, idioms, and their domains: Evidence from adjectival participles. In Raffaella Folli, Christina Sevdali and Robert Truswell (eds.), *Syntax and its limits*. Oxford University Press.
- Arad, Maya. 2005. *Roots and patterns*, Studies In Natural Language And Linguistic Theory. Dordrecht: Springer.
- Arregi, Karlos and Andrew Nevins. 2014. A monoradical approach to some cases of disuppletion. *Theoretical Linguistics* 40(3). 311–330.
- Asher, Nicholas. 2011. *Lexical meaning in context*. Cambridge University Press.
- Baker, Mark C. 2003. *Lexical categories: verbs, nouns, and adjectives*, Cambridge Studies in Linguistics. Cambridge University Press.
- Bale, A., H. Khanjian and M. Gagnon. 2010. Cross-linguistic representations of numerals and number marking. In *Semantics and linguistic theory (salt)*, vol. 20, 1–15.
- Bhattacharya, Tanmoy. 1999. *The structure of the bangla dp*. Ph.D. thesis, University of London.
- Bilgrami, Akeel. 1998. Why holism is harmless and necessary. *Philosophical Perspectives* 12. 105–126.
- Bliss, H. 2004. The semantics of the bare noun in turkish. *Calgary Working Papers in Linguistics* 25. 1–65.

- Bobaljik, Jonathan David. 2012. *Universals in comparative morphology: Suppletion, superlatives, and the structure of words*. Cambridge: MIT Press.
- Bobaljik, Jonathan David and Heidi Harley. 2017. Suppletion is local: evidence from Hiaki. In Heather Newell, Máire Noonan, Glyne Piggott and Lisa deMena Travis (eds.), *The structure of words at the interfaces*. Oxford.
- Boleda, Gemma, Marco Baroni, Nghia The Pham and Louise McNally. 2013. Intensionality was only alleged: On adjective-noun composition in distributional semantics. In *Proceedings of iwcs 2013*. Potsdam.
- Bolinger, Dwight. 1967. Adjectives in english: attribution and predication. *Lingua* 18. 1–34.
- Borer, Hagit. 2005. *In name only*. Oxford University Press.
- Borer, Hagit. 2014. Derived nominals and the domain of content. *Lingua* 141. 71–96.
- Bouchard, Denis. 2002. *Adjectives, number, and interfaces: why languages vary*, North Holland Linguistic Series: Linguistic Variations. Elsevier.
- Željko Bošković. 2012. Phases in DPs and NPs. In À. Gallego (ed.), *Phases: developing the framework*, 343–384. Berlin: De Gruyter Mouton.
- Cappelle, Bert, Pascal Denis and Mikaela Keller. 2018. Facing the facts of fake: a distributional semantics and corpus annotation approach. *Yearbook of the German Cognitive Linguistics Association* 6(1).
- Carlson, Greg. 1977. *Reference to kinds in english*. Ph.D. thesis, University of Massachusetts, Amherst.
- Carlson, Greg. 2003. Interpretive asymmetries in major phrases. In *Asymmetry in grammar: Syntax and semantics*, 301–314. John Benjamins.
- Chierchia, Gennaro. 1998. Reference to kinds across languages. *Natural Language Semantics* 6(4). 339–405.
- Chomsky, Noam. 2001. Derivation by Phase. In Michael Kenstowicz (ed.), *Ken Hale: A life in language*, 1–52. Cambridge, MA: MIT Press.
- Cinque, Guglielmo. 2010. *The syntax of adjectives: a comparative study*. MIT Press.
- Cinque, Guglielmo. 2014. The semantic classification of adjectives: a view from syntax. *Studies in Chinese Linguistics* 35(1). 3–32.
- Coulson, Seana and Gilles Fauconnier. 1999. Fake Guns and Stone Lions: Conceptual Blending and Privative Adjectives. In & L. Michaelis B. Fox, D. Jurafsky (ed.), *Cognition and Function in Language*. CSLI.
- Danon, Gabi. 2012. Two structures for numeral-noun constructions. *Lingua* 122. 1282–1307.

- Dayal, Veneeta. 1992. The singular-plural distinction in hindi generics. In *Semantics and linguistic theory (salt)*, vol. 2, 39–58.
- Dayal, Veneeta. 2004. Number marking and indefiniteness in kind terms. *Linguistics and Philosophy* 27. 393–450.
- Dayal, Veneeta. 2012. Bangla classifiers: mediating between kinds and objects. *Rivista di Linguistica* 24(2). 195–226.
- Del Pinal, Guillermo. 2015. Dual Content Semantics, privative adjectives, and dynamic compositionality. *Semantics and Pragmatics* 8(7). 1–53.
- Del Pinal, Guillermo. 2018. Meaning, modulation, and context: a multidimensional semantics for truth-conditional pragmatics. *Linguistics and Philosophy* 41. 165–207.
- Déprez, Vivane. 2001. On the syntactic and semantic nature of haitian bare nps. In D. Cresti, C. Tortora and T. Saterfield (eds.), *Current issues in romance linguistics: Selected papers from the 28th linguistics symposium on romance languages*. Amsterdam: John Benjamins.
- Déprez, Viviane. 2005. Morphological number, semantic number, and bare nouns. *Lingua* 115(6). 857–883.
- Despić, Miloje. 2019. (Non-)Intersective adjectives and root suppletion. *The Linguistic Review* 36(3).
- Despić, Miloje and Yael Sharvit. 2011. Some "non-intersective" adjectives are genuinely noun-taking. In Suzi Lima, Kevin Mullin and Brian Smith (eds.), *Proceedings of the Northeast Linguistic Society* 39, 281–294.
- Dölling, J. 1993. Commonsense ontology and semantics of natural language. *Zeitschrift für Sprachtypologie und Universalienforschung* 46. 133–141.
- Egg, Markus. 2008. Reference to embedded eventualities. In Johannes Dölling, Tanya Heyde-Zybatow and Martin Schäfer (eds.), *Event structures in linguistic form and interpretation*, 149–172. Berlin: Walter de Gruyter.
- Embick, David and Alec Marantz. 2008. Architecture and blocking. *Linguistic Inquiry* 39. 1–53.
- Ferris, Connor. 1993. *The meaning of syntax: A study in the adjectives of english*. London: Longman.
- Fodor, Jerry and Ernest Lepore. 1992. *Holism: A shopper's guide*. Blackwell.
- Franks, B. 1995. Sense Generation: A "quasi-classical" approach to concepts and concept combination. *Cognitive Science* 19. 441–505.
- Frege, Gottlob. 1914. Logic in mathematics. In H. Hermes, F. Kambartel and F. Kaulbach (eds.), *Gottlob Frege: Posthumous Writings*, 72–84. Chicago: University of Chicago Press. Trans. P. Long and R. White.

- Gallin, Daniel. 1975. *Intensional and higher-order modal logic: with applications to montague semantics*. Elsevier.
- Gleitman, Lila R., Andrew C. Connolly and Sharon Lee Armstrong. 2012. Can prototype representations support composition and decomposition? In Markus Werning, Wolfram Hinzen and Edouard Machery (eds.), *The oxford handbook of compositionality*, Oxford Handbooks in Linguistics, 418–436. Oxford University Press.
- Görgülü, E. 2012. *Semantics of nouns and the specification of number in turkish*. Ph.D. thesis, Simon Fraser University.
- Guerrini, Janek. 2018. The link between misinterpretation, intentionality, and mental agency in the natural language interpretation of “fake”. *Rivista Italiana di Filosofia Analitica Junior*.
- Halle, Morris and Alec Marantz. 1993. Distributed Morphology and the pieces of inflection. In Ken Hale and Samuel Jay Keyser (eds.), *The view from Building 20: Essays in linguistics in honor of Sylbain Bromberger*, 111–176. Cambridge, MA: MIT Press.
- Harley, Heidi. 2005. How do verbs take their names? Denominal verbs, manner incorporation and the ontology of roots in English. In Nomi Erteschik-Shir and Tova Rapoport (eds.), *The syntax of aspect: deriving thematic and aspectual interpretation*, 42–64. Oxford: Oxford University Press.
- Harley, Heidi. 2014. On the identity of roots. *Theoretical Linguistics* 40(3-4).
- Hausser, R.R. 1984. Surface compositional grammar. *Studies in Theoretical Linguistics* 4.
- Heim, Irene. 2000. Degree operators and scope. In Brendan Jackson and Tanya Matthews (eds.), *Proceedings of Semantics and Linguistic Theory X*, 40–64. Ithaca, NY: Cornell University.
- Hintikka, J. 1983. *The game of language: Game-theoretical semantics and its applications*. Dordrecht: Reidel.
- Janssen, Theo M.V. 1983. *Foundations and Applications of Montague Grammar*. Amsterdam: Mathematisch Centrum.
- Kamp, Hans. 1975. Two theories about adjectives. In E.L. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge University Press.
- Kamp, Hans and Barbara Partee. 1995. Prototype theory and compositionality. *Cognition* 57. 129–191.
- Krifka, Manfred, Jeffrey Pelletier, Greg Carlson, Alice ter Meulen, Godehard Link and Genaro Chierchia. 1995. Genericity: an introduction. In *The generic book*. The University of Chicago Press.
- Larson, Richard K. 1998. Events and modification in nominals. In Devon Strolovitch & Aaron Lawson (ed.), *Proceedings of SALT 8*, 145–167.

- Larson, Richard K. 1999. Semantics of adjective modification. LOT Winter School Lectures, Amsterdam [http://semlab5.sbs.sunysb.edu/rlarson/LOT\(99\)/Contents.html/index.html](http://semlab5.sbs.sunysb.edu/rlarson/LOT(99)/Contents.html/index.html).
- Larson, Richard K. and Franc Marušič. 2004. On indefinite pronoun structures with APs: reply to Kishimoto. *Linguistic Inquiry* 35. 268–287.
- Larson, Richard K. and Naoko Takahashi. 2004. Order and interpretation in prenominal relative clauses. In M. Kelepir and B. Öztürk (eds.), *Proceedings of the Workshop on Altaic Formal Linguistics II*, vol. 54, MIT Working Papers in Linguistics, 101–120. Cambridge, MA: MITWPL.
- Leffel, Timothy James. 2014. *The semantics of modification: adjectives, nouns, and order*. Ph.D. thesis, New York University.
- Link, Godehard. 1983. The logical analysis of plurals and mass terms: a lattice-theoretical approach. In *Formal semantics: the essential readings (2002)*, 127–146. Blackwell.
- Longobardi, Giuseppe. 1994. Reference and proper names: A theory of n-movement in syntax and logical form. *Linguistic Inquiry* 25(4). 609–665.
- Maienborn, Claudia. 2020. Revisiting Olga, the beautiful dancer: An intersective A-analysis. In *Proceedings of SALT 30*, 63–82.
- Marantz, Alec. 2013. Locality domains for contextual allomorphy across the interfaces. In Ora Matushansky and Alec Marantz (eds.), *Distributed Morphology Today: Morphemes for Morris Halle*. MIT Press.
- Martin, Joshua and Daria Bikina. 2022. Intersectivity at the interface: the syntax and semantics of Russian adjectives. In *Proceedings of FASL 30*.
- McNally, Louise and Gemma Boleda. 2004. Relational adjectives as properties of kinds. *Empirical issues in syntax and semantics* 5. 179–196.
- Mel'čuk, Igor A. 1994. Suppletion: Toward a logical analysis of the concept. *Studies in Language* 18(2). 339–419.
- Mitrović, Moreno and Phoevos Panagiotidis. 2020. Adjectives exist, adjectivisers do not: a bicategorical typology. *Glossa: a journal of general linguistics* .
- Moltmann, Friederike. 2007. Events, tropes and truthmaking. *Philosophical Studies* 134(3). 363–403.
- Moltmann, Friederike. 2009. Degree structure as trope structure: A trope-based analysis of positive and comparative adjectives. *Linguistics and Philosophy* 32. 51–94.
- Moltmann, Friederike. 2013. On the distinction between abstract states, concrete states, and tropes. In Alda Mari, Claire Beyssade and Fabio Del Prete (eds.), *Genericity*, 292–311. Oxford: Oxford University Press.

- Montague, Richard. 1970. English as a formal language. In Bruno Visentini et al (ed.), *Linguaggi nella Società e nella Tecnica*, 189–224. Milan: Edizioni di Comunità.
- Mueller-Reichau, Olav. 2013. *Sorting the world: on the relevance of the kind/object distinction to referential semantics*, Linguistics and Philosophy. Berlin, Boston: De Gruyter.
- Nayak, Neha, Mark Kowarsky, Gabor Angeli and Christopher D. Manning. 2014. A dictionary of nonsubsecutive adjectives. Tech. rep., Technical Report CSTR 2014-04, Department of Computer Science, Stanford University.
- Oliver, Michael. 2014. Interpretation as Optimization: Constitutive material adjectives. *Lingua* 149. 55–73.
- Pagin, Peter and Francis Jeffrey Pelletier. 2007. Content, context, and composition. In G. Preyer and G. Peter (eds.), *Context-sensitivity and semantic minimalism: new essays on semantics and pragmatics*, 25–62. Oxford University Press.
- Panagiotidis, Phoevos. 2011. Categorical features and categorizers. *The Linguistic Review* 28. 325–346.
- Panagiotidis, Phoevos. 2014. A minimalist approach to roots. In Peter Kosta; Steven Franks; Lilia Schürcks; Teodora Radeva-Bork (ed.), *Minimalism and beyond: radicalizing the interfaces*, 287–303. Amsterdam: John Benjamins.
- Panagiotidis, Phoevos. 2015. *Categorical features: a generative theory of word class categories*. Cambridge University Press.
- Partee, Barbara. 1984. Compositionality. In F. Landman and F. Veltman (eds.), *Varieties of Formal Semantics* (GRASS 3), 281–311. Dordrecht: Foris.
- Partee, Barbara. 1997. Montague grammar. In J. van Benthem and Alice ter Meulen (eds.), *Handbook of logic and language*, 5–91. The MIT Press.
- Partee, Barbara. 2007. Compositionality and coercion in semantics: The dynamics of adjective meaning. In *Cognitive Foundations of Interpretation*. Royal Netherlands Academy of Arts and Sciences.
- Partee, Barbara. 2009. Formal Semantics, Lexical Semantics, and Compositionality: The Puzzle of Privative Adjectives. *Philologia* 7. 11–21.
- Partee, Barbara. 2010. Privative Adjectives: Subsecutive Plus Coercion. In Thomas Zimmermann Rainer Bauerle, Uwe Reyle (ed.), *Presuppositions and Discourse: Essays Offered to Hans Kamp*, 273–285. Brill.
- Pavlick, Ellie and Chris Callison-Burch. 2016. So-Called Non-Subsecutive Adjectives. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, 114–119.
- Pietroski, Paul M. 2018. *Conjoining meanings: semantics without truth values*. Oxford University Press.

- Prasada, Sandeep. 2008. Aspects of a fully psychological theory of sortal representation. Unpublished ms. New York: Hunter College, CUNY.
- Prinz, Jesse J. 2012. Regaining composure: a defence of prototype compositionality. In Markus Werning, Wolfram Hinzen and Edouard Machery (eds.), *The oxford handbook of compositionality*, Oxford Handbooks in Linguistics. Oxford University Press.
- Pustejovsky, J. 1991. The generative lexicon. *Computational Linguistics* 17(4). 409–441.
- Putnam, Hillary. 1970. Is semantics possible? *Metaphilosophy* 1(3). 187–201.
- Quilty-Dunn, Jake. 2021. Polysemy and thought: toward a generative theory of concepts. *Mind & Language* 36(1). 158–185.
- Rayo, A. 2013. A plea for semantic localism. *Nous* 47. 647–679.
- Recanati, François. 2010. *Truth-conditional pragmatics*. Oxford University Press.
- Reichard, Ulrich. 2013. Inference and grammar: Intersectivity, subsectivity, and phases. In Alison Henry Catrin S. Rhys, Pavel Iosad (ed.), *Microvariation, minority languages, minimalism and meaning: Proceedings of the irish network in formal linguistics*, 222–244. Cambridge Scholars Press.
- Reyes, J. Macnamara;M. La Palme Reyes;G.E. 1994. Reference, kinds, and predicates. In J. Macnamara;G.E. Reyes (ed.), *Logical foundations of cognition*, 91–142. OUP.
- Sadler, Louisa and Douglas J. Arnold. 1994. Prenominal adjectives and the phrasal/lexical distinction. *Journal of Linguistics* 30. 187–226.
- Sauerland, U., J. Anderssen and K. Yatsushiro. 2005. The plural is semantically unmarked. In S. Kepser and M. Reis (eds.), *Linguistic evidence: empirical, theoretical, and computational perspectives*, 413–434. Berlin/New York: Mouton de Gruyter.
- Sauerland, Uli. 2003. A new semantics for number. In *Proceedings of SALT 13*, 258–275.
- Sağ, Yağmur. 2019. *The semantics of number marking: reference to kinds, counting, and optional classifiers*. Ph.D. thesis, Rutgers University.
- Sağ, Yağmur. 2021. Bare singulars and singularity in turkish. *Linguistics and Philosophy* .
- Siegel, Muffy. 1976. *Capturing the adjective*. Ph.D. thesis, University of Massachusetts, Amherst, MA.
- Simpson, Andrew and Saurov Syed. 2016. Blocking effects of higher numerals in bangla: a phase based approach. *Linguistic Inquiry* 47(4). 754–763.
- Spector, Benjamin. 2007. Aspects of the pragmatics of plural morphology: on higher-order implicatures. In U. Sauerland and P. Stateva (eds.), *Presupposition and implicature in compositional semantics*, 243–281. Palgrave Macmillan.

- Svenonius, Peter. 1994. On the structural location of the attributive adjective. In E. Duncan, D. Farkas and P. Spaelti (eds.), *Proceedings of the twelfth west coast conference on formal linguistics (wccfl)*, 439–454. Stanford, CA: CSLI Publications.
- Svenonius, Peter. 2004. On the edge. In D. Adger, D.C. Cat and G. Tsoulas (eds.), *Syntactic edges and their effects*. Dordrecht: Kluwer.
- Svenonius, Peter. 2008. The position of adjectives and other phrasal modifiers in the decomposition of dp. In *Adjectives and adverbs: syntax, semantics, and discourse*, 16–42. Oxford University Press.
- Syed, Saurov. 2014. Adjective-fronting as evidence for focus and topic within the bangla nominal domain. In Pritha Chandra and Richa Srishti (eds.), *The lexicon-syntax interface: perspectives from South Asian languages*, 53–70. John Benjamins.
- Syed, Saurov. 2015. Focus-movement within the dp: Bangla as a novel case. In *Proceedings of the 32nd west coast conference on formal linguistics (wccfl)*, 332–341. Somerville, MA: Cascadilla Proceedings Project.
- Syed, Saurov and Andrew Simpson. 2017. On the dp/np status of nominal projections in bangla: consequences for the theory of phases. *Glossa: a journal of general linguistics* 2(1). 1–24.
- Szabó, Zoltán Gendler. 2012. The Case for Compositionality. In Markus Werning, Wolfram Hinzen and Edouard Machery (eds.), *The oxford Handbook of Compositionality*, Oxford Handbooks in Linguistics, 64–80. Oxford University Press.
- Talić, Aida. 2017. Long-form and short-form prenominal adjectives are not reduced relative clauses in Bosnian/Croatian/Serbian. *Linguistic Inquiry* 48(1). 194–212.
- Taylor, John R. 2003. *Linguistic categorization*. Oxford: Oxford University Press, 3rd edn.
- Vergnaud, Jean-Roger and Maria Luisa Zubizarreta. 1992. The definite determiner and the inalienable constructions in french and in english. *Linguistic Inquiry* 23(4). 595–652.
- Williams, Donald Cary. 1953. The elements of being. *Review of Metaphysics* 7(2). 3–18.
- Winter, Yoad and Joost Zwarts. 2012. Event orientated adnominals and compositionality. In *Proceedings of annual meeting of the israeli association of theoretical linguistics*. Tel-Aviv.
- Xu, F. 2005. Categories, kinds, and object individuation in infancy. In C. Gershkoff-Stowe; D. Rakison (ed.), *Building object categories in developmental time. papers from the 32nd carnegie symposium on cognition*, 63–89. New Jersey: Lawrence Erlbaum.
- Zadrozny, Wlodek. 1994. From compositional to systematic semantics. *Linguistics and Philosophy* 17(4). 329–342.
- Zamparelli, Roberto. 2000. *Layers in the determiner phrase*. Routledge.

Zobel, Sarah. 2017. The sensitivity of natural language to the distinction between class nouns and role nouns. In Dan Burgdorf, Jacob Collard, Sireemas Maspong and Brynhildur Stefansdottir (eds.), *Semantics and linguistic theory (salt)*, vol. 27, 438–458. Ithaca, NY: CLC.

Zweig, E. 2009. Number-neutral bare plurals and the multiplicity implicature. *Linguistics and Philosophy* 32(4). 353–407.