

The semantics of modification: Adjectives, nouns, and order

by

Timothy James Leffel

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Linguistics
New York University
September, 2014

Professor Chris Barker

this thesis is dedicated to Rags, the finest four-legged friend a fool could ever hope to have.



miss you pup <3

Acknowledgments

I always hoped my dissertation acknowledgments would be long and poetic and profoundly meaningful to all the people I have intellectual and personal debts to. But I've been writing a lot lately and my gas tank is running on E. So instead they are going to be short(ish!) and formulaic but hopefully still meaningful.

Thanks Dad, for exposing me to logic and philosophy and the pursuit of knowledge at such a young age—I look back fondly on the dinners you spent scrawling syllogisms on napkins and paper tablecloths to teach me propositional logic. As you can see, I haven't forgotten modus ponens! Thanks Mom, for always making me do my math homework as a kid, despite my claims that I hated it. Like you and especially your parents, I eventually learned that math is a beautiful thing—I just had to find the right kind. Thanks Amy, for making me a proud big brother. You are an intelligent, strong, and principled young woman. I doubt if many people look up to their younger siblings, but I do. I especially appreciate you keeping Mom and Dad sane over the years, in spite of my erratic and sometimes worry-inducing behavior. You're the best!

Thanks George Schumm, for talking me down from dropping your modal logic class freshman year. I'll never forget your advice (and I paraphrase): “just because something's hard or you don't understand it right away doesn't mean you can't figure it out with enough hard work. Hell, you and me could sit down and figure out quantum mechanics if we had to; it'd just take a lot of time and be really hard.” Thanks Carl Pollard, for showing me the exciting connections between logic and language. No doubt in my mind that I wouldn't be a linguist today were it not for the time and effort you invested in me as a young college student.

Thanks Chris Barker, for being an advisor, mentor, and friend to me over the last five years. Thanks for the stimulating discussions, consistently solid and thoughtful advice (both professional and personal), comments on papers, excellent dinners/drinks at your place (thanks also to Svetlana for that!), and for guidance through the ups and downs of grad school. No one is born knowing how to write a thesis or how to act in various professional environments, but thanks to your dedication and help I think that I am well on my way. Looking forward to interacting with you for years to come, as a (near) colleague but ultimately still as your student. Don't think I could have done it without you.

Thanks Anna Szabolcsi, for serving on every committee I've had at NYU, for delivering extremely thought-provoking feedback on every major project I've done here, for teaching some of the most exciting seminars I've been a part of, and for the vision of what NYU semantics could be that you must have come here with years ago. I can't imagine a better environment in which to receive my semantics training. It's

been really really real. Thanks Philippe Schlenker, for being so generous with your (seemingly limitless) time and resources and energy. I truly admire your ability to be several places at the same time and to provide great input on everything every student you've encountered has ever written (including the projects of mine you've directed or committee-d!). The field is lucky to have someone like you who's willing to challenge basic assumptions and seriously entertain exotic ideas with such intellectual rigor. Thanks Liina Pylkkänen, for initiating our collaboration on RestrMod, for teaching me how to convert an idea into a behavioral experiment, for helping me to understand in a much deeper and more hands-on way the nature of scientific progress, for all the input you've given on my dissertation project, and for your role in the amazing post-graduate opportunity I'm now embarking on. Couldn't be happier with it! Thanks Marcin Morzycki, for providing shockingly (in a very good way) thorough commentary and feedback on my dissertation, all without ever having met you in person. I'm looking forward to that changing very soon—see you around the world of semantics! Thanks also to everyone else—audiences, colleagues, etc.—for the collectively invaluable input you've given me on this and all my other research projects.

Vielen Dank to Malte Zimmermann and the SFB 632 folks—that research group is really something special. I'm extremely grateful that you took me on for a while; I learned a lot about languages and the world during my time in Berlin/Potsdam. See you guys around!

Thanks to all my grad student pals from NYU. You guys have greatly enriched the last five years of my life. Special thanks to my first-year class, and to Sally, Simon, Mike, Jeremy, and Dylan, a few of the folks from whom I've benefited most intellectually in addition to personally. Thanks also to numerous other NYU professors, students, and visitors—it would take forever to name you all so I'll just say you know who you are. Wouldn't have been the same being in the department without each one of you :)

And last but far from least, thanks to my friends. Gabe, Zach, Ryan, Danny, Jon, so many more I want to name but can't without violating my own pet peeve of acknowledgments containing very long lists of names. My Westerville crew, 15th/Lakeview boys, COJ buddies, Berlin people, my friends in New York. Johnny, you've never had a younger sibling and I've never had an older one—I imagine the friendship we've developed in the last five years approximates what it feels like to have a big brother. Thanks for making our household such a great environment for both work and play. Special thanks to Dootza, for your love and encouragement and perseverance and patience with me and your loyalty and your giggles and your intellect and your ability to turn my frown upside down. Love you bb <3

Dissertation Abstract

This dissertation motivates, defines, and applies a simple but general theory of nominal modification for natural languages. Chapter 2 presents two sub-theories that jointly constitute the bulk of the proposal: the first specifies a morpho-semantics for nouns and inflectional features. In a sentence, it says that lexical count nouns denote predicates of kinds, and that the composition of inflectional features with a noun converts the noun's kind-based meaning into a predicate of individuals, the extension of which reflects whether that noun is singular or plural, masculine or feminine, etc. The second sub-theory specifies the syntactic structure of nominal modification. Non-appositive modifiers are integrated into nominals either via adjunction to the nominal head prior to the introduction of inflectional features (*head-adjunction*); or via adjunction to the inflected noun, whose denotation has already been converted into a predicate of individuals (*phrasal adjunction*). Directionality parameters are proposed for each kind of adjunction: in English, head-adjunction is to the left only, while phrasal adjunction is bi-directional; in Italian head-adjunction is bi-directional while phrasal adjunction is to the right only. As with all parametric theories, typological predictions can be extracted from the analysis.

The theory of Chapter 2 is applied to a selection of modification-related phenomena in the remainder of the dissertation, focusing on “Bolinger contrasts” and the approach to modification in DP that they have inspired. This approach, “the two-domains theory,” provides indispensable basic insights, but I argue is in need of more concrete and well-motivated semantic foundations. Through detailed investigation of the restrictive/non-restrictive opposition in nominal modifiers (Chapter 3), the direct/implicit relative opposition in certain modal attributive adjectives (Chapter 4), and a selection of additional modification-related puzzles (Chapter 5), I aim to establish the theory in Chapter 2 as a concrete, principled, and cross-linguistically adaptable framework for analysis of the lexical- and morpho-semantics of nouns, and of noun phrase-internal composition. By accounting for the distribution of (some) Bolinger contrasts with a more conservative set of assumptions, specific grammatical enrichments and silent elements postulated by two-domains theories are rendered largely unnecessary. While data from English are examined most thoroughly throughout, the general architecture of the theory I propose can be instantiated to yield a theory of modification for any language. In Chapters 3 and 5, discourse coherence relations are argued to play an essential role in certain quasi-grammatical phenomena related to modification, suggesting that the study of modification requires attention to discourse-level pragmatic processes.

Contents

| | |
|--|-----------|
| Acknowledgments | iv |
| Dissertation Abstract | vi |
| 1 Introduction | 1 |
| 1.0 What this dissertation is about | 1 |
| 1.1 Bolinger contrasts and “two-domains” theories | 2 |
| 1.1.1 The core data patterns | 2 |
| 1.1.2 Existing approaches: the dual source model and the notion of “two domains” | 6 |
| 1.1.3 Preview: simulating two domains without dual sources | 8 |
| 1.2 Theoretical assumptions and formal preliminaries | 9 |
| 1.2.1 Syntactic assumptions | 10 |
| 1.2.2 Semantic assumptions | 12 |
| 1.3 Dissertation overview and main chapter summaries | 15 |
| 1.3.1 Overview of the dissertation in 289 words | 15 |
| 1.3.2 Chapter 2: Motivating and building a theory of nominal modification | 16 |
| 1.3.3 Chapter 3: Restrictive versus non-restrictive modification | 17 |
| 1.3.4 Chapter 4: Direct versus implicit relative modal adjectives | 19 |
| 1.3.5 Chapter 5: Modified names, subtriggering, and manner implicatures | 20 |
| 2 The internal syntax and semantics of noun phrases | 23 |
| 2.0 Introduction | 23 |
| 2.1 Motivation for a fine-grained theory of nominal reference | 24 |
| 2.2 Some recent advances in DP syntax and semantics | 26 |
| 2.2.1 Kind- and individual-reference in common noun denotations | 26 |
| 2.2.2 Event- and kind-relativity in adjectives and nouns | 28 |
| 2.2.3 The semantics of number and the unmarked status of PL | 30 |
| 2.2.4 DP layers and the type/token distinction in nominal reference | 32 |
| 2.2.5 Kind reference and cross-linguistic considerations | 36 |
| 2.2.6 Summary of proposals to enrich NP/DP syntax and semantics | 40 |
| 2.3 The parallel between type-reference in N and generic readings of inner modifiers | 40 |
| 2.4 Building a theory of composition in DP, Part 1: Nouns and determiners | 43 |
| 2.4.1 Basic assumptions about noun meanings and plural/singular | 43 |

| | | |
|----------|---|------------|
| 2.4.2 | A dataset for theory building | 46 |
| 2.4.3 | The interpretation of plurals | 47 |
| 2.4.4 | The interpretation of non-plurals | 50 |
| 2.4.5 | A theory of nouns, summarized | 54 |
| 2.5 | Building a theory of composition in DP, Part 2: Modification | 55 |
| 2.5.1 | A theory of nominal modification | 56 |
| 2.5.2 | Individual-level versus stage-level attributive adjectives | 58 |
| 2.5.3 | Why ban rightward head-adjunction (in English)? | 62 |
| 2.6 | Conclusion | 63 |
| 3 | Case study: restrictive versus non-restrictive modification | 64 |
| 3.0 | Introduction | 64 |
| 3.1 | What are restrictive and non-restrictive modification (not)? | 65 |
| 3.1.1 | Restrictive versus non-restrictive relative clauses | 66 |
| 3.1.2 | Restriction = _{df} $[[\text{modifier}(\text{noun})] \subset [\text{noun}]]?$ | 67 |
| 3.1.3 | Heim & Kratzer (1998): Non-restrictiveness \approx Apposition? | 69 |
| 3.1.4 | Umbach (2006): Restriction \approx Focus? | 72 |
| 3.1.5 | Non-restrictive modifiers as conventional implicature triggers? | 73 |
| 3.1.6 | Further considerations: relevance and predicational variability | 75 |
| 3.1.7 | Wrapping up §1 | 80 |
| 3.2 | Defining (<i>non-</i>) <i>restrictive</i> in terms of discourse coherence establishment | 80 |
| 3.2.1 | Crash course on theoretical discourse semantics | 80 |
| 3.2.2 | What are the minimal discourse units? An argument from sentence-internal coherence | 87 |
| 3.2.3 | Defining a simple version of discourse semantics | 89 |
| 3.2.4 | Defining <i>restrictive</i> and <i>non-restrictive</i> in terms of discourse coherence | 95 |
| 3.2.5 | Testing the definitions on some basic cases | 97 |
| 3.3 | Revisiting theories of R/NR and the two-domains model | 100 |
| 3.3.1 | Some reasons to suspect that R/NR is of a different nature than Bolinger contrasts | 101 |
| 3.3.2 | Two-domains approaches to the data, and where they go wrong | 106 |
| 3.4 | A new syntax-semantics for non-restrictive adjectives | 113 |
| 3.4.1 | Concluding remarks on R/NR and its relation to the two-domains model | 120 |
| 3.5 | Discussion and conclusion | 120 |
| 3.5.1 | Psycholinguistic connections | 120 |
| 3.5.2 | A constrained theory of coherence relations for non-restrictive modification? | 122 |
| 3.5.3 | Operationalizing “non-restrictive” for behavioral stimulus creation | 123 |
| 4 | Case study: an ambiguity in modal adjectives | 126 |
| 4.0 | Introduction | 126 |
| 4.1 | Background on <i>possible</i> : distribution and approaches | 128 |
| 4.1.1 | Restrictions on the distribution of indirect readings | 128 |

| | | |
|----------|---|------------|
| 4.1.2 | The ambiguity as evidence for the two-domains theory | 133 |
| 4.1.3 | Summary and critical review of existing analyses | 134 |
| 4.2 | The independence of direct/indirect <i>possible</i> from the two-domains model | 137 |
| 4.2.1 | Direct modal adjectives are not restricted to inner positions | 137 |
| 4.2.2 | Outer position and reduced relative structure do <i>not</i> imply ‘implicit relative reading’ | 140 |
| 4.3 | The internal structure of implicit relative modal adjectives in English | 141 |
| 4.3.1 | Modality type restrictions as evidence for covert structure | 141 |
| 4.3.2 | Arguments that ▲ is always a degree relative | 142 |
| 4.3.3 | The syntax of DP-internal infinitives, and explaining the locality restriction | 147 |
| 4.4 | Formalizing the new analysis of implicit relative modals in English | 151 |
| 4.4.1 | Deriving implicit relative readings in universals | 152 |
| 4.4.2 | Implicit relative readings with superlatives (‘modal superlatives’) | 154 |
| 4.4.3 | Definites: why are <i>possible</i> and <i>necessary</i> different? | 157 |
| 4.4.4 | Deriving the Bolinger distribution without a two-tiered DP | 161 |
| 4.5 | Discussion and conclusion | 165 |
| 5 | Related issues and applications | 166 |
| 5.0 | Introduction | 166 |
| 5.1 | Composition problems and modified names (and beyond) | 167 |
| 5.1.1 | Approaches to the composition problem | 168 |
| 5.1.2 | Complex names | 170 |
| 5.1.3 | Some putative cousins of non-restrictive modification | 175 |
| 5.2 | Subtriggering: modifier position and discourse coherence | 177 |
| 5.2.1 | <i>Any</i> in positive contexts | 178 |
| 5.2.2 | Ambiguity versus discourse coherence in subtriggering | 180 |
| 5.2.3 | A counterexample to the extra-parameter analysis (Dayal 1995,1998) | 182 |
| 5.3 | Structural implicatures and non-restrictive modifiers | 185 |
| 5.3.1 | Scalar implicatures, manner implicatures, and structural alternatives | 185 |
| 5.3.2 | The puzzle: Manner implicatures of (non-)restrictive modifiers | 190 |
| 5.3.3 | Three analysis sketches | 191 |
| 5.3.4 | A final puzzle: Implicatures of pre- versus postnominal modifiers | 193 |
| 5.3.5 | Predictions of the three analyses for the linear order puzzle | 196 |
| 6 | Conclusions | 201 |
| | Bibliography | 204 |

Chapter 1

Introduction

1.0 What this dissertation is about

Since the early work of Dwight Bolinger and Zeno Vendler, linguists have recognized the existence of elusive but undeniable connections between the relative order of words in a modified noun phrase and the range of interpretations such a phrase can have.¹ Moreover, some manifestation of this kind of connection exists in a wide range of the world's languages, suggesting that beneath such patterns lie deep and fairly general facts about language. The goal of this dissertation is to use order-sensitive ambiguities in modification structures—"Bolinger contrasts"—to inform the theory of modification in natural language, as well as the theory of nominal structure and interpretation more generally.

A particularly clear example of a Bolinger contrast, due to Larson (2000a), goes like this: the sentence *Mary interviewed every possible candidate* can mean one of two things, (i) that Mary interviewed everyone who was possibly a candidate; or (ii) that Mary interviewed every actual candidate that it was possible to interview. The effect of word-order on interpretation reveals itself when the linear order of *candidate* and *possible* is reversed: *Mary interviewed every candidate possible* naturally means something like (ii), but is incompatible with interpretation (i). This example is not an isolated case; English displays a range of modification ambiguities with distributional profiles very similar to this one's. And English is by no means unique in this respect. Most if not all Romance languages disallow certain interpretations of certain prenominal attributive adjectives that are readily available in postnominal position (and frequently vice versa). Korean relative clauses are interpreted intersectively when appearing between a determiner and a noun, but appositively when appearing noun phrase-initially (Kim 1997). The list goes on. I would go as far

¹ By "modified noun phrase" will be meant any noun phrase containing one or more attributive adjectives (*the big dog*), relative clauses (*the dog that I petted*), locative PPs (*the dog in the yard*), etc.

as to conjecture that no language permitting both pre- and postnominal modification is such that the relative order of a noun and its modifiers can be freely permuted without ever causing any interpretive consequences.

To the extent that the proposals of this dissertation can be summarized in a single sentence, the central thesis I will advance is that the defining properties of Bolinger contrasts—both distributional and semantic—follow from a small set of independently motivated assumptions about the lexical semantics and morpho-semantics of nouns, along with the distinction between phrasal- and head-adjunction as articulated by Sadler & Arnold (1994). This idea stands in contrast to influential “two-domains” theories of Bolinger contrasts, according to which certain regions of the noun phrase automatically associate specific interpretive properties with modifiers occupying those regions (how this association is achieved differs from theory to theory; see especially Larson 1998,1999; Larson & Marušič 2004; Cinque 2010,2014). I argue that the results of these theories for English can be largely replicated—and in some cases extended—with a more conservative set of principles that govern the interpretation of nouns instead of modifiers themselves. Some cross-linguistic differences in Bolinger contrasts are also captured with these assumptions, supplemented by a small set of syntactic parameter settings.

1.1 Bolinger contrasts and “two-domains” theories

Analyses of individual Bolinger contrasts have been advanced by semanticists, while a number of more general theories of Bolinger contrasts have been advanced by syntacticians. I will argue throughout this dissertation that both approaches are in fact necessary (and complementary), but also that there are critical bodies of data that no existing theory can fully capture. Semantically-oriented analyses are typically designed for a specific delimited dataset, while syntactically-oriented theories tend to be more general but lack concrete, testable hypotheses about meaning. This section will survey the landscape of theories and results about Bolinger contrasts and the syntax-semantics mapping in nominal modification more generally. By identifying gaps in current theoretical understanding, we simultaneously establish motivation for a theory of the kind I will develop in subsequent chapters.

1.1.1 The core data patterns

The most comprehensive study of Bolinger contrasts to date is the recent monograph of Cinque (2010). In that study, Cinque describes nine contrasts which are argued to form a syntactic natural class in English and in Italian. Each contrast can be described in terms of three general characteristics: *(i)* a semantic dimension along which two truth-conditionally distinct interpretations exist; *(ii)* a set of syntactic configurations; and

(iii) possible values of (i) for each specification of (ii). Take the example from §1.0, for instance, reproduced in 1.1 with labels for the opposing readings.

- (1.1) a. Mary interviewed every possible candidate. (Larson 2000a:1)
✓*Direct*: “Mary interviewed everyone who was possibly a candidate.”
✓*Indirect*: ‘Mary interviewed every candidate it was possible to interview.’
- b. Mary interviewed every candidate possible.
✗*Direct*, ✓*Indirect*

For this contrast, the parameter would be “direct versus indirect reading of *possible*,” and the configuration-value pairings would be Adj-Noun → {direct, indirect}; and Noun-Adj → {direct}. This strategy has been shown to reveal some interesting cross-linguistic generalizations. For example, the contrast in 1.1 can be replicated in Spanish and Italian along the same semantic dimension, but with different configuration-value pairings. The following minimal pair, due to Demonte (2008), suggests that in Spanish the Adj-Noun order has only the direct interpretation while the Noun-Adj order has only the indirect interpretation.

- (1.2) a. Atendió a todos los visitantes posibles.
attended.3sg to all the visitors possible
‘He received all the visitors it was possible for him to receive.’
- b. Atendió a todos los posibles visitantes.
attended.3sg to all the possible visitors
‘He attended to all the people that were possible visitors.’ (Spanish)

Even apart from cross-linguistic considerations, though, there are a number of reasons why Bolinger contrasts are interesting empirically and theoretically. Empirically, this topic is of interest because a wide range of logically independent interpretive properties of modifiers appear to pattern alike, as a function of structural position in the noun phrase. Theoretically, these contrasts raise questions like the following: are adjectives lexically ambiguous or are the phrases they appear in syntactically ambiguous? If the latter, then precisely what piece of the structure is responsible for endowing adjectives with the semantic properties they have in a particular configuration?

Another example, which along with the direct/indirect opposition in 1.1 constitutes the primary empirical target of this dissertation, is the opposition between so-called “restrictive” and “non-restrictive” interpretations of attributive adjectives. As can be seen from the oft-cited minimal pair in 1.3, the distribution of the non-restrictive reading coincides with the distribution of the direct reading of *possible* in example 1.1: only

the order in (1.3a) is compatible with a situation in which all of the relevant words are taken to be unsuitable.

- (1.3) a. Every unsuitable word was deleted. (Larson & Marušič 2004:275)
✓*Restrictive*: “Every word that was unsuitable was deleted.”
✓*Nonrestrictive*: “Every word was deleted; they were unsuitable.”
- b. Every word unsuitable was deleted.
✓*Restrictive*, ✗*Nonrestrictive*

The full range of Bolinger contrasts observed in the literature is vast, and by no means limited to these two particular phenomena. The pairs in 1.4, for example, have slightly different possible meanings, which many researchers have related to the i(ndividual)-level/s(tage)-level distinction.

- (1.4) a. the visible stars, the stars visible (Bolinger 1967:3-4)
b. the guilty people, the people guilty
c. the stolen jewels, the jewels stolen
d. the responsible individuals, the individuals responsible

A similar and perhaps clearer contrast exists for temporal modifiers: the alternation in 1.5 reflects the fact that *Thursday meeting* can refer to a regular meeting whose default day is Thursday as well as to a specific meeting on a specific Thursday, whereas *meeting Thursday* can only refer to a single Thursday’s meeting.

- (1.5) a. Our Thursday meeting will take place on Wednesday (this week).
b. #Our meeting Thursday will take place on Wednesday (this week).

(Non-)restrictiveness, direct/indirect modal modification, and stage-/individual-level are logically independent parameters of meaning, but the alternations appear to pattern syntactically in a parallel fashion: given an appropriate set of lexical items, in English both readings are (usually) available when the adjective appears prenominal, but only one reading is available postnominally.

Another striking property of Bolinger contrasts is that their syntactic manifestation varies across languages, but does so systematically. Some descriptions of Spanish and Italian report a “mirror-image” distribution, in which it is the *postnominal* modifier position that is ambiguous, and the readings that are only available prenominal in English (non-restrictive, i-level, etc.) are the only readings prenominal. From Cinque 2010:Ch2:

- | | |
|---|--|
| <p>(1.6) a. le <u>stelle invisibili</u> di Andromeda the stars invisible of A. ✓<i>Individual-level</i>, ✓<i>Stage-level</i></p> <p>b. le <u>invisibili stelle</u> di Andromeda ✓<i>Individual-level</i>, ✗<i>Stage-level</i></p> | <p>(1.7) a. le <u>lezioni noiose</u> di Ferri the classes boring of F. ✓<i>Restrictive</i>, ✓<i>Nonrestrictive</i></p> <p>b. Le <u>noiose lezioni</u> di Ferri ✗<i>Restrictive</i>, ✓<i>Nonrestrictive</i></p> <p>(Italian, Cinque 2010:7-8)</p> |
|---|--|

Bolinger effects are also observed with multiple modifiers. For example, (1.8b) is odd because it implies that generally non-visible stars could be temporarily visible. That (1.8b) does not have a reading equivalent to (1.8a) shows that in certain cases, s-level readings must precede i-level readings.

- | | |
|--|--|
| <p>(1.8) a. The <u>invisible visible</u> stars include Capella.</p> <p>b. #The <u>visible invisible</u> stars include Capella.</p> | <p>(Larson 1998, attributed to B. Citko)</p> |
|--|--|

Perhaps more clearly, the same point can be made using the “intersective/adverbial” alternation, which exists for certain manner adjectives (Larson 1998). The basic ambiguity is illustrated in 1.9. The pair in 1.10 shows that the adverbial reading of *beautiful* (“dances beautifully”) is blocked if a non-adverbial adjective intervenes between *beautiful* and the noun. This is precisely the distribution of i-level versus s-level readings of *visible* suggested by example 1.8.

- (1.9) Olga is a beautiful dancer.
- Adverbial*: “Olga dances beautifully.” *Intersective*: “Olga is beautiful and a dancer.”

- | | |
|--|---|
| <p>(1.10) a. Olga is a blond beautiful dancer.</p> <p>b. Olga is a beautiful blond dancer.</p> | <p>(both readings of <i>beautiful</i> are available)</p> <p>(only intersective reading of <i>beautiful</i> available)</p> |
|--|---|

Theoretically, observations like the ones surveyed here have been taken to constitute powerful evidence for the “two-domains” model of nominal modification, which holds that regions of DP grammatically associate specific interpretive properties with modifiers occupying those regions (Larson (1998,1999); Cinque (2010); a.o.). Implementations of the two-domains model posit a range of grammatical enrichments that are meant to explain why, for instance, pronominal modifiers in English tend to be ambiguous between the two readings but pronominal modifiers in (some) Romance languages do not. Bolinger effects pattern non-uniformly across languages, so these conclusions have significant implications for nominal syntax and semantics cross-linguistically. The next subsection summarizes in some detail the respective theories of Larson and Cinque.

1.1.2 Existing approaches: the dual source model and the notion of “two domains”

In a series of papers and lectures (Larson 1998; Larson 1999; Larson 2000a; Larson 2000b; Larson & Cho 2003; Larson & Marušič 2004; Larson & Takahashi 2007; a.o.), Richard Larson developed an important version of the two-domains theory of nominal modification. Larson’s central proposal is that there are two distinct regions inside the DP in which attributive adjectives may be syntactically positioned, and that the semantics of an attributive adjective is (partially) determined by which region of the DP it occupies. Cinque (2010) advanced a syntactically-oriented alternative formulation of the two-domains theory, drawing heavily on English data discussed by Larson and supplemented with corresponding examples from Italian.

Both Larson and Cinque distinguish “direct” from “indirect” modification, a distinction initially introduced by Sproat & Shih (1988) with a similar meaning. On Larson’s and Cinque’s conception, direct modifiers are structurally close to the noun and are associated with non-restrictive/*i*-level/etc. readings. Indirect modifiers are more distant from the noun and are associated with the opposing set of readings. While the formalisms employed to achieve this result are different, Larson and Cinque’s theories both guarantee that the syntactic position of a modifier—and its structural relationship to the noun—constrains its semantics.

On Cinque’s (2010,2014) “dual-source hypothesis,” indirect modifiers are always adjoined reduced relative clauses (with presumably intersective semantics). Direct modifiers are analyzed as specifiers of functional heads which endow adjectives with interpretive properties such as non-restrictive/*i*-level/etc., parallel to the theory of adverbs advanced in Cinque 1999. This means that the direct readings are derived by silent functional items, while the indirect readings are equivalent to predicative uses (since indirect adjectives are predicates of covert relatives).

On Larson’s (1998,1999) theory, direct modifiers are inside NP while indirect modifiers are external to NP. DP contains two silent event quantifiers with different quantificational forces and fixed structural positions. The interpretation of an adjective, on this theory, depends on which quantifier it is in the scope of: the generic quantifier can affect inner adjectives only, and the existential quantifier, outer adjectives.

(1.11) *Two theories of nominal modification based on the two-domains model*

- a. $[_{DP} D [AP_{indirect} [[_{XP} AP_{direct} [_{X'} X NP]] AP_{indirect}]]]$ (Cinque 2010)
(functional head X imposes semantics on AP_{direct})
- b. $[_{DP} \exists e [AP_{indirect} [\Gamma e [_{NP} AP_{direct} N]] AP_{indirect}]]$ (Larson & Marušič 2004)
(Only AP_{direct} is in the scope of generic event quantifier Γe)

There is only one postnominal position in (1.11a)/(1.11b), which on both theories is meant to explain why English Noun-Adj sequences do not display the same kind of ambiguities as do Adj-Noun sequences. The existence of two prenominal domains is meant to explain ordering constraints like those in 1.8 and 1.10.

Cinque (2010) extends the two-domains analysis to Italian, which he argues displays the “mirror-image” distribution of English (see above): the (default) postnominal position hosts both readings, while the prenominal position hosts only the direct readings (diagrams simplified here).

(1.12) *Mirror-image distribution of attributive adjective interpretations (Cinque 2010)*

- a. *English:* [DP AP_{indirect} AP_{direct} NP AP_{indirect}]
- b. *Italian:* [DP AP_{direct} NP AP_{direct} AP_{indirect}]

In sum, then, a two-domains theory of modification holds that attributive adjectives can occupy two kinds of structural positions: the first kind of position (for “direct modifiers”) is close to the noun and prenominal (in English); the second (for “indirect modifiers”) can be much higher than the noun and prenominal or postnominal (in English). These two hypothesized structural tiers will be referred to as “domains of modification” throughout. Importantly, on both Larson’s and Cinque’s versions of the theory, the direct and indirect prenominal domains are rigidly ordered with respect to one another, so that a sequence of the form

(1.13) * [Adj_{direct} [Adj_{indirect} [Adj_{direct}...N...]]]

is predicted to be impossible in all languages. On the other hand, a sequence of two modifiers in the same domain followed by a modifier from a different domain is a possible configuration:

(1.14) [Adj_{indirect} [Adj_{direct} [Adj_{direct}...N...]]]

One way to think about the two-domains theory is by analogy to how the structure of clauses is broken up in syntactic theory: in the C-domain are elements semantically associated with clause-typing and information structure; in the I/T-domain are inflectional elements associated with tense and agreement; and in the V-domain are elements semantically associated with argument structure and thematic roles. In the present case, domains of modification are distinguished from one another by the semantic properties associated with them: indirect modifiers tend to have intersective interpretations, whereas the interpretation of a direct modifier can be more idiosyncratic and does not always fall out in an obvious way from ordinary composition.

1.1.3 Preview: simulating two domains without dual sources

The main proposal of this dissertation is the following: regardless of whether or not there are two distinct domains of modification in the noun phrase, the distribution of Bolinger contrasts follows (for the most part) from the lexical semantics and morpho-semantics of nouns. Put another way, we do not need to introduce silent structure or covert operators into the noun phrase to understand how or why (English) prenominal modification structures are ambiguous in ways that corresponding postnominal structures are not. Instead, I argue that our understanding of structure-sensitive ambiguities in modification will automatically progress as we gain deeper understanding of composition in noun phrases that do *not* contain modifiers.

Two-domains theories postulate silent structures or operators that affect the interpretation of modifiers only. For example on Cinque's theory, direct modification heads introduce direct modifiers and endow them with certain properties. But a direct modification head never affects the interpretation of a noun, and is not even present in the nominal structure unless a direct modification adjective is. Similarly, Larson's event quantifiers bind variables in the representations of adjectives but only sometimes do these operators simultaneously bind corresponding variables in the noun. On the theory that I develop here, however, there is a single set of processes that interpret nouns and noun phrases, and these processes apply uniformly regardless of whether a modifier is present or not. This approach is more theoretically conservative, but still captures the core generalizations about Bolinger contrasts. Just like in two-domains theories, the position in which a modifier is introduced into the structure does affect how that structure is interpreted; however, on the theory I advance here, this interaction is a side-effect of how nouns and noun phrases are interpreted—there are no specific principles governing what kind of reading a given modifier in a given structure should have. It is in this sense that the proposals of this dissertation “simulate” the results of the two-domains model of modification without the introduction of mechanisms solely governing the interpretation of modifiers.

The existence of “reduced relative clause” structure for attributive adjectives is a contested issue among syntactically-inclined semanticists, and a popular position within syntactic theory (Kayne 1994; Leu 2008; Cinque 2010). Throughout the dissertation, we encounter situations in which researchers have posited reduced relative structure to solve various puzzles. In some cases, I suggest that such puzzles can be solved without appeal to this kind of covert syntax. But in others, the existence of covert relative structure appears to have both empirical motivation and a crucial role to play in the derivation of generalizations. The results of this dissertation are therefore partially but not entirely independent from the issue of whether adjectival

modifiers contain relative clause structure.

1.2 Theoretical assumptions and formal preliminaries

A natural language like English can be viewed as an abstract combinatorial system analogous to formal languages of logic and mathematics. This is the intuition upon which the foundations of generative syntax and formal semantics for natural language were built, as reflected in the early work of Chomsky (1955) and Montague (1974), respectively. Three important similarities between natural and formal languages are noted here. First, both kinds of systems have a determinate *syntax*; i.e. a set of primitive symbols which can be recursively combined with one another according to some set of rules or constraints. Second, primitive symbols (e.g. morphemes or logical constants) are conventionally associated with *non*-linguistic objects or concepts in the world (or in imagination). Third and finally, both natural and formal languages are *compositional* in the sense that the meaning of any syntactically well-formed expression is uniquely determined by the meanings of its constituent units, the syntactic structure of the expression, and the operations used to combine the constituent units' meanings. These similarities make it possible to construct models of natural language syntax and semantics with precise and well-understood tools from formal language theory and mathematical logic. Furthermore, the syntactic and semantic components can be designed to interface with one another, resulting (ideally) in a precise and predictive theory of a language's syntax-semantics interface.

A concrete example illustrating the “intersective” analysis of adjectival modification: suppose a new expression of category N (“an N”) can be formed by concatenating an A with an N. To compute the meaning of a complex N $[_N [_A a] [_N n]]$ constructed by this rule, we need to specify the lexical meanings of a and n as well as a *composition principle* that can be used to combine their meanings—roughly, a way of interpreting concatenation. To this end, suppose lexical A's and N's are required to have sets as their meanings, and suppose there is a composition principle of *intersection* (**int**) which says that the meaning of the concatenation $\alpha \frown \beta$ of two set-denoting expressions α and β is equal to the intersection of α 's denotation and β 's denotation. Then we can compute a complex N's meaning as follows:

$$\begin{aligned}
 \llbracket [_N [_A a] [_N n]] \rrbracket & && (\llbracket x \rrbracket = \text{the meaning of 'x'}) \\
 = \llbracket a \frown n \rrbracket & && (a \text{ and } n \text{ are sisters}) \\
 = \mathbf{int}(\llbracket a \rrbracket)(\llbracket n \rrbracket) & && (a \text{ and } n \text{ denote sets}) \\
 = \llbracket a \rrbracket \cap \llbracket n \rrbracket & && (\text{by definition of } \mathbf{int})
 \end{aligned}$$

According to this toy example, adjective and noun meanings are things that are true or false of objects,


and thus **int** predicts that $[_N [_A a][_N n]]$ is true of an object if and only if both $[_A a]$ and $[_N n]$ are true of that object. Such a prediction is then tested by checking commonsense intuitions about entailment: does something count as a $[_N [_A yellow] [_N duckling]]$ if and only if it counts as a duckling and is yellow? Plausibly. But this simplistic model makes incorrect predictions for cases like *rubber duckling* or *fake duckling*, which intuitively aren't kinds of ducklings at all. The theory must therefore be revised to reflect these intuitions.

Analyses in contemporary formal semantics have the methodological shape of this example: define a set of primitive category-word pairs (build a *lexicon*); assign meanings to those words (give *lexical semantics*); adopt a set of syntactic principles that generate complex hierarchical expressions from the words and their categories (write a *grammar*); and specify how to compute the meaning of an expression containing sub-expressions whose meanings are known (define the *composition principles*). The remainder of this section defines the formalisms that I will be using in this dissertation as models of the lexicon, lexical semantics, grammar, and composition principles.


1.2.1 Syntactic assumptions

This dissertation is not a work in theoretical syntax, but issues relating to phrase structure, constituency, morphology, and agreement are important in many of the analytical issues to be encountered. Therefore, I will uniformly use the terminology and notation of X-bar theory, a constrained format for representing syntactic structure that has been the dominant theoretical paradigm in generative linguistics for several decades now. For concreteness, the allowable forms for phrase- and bar-level projections are specified in 1.15 and 1.16, respectively.

(1.15) **Phrasal projections:**



(1.16) **Bar-level projections:**



The constraints imposed by 1.15 and 1.16 are complemented by *selectional properties* of lexical items. For example the transitive verb *talk* selects for PP, and not DP/NP, complements. This is why the sentence

**John* [_{VP} *talked* [_{NP} *Mary*]] is ungrammatical while *John* [_{VP} *talked* [_{PP} *to Mary*]] is not, even though both conform to 1.15-1.16. Selectional information can be about syntactic category, thematic role, inflectional morphology, etc. Thus the selectional properties of the D *these* explains why *these dogs* is well-formed but **these dog* is not.

Cross-linguistic variation is captured in X-bar theory by so-called parameter settings. As indicated by rule 1.16, English tends to be a head-initial language in the sense that heads tend to linearly precede their complements. Other languages have head-final tendencies, resulting in e.g. postpositions instead of prepositions, OV instead of VO word-order, postnominal instead of prenominal articles, etc. There is a vast literature on cross-linguistic syntax in the framework of X-bar theory, spanning decades, investigating dozens of languages, and proposing countless reformulations of the basic schemata in 1.15-1.16. Given the impossibility of taking into account the results of all research in this tradition, here I will simply stick to the above rules, revising and/or expanding the list whenever motivation to do so is encountered. This includes discussions of cross-linguistic variation; in Chapter 2, for example, I propose that the operation of head-adjunction is more constrained in some languages than it is in others.

To see how the concepts introduced above work in action, we close this discussion by providing a syntactic derivation of an English sentence of moderate complexity. The sentence is 1.17.

(1.17) I believe that the cop poked John with a stick.

Clauses are projections of the functional category T(ense), subjects occupy the specifier of T, and T selects for a VP complement. T is an abstract, phonetically empty functional item (although it is morphologically realized on inflected verbs). Thus the matrix clause has the shape

(1.18) [_{TP} [_{NP} I] [_{T'} [_T PRS] [_{VP} [_V believe] [...]]]]


Finite embedded clauses such as the complement of *believe* in 1.17 are analyzed as projections of C(omplementizer), which takes TP as its complement. The PP *with his stick* is adjoined to the embedded VP. Furthermore, I adopt the currently popular assumption that noun phrases containing determiners like *the* or *every* are projections of D instead of N (see Szabolcsi 1987; Abney 1987). According to this assumption, *the dog* has the structure [_{DP} [_D the] [_{NP} [_{N'} [_N dog]]]] instead of [_{NP} [_D the] [_N dog]]. Irrelevant bar-level projections will often be omitted, e.g. the structure of *the dog* can be abbreviated as [_{DP} [_D the] [_N dog]]. Thus the structure of the constituent abbreviated '[...]' in 1.18 is

(1.19) [CP [C that] [TP [DP [D the] [NP cop]] [T' [T PAST]
 [VP [VP [V poked] [NP John]] [PP [P with] [DP [D a] [NP stick]]]]]]]]]]

And this concludes the background on theoretical syntax. Additional constraints, principles, assumptions, etc., will be motivated and introduced as needed throughout subsequent chapters.

1.2.2 Semantic assumptions

In the syntax-semantics interface I will assume here, we superimpose semantic derivations on top of syntactic derivations by labeling each terminal node in a tree with a logical term that models the lexical semantics of the item occupying that terminal node. A simplified toy example to clarify the strategy: *John* has as its lexical semantics the non-logical constant **john**, and the VP *left* denotes a one-place predicate that is true of an individual if and only if there is a past time during which that individual left—in the semantic language defined below this is expressed by the term $\lambda x[\text{PAST}[\mathbf{left}(x)]]$. Because the argument type for the VP-meaning matches the type of the subject-meaning, the two terms can compose via function-argument application, yielding the boolean formula shown below as the meaning of full sentence *John left*.

(1.20) 

The resulting term is then interpreted relative to a particular context of use, and receives a value of true or false depending on the properties of the model against which it is interpreted.

Throughout the dissertation I formulate semantic analyses in a slightly modified version of the higher-order language Ty_2 , defined by Gallin (1975). Ty_2 was introduced as a reformulation of Montague's (1974) intensional logic, the primary difference being that Ty_2 has three distinct basic types instead of only two. The basic types of Ty_2 are e (for individuals), s (for states or worlds), and t (truth-values). Terms are used to model natural language meanings and types constrain the possible denotations of terms. The set **type** of Ty_2 types is the closure of the set of basic types $\{e, s, t\}$ under the binary type-constructor $\langle \cdot, \cdot \rangle$. For example e ; $\langle s, \langle e, t \rangle \rangle$; $\langle t, s \rangle$; and $\langle e, \langle e, \langle \langle \langle s, e \rangle, t \rangle, e \rangle, s \rangle, t \rangle$ are all types.

Terms consist of variables, constants, and whatever can be constructed from them via the term-formation rules specified in 1.21. Beginning with the simplest terms, for each type α , we assume a denumerable stock of variables x_1, \dots, x_n, \dots and non-logical constants c_1, \dots, c_n, \dots of type α . The full set of terms is built

recursively from variables, constants, and syncategorematic expressions such as brackets and λ .² Where α, β are any types and $A : \alpha$ is read “ A is a term of type α ,”

(1.21) **Definition.** The terms of Ty_2 are:

- | | | |
|----|--|----------------------------|
| a. | every variable $x_n : \alpha$ and constant $c_n : \alpha$ is a term of type α ; | variables/constants |
| b. | if $B : \langle \alpha, \beta \rangle$ and $A : \alpha$, then $B(A) : \beta$; | applications |
| c. | if $A : \alpha$ and $x : \beta$ is a variable, then $\lambda x[A] : \langle \beta, \alpha \rangle$; | abstractions |
| d. | if $A : t$ and $B : t$, then $\neg A$, $[A \rightarrow B]$, $[A \wedge B]$, $[A \vee B]$ are of type t ; | connectives |
| e. | if $\varphi : t$ and $x : \alpha$ a variable, then $\forall x[\varphi] : t$ and $\exists x[\varphi] : t$; and | quantifiers |
| f. | nothing else is a term. | closure clause |

And this is the syntax of Ty_2 . We can now build any formula of first-order logic, as well as lambda-abstracts of any type. In the class of interpretations we now specify for this language, first-order formulas will have the same meanings they have in classical model theory,³ and a lambda abstract of the form $\lambda x_e[P_{\langle e,t \rangle}(x)]$ will denote the function from things of type e (individuals) to things of type t (truth-values) that returns true for an argument a iff a has the property denoted by P .

First, fix a non-empty set D , thought of as the “domain of discourse” or the set of entities. In a given interpretation, all terms of type e will have as values elements of D . A second non-empty set W is thought of as the set of “possible worlds” or “situations,” and it supplies values for terms of type s . Terms of type t —aka *formulas*—take as values elements of a third set, $\text{Bool} = \{\text{true}, \text{false}\}$. This reflects the Fregean intuition that declarative sentences refer to truth-values. Terms of type $\langle s, t \rangle$ are *propositions*, expressions that become formulas when fed a world-argument. As usual, the proposition expressed by $p : \langle s, t \rangle$ is therefore conceptualized as the set of worlds w in which $p(w) = \text{true}$ (see Pollard 2008 for a modern alternative).

To specify denotation-types for expressions with non-basic types, we define the class of model-theoretic frame structures specified by Gallin (1975) for Ty_2 . A *frame* specifies for any type α the set of objects that are possible denotations of α -type terms. The frames used here are those satisfying the following definition:

(1.22) **Definition.** A family of sets $(M_\alpha)_{\alpha \in \text{type}}$ is a Ty_2 frame iff

² The logical constants \wedge, \neg, \forall , etc. can actually be defined within this system via only variables, λ , and \top or \perp (see Henkin 1963; Pollard 2008), but they are introduced syncategorematically here for simplicity’s sake. The fact that they *can* be defined in this way allows us to skip meta-theoretical definitions of \neg, \wedge , etc. whenever convenient.

³ Many of the details that guarantee this are omitted here for reasons of space, including the Ty_2 axioms imposing classical interpretations on the connectives and quantifiers.

- a. $M_e = D;$ (e -terms denote individuals)
- b. $M_s = W;$ (s -terms denote worlds/situations)
- c. $M_t = \text{Bool};$ and (t -terms denote truth-values)
- d. $M_{\langle\alpha,\beta\rangle} = (M_\beta)^{M_\alpha}$ ($\langle\alpha,\beta\rangle$ -terms denote functions from M_α to M_β)

A *model* is a frame along with a function that assigns to each non-logical constant an element in its appropriate denotational domain. Officially:

(1.23) **Definition.** A Ty_2 model is a pair $\mathcal{M} = \langle (M_\alpha)_{\alpha \in \text{type}}, val \rangle$ whose first component is a Ty_2 frame and whose second component is a function val that assigns to each constant $c : \alpha$ an element of M_α for all $\alpha \in \text{type}$.

Also important is the notion of a variable assignment: a *variable assignment* on a frame is a function that maps each variable of Ty_2 to an object in its type's domain. For example if g is an assignment and $x : \alpha$, then $g(x) \in M_\alpha$.

To specify the *value of a term relative to model \mathcal{M} and assignment g* , we define a conservative extension of $val \cup g$ called “the meaning function,” denoted by $\llbracket \cdot \rrbracket^{\mathcal{M},g}$. Since every Ty_2 term is a constant, a variable, an application, or an abstraction, the following clauses specify the interpretation of any term not containing the logical connectives (but see fn.3).

(1.24) **Definition.** The value of a Ty_2 term relative to \mathcal{M} and g

- a. $\llbracket x : \alpha \rrbracket^{\mathcal{M},g} = g(x)$
- b. $\llbracket c : \alpha \rrbracket^{\mathcal{M},g} = val(c)$
- c. $\llbracket \lambda x_\alpha [A_\beta] \rrbracket^{\mathcal{M},g} = \text{the function } f : M_\alpha \rightarrow M_\beta \text{ s.t. for any } o \in M_\alpha, f(o) = \llbracket A \rrbracket^{\mathcal{M},g[x/o]}$
- d. $\llbracket A_{\langle\alpha,\beta\rangle} (B_\alpha) \rrbracket^{\mathcal{M},g} = \llbracket A \rrbracket^{\mathcal{M},g}(\llbracket B \rrbracket^{\mathcal{M},g})$

Since truth-values are literally part of the model theory (being elements of M_t), the definition of truth is simple: a closed formula $\varphi : t$ is true relative to \mathcal{M} iff $\llbracket \varphi \rrbracket^{\mathcal{M},g} = \text{true}$ for all g . The usual logical constants are understood classically, so that $\neg\varphi$ is true iff φ is false; $[\varphi \wedge \psi]$ is true iff both φ and ψ are; $\forall x_\alpha [\varphi]$ is true iff for every $o \in M_\alpha$ and every assignment g , $\varphi[g(x)/o]$ is true; and so on.

To make formulas more readable, we frequently and implicitly use the conversion rule β -reduction. Conversion rules play an important role in allowing formulas to be parsable by the human eye/mind. β -

reduction reflects the fact that in the lambda-calculus, applications are interpreted as function-argument applications. Roughly, β -reduction can apply only when A contains no free occurrences of y .

(β) $\lambda x_\alpha[A](y_\alpha)$ is equivalent to $A[x/y]$.

To analyze a fragment of a natural language like English, we can introduce lexical entries using the notation $\llbracket \text{word} \rrbracket = [\text{Ty}_2 \text{ term}]$, where $\llbracket \cdot \rrbracket$ with no superscripts is metalanguage notation for “the Ty_2 term that represents the meaning of natural language expression \cdot .” The terms used to state hypotheses about word meanings will often involve novel constants like **left** : $\langle e, t \rangle$ or **john** : e . Since these constants are not part of the language defined above, we must specify exactly how their interpretations should be constrained. For example, if we want to talk about height formally, we could introduce a constant **tall** : $\langle e, t \rangle$, and henceforth restrict attention to only those models \mathcal{M} in which **tall** denotes the function $f : D \rightarrow \text{Bool}$ that maps all and only the tall elements of D to true. Using the aforementioned notation, we can then write a lexical entry associating the English adjective *tall* with this new constant: $\llbracket \text{tall} \rrbracket = \lambda x[\text{tall}(x)]$. These conventions will frequently be made use of. In most cases novel constants’ intended meanings will be perfectly clear; e.g. **dog** should be true of dogs, $\lambda x[\lambda y[\text{loves}(x)(y)]]$ holds of (a, b) when b loves a , etc. When a constant’s meaning is not clear, I will specify what its intended meaning is, usually informally.

1.3 Dissertation overview and main chapter summaries

1.3.1 Overview of the dissertation in 289 words

This dissertation motivates, defines, and applies a simple but general theory of nominal modification for natural languages. Chapter 2 presents two sub-theories that jointly constitute the bulk of the proposal: the first specifies a morpho-semantics for nouns and inflectional features, according to which lexical count nouns denote predicates of kinds, and the composition of inflectional features with a noun converts the noun’s kind-based meaning into a predicate of individuals. The second sub-theory specifies the syntax of nominal modification, which can have one of two forms: modifiers are adjoined to uninflected kind-denoting nominal heads (*head-adjunction*); or they are adjoined to inflected, individual-denoting noun phrases (*phrasal adjunction*). Directionality parameters are proposed for each kind of adjunction; e.g. in English, head-adjunction is to the left only while phrasal adjunction is bi-directional.

The theory of Chapter 2 is applied to a selection of modification-related phenomena in the remainder of the dissertation, focusing on Bolinger contrasts and comparing results to those of two-domains based theo-

ries of modification. Through detailed investigation of the restrictive/non-restrictive opposition in nominal modifiers (Chapter 3), the direct/implicit relative opposition in certain modal attributive adjectives (Chapter 4), and a selection of additional modification-related puzzles (Chapter 5), I aim to establish the theory in Chapter 2 as a concrete, principled, and cross-linguistically adaptable framework for analysis of the lexical- and morpho-semantics of nouns, and of noun phrase-internal composition. By accounting for the distribution of (some) Bolinger contrasts with a more conservative set of assumptions, specific grammatical enrichments and silent elements postulated by two-domains theories are rendered largely unnecessary. Parts of Chapters 3 and 5 argue that discourse coherence relations play an essential role in certain quasi-grammatical phenomena related to modification, suggesting that modification cannot be fully understood independently of discourse-level pragmatic processes.

1.3.2 Chapter 2: Motivating and building a theory of nominal modification

Chapter 2 sets the stage for the remainder of the dissertation by motivating and defining a new theory of the syntax-semantics mapping in simple noun phrases. §2.1 and §2.2 argue that the traditional analysis of noun meanings as predicates of individuals cannot be maintained in full generality. The motivation for seeking an alternative comes from a variety of sources: I argue in §2.1 that the inflectional properties of nouns, widely considered to be “uninterpretable” features, should be reflected in the denotations of nouns. It follows from this that the introduction of inflectional morphology on a noun has a non-trivial effect on semantics.

§2.2 summarizes a body of previous research on the semantics of nouns across languages, and on the relationship between kind- versus individual-reference on the one hand and the syntactic projections of N versus D on the other. The most important idea emerging from this discussion is that reference to kinds or types is encoded within the lower NP region of the nominal, while reference to individuals or tokens is a function of the higher DP region of the nominal. This correlation is shown in §2.3 to correspond exactly to an asymmetry in the interpretation of nominal modifiers, which has been informally described by linguists since at least Bolinger (1967) in the following terms: generic, characterizing, or permanent properties tend to be expressed by modifiers that are structurally close to head nouns, while more distant modifiers tend to express accidental, non-essential, or temporary properties. This parallel motivates the following analysis, which constitutes the foundation for the analyses throughout the dissertation: inner modifiers compose with nouns within the kind-denoting structural level and as a result have the first set of interpretive properties (generic, characterizing); and outer modifiers compose with nouns in the higher, token-denoting level and

as a result have the second set of interpretive properties (accidental, temporary). §2.4-2.5 formalize this parallel by encoding the type-to-token conversion in the semantics of inflectional morphology, and defining a syntax that allows (some) modifiers to compose with both inflected and uninflected nouns.

§2.4 introduces a theory of noun phrase structure and meaning. In a sentence, it says that lexical count nouns denote predicates of kinds, and that (for languages that have it) the composition of inflectional features with a noun converts the noun's kind-based meaning into a predicate of individuals, the extension of which reflects whether that noun is singular or plural, masculine or feminine, etc. I show that this analysis accounts for a range of generalizations reported in the literature about cross-linguistic differences in kind- and individual-reference. The generalizations concern the grammaticality and interpretation of bare nominals and simple definite descriptions in English, German, Italian, Spanish, and French.

§2.5 extends the theory of §2.4 to NPs and DPs containing modifiers. Non-appositive modifiers are integrated into nominals in one of two ways, a dichotomy partially inspired by Sadler & Arnold (1994): via adjunction to the nominal head prior to the introduction of inflectional features (*head-adjunction*); or via adjunction to the inflected noun, whose denotation has already been converted into a predicate of individuals (*phrasal adjunction*). The modification sub-theory also contains directionality parameters for each kind of adjunction: in English, head-adjunction is to the left only, while phrasal adjunction is bi-directional. But in Italian, I argue that it is head-adjunction that is bi-directional, while phrasal adjunction is to the *right* only. This approach yields a constrained typology of languages according to how nominal modification is syntactically encoded, similar to but more general than Cinque's (2010) dual-source hypothesis.

Finally, in the remainder of the chapter, I show how the mirror-image distribution of modifier interpretation in English versus Italian is treated within this framework. Using the stage-level/individual-level Bolinger contrast as a testing ground, I show that the proposals of §2.4-2.5 jointly derive the observed distribution of readings in English, as well as ordering restrictions between prenominal stage- and individual-level adjectives.

1.3.3 Chapter 3: Restrictive versus non-restrictive modification

Chapter 3 investigates the phenomenon of non-restrictive adjectival modification, focusing on connections to the two-domains theory of modification and to the field of discourse semantics/pragmatics. In the first part of the chapter, §3.1, I present and assess five existing conceptions within linguistic theory of what "non-restrictive modification" is, arguing that all fail to match intuitions for at least some body of critical

examples. Syntactic characterizations apply only to relative clauses, and hence are ill-suited to categorize attributive adjectives. Furthermore, the semantic property of failing to restrict a noun denotation is not a sufficient condition for a nominal modifier to be considered non-restrictive in the intuitive sense. These observations demonstrate that non-restrictiveness is neither an exclusively syntactic nor an exclusively semantic property. §3.1 concludes with a collection of examples strongly suggesting that non-restrictive modifiers are required to provide “relevant” information in a way that restrictives need not; this generalization supports the hypothesis that a modifier in a felicitous utterance is non-restrictive only if the modifier stands in a certain kind of pragmatic relationship to some other proposition in the same utterance or conversation. In other words, the correct characterization of non-restrictiveness must relate to discourse-level factors in addition to properly semantic factors.

§3.2 proposes new formal definitions for *restrictive* and *non-restrictive* which have both semantic and pragmatic components. The pragmatic component makes reference to discourse coherence relations (Hobbs 1985; Kehler 2002; Asher & Lascarides 2003) as well as to the distinction between active and passive conversational contexts (Kripke 2009). In particular, I introduce the central hypothesis that a modifier is non-restrictive (and felicitous) if and only if two conditions hold: (i) that the speaker believes the modified noun phrase to have the same denotation as the unmodified noun; and (ii) that the modifier can be used to deduce an implication and rhetorically relate it to some salient proposition in the active discourse context. This proposal explains the observation from §3.1.6 that in contrast to restrictive modifiers, non-restrictives often feel deviant if their content is not in some sense “relevant.” Restrictive modification is defined as the denial of property (i), and thus may or may not contribute to coherence. These definitions are meant to characterize in a precise but general way the intuitions behind working linguists’ usage of the terms.

Throughout §3.2, I use terminology and tools from a framework for discourse semantics—specifically a hybrid extension of Asher & Lascarides’ and Kehler’s. Formal theories of discourse relations are still relatively obscure within model-theoretic semantics, and therefore I devote §3.2.1-3.2.2 to general discussion of how discourse-level phenomena can and should be analyzed within the parameters of modern semantic theory. In §3.2.3 I introduce a simple, stripped-down framework for discourse semantics inspired by Asher & Lascarides’s (2003) architecture. The formalism is designed to interface in a straightforward way with standard truth-conditional semantics, but is flexible enough to be adapted to mostly any static or dynamic semantic framework (though the system itself contains a context-update procedure).

§3.3 returns to the topic of DP structure, and asks whether the restrictive/non-restrictive Bolinger con-

trast (R/NR) should indeed be viewed as such in light of the considerations of §3.2. In §3.3.1 I show with empirical and conceptual arguments that R/NR differs fundamentally from other Bolinger contrasts, despite its superficially similar appearance. In §3.3.2 I scrutinize and demonstrate some difficulties for Larson’s (1998) and Cinque’s (2010) two-domains based theories of R/NR, irrespective of how (*non-*)*restrictive* is understood. §3.4 then proceeds to apply the theory of NP/DP from Chapter 2 and the characterizations from §3.2 to the syntax, semantics, and pragmatics of R/NR. The main conclusion of the new analysis is that R/NR only superficially has the distribution of Bolinger contrasts as described by Cinque (2010), and that restrictive/non-restrictive does not track the inner/outer distinction as related contrasts do. The analysis as implemented here depends upon two additional assumptions, though in principle alternatives could be sought: (i) the projection of DP shells in the syntax of modified nominals, following Larson (1991); and (ii) a novel type adjustment principle, which encodes side-issue entailments—including the contribution of non-restrictive adjectives and appositives—as definedness conditions (like the partial-function analysis of presupposition). The mechanism in (ii) composes a noun (phrase) and a modifier to result in an expression that functions syntactically like a noun (phrase) but is defined only if the modifier property holds of the referent of the noun (phrase). As in Chapter 2, variation across binary syntactic parameters is argued to underly different empirical patterns across languages (here, for example, the availability of non-restrictive readings of adjectives in English versus Italian). §5 considers some connections of the theory developed here to related issues, including psycholinguistic studies of modification. In particular, eye-tracking researchers working in the visual world paradigm study a notion of “contrast” in the interpretation of attributive adjectives that closely corresponds to the theoretical linguist’s notion of “restriction” as conceived of here. I discuss some possible ways that the proposals of this chapter could be evaluated using experimental and behavioral methodologies, as well as a couple of theoretical loose ends.

1.3.4 Chapter 4: Direct versus implicit relative modal adjectives

Chapter 4 is a second case study in adjective ambiguity, which targets the “direct/implicit relative” alternation in modal adjectives illustrated in 1.1 above. This contrast is of particular interest because the availability of the implicit relative reading of e.g. *possible* depends not only upon syntactic position, but also upon the lexical semantics of the DP’s determiner: for example *every* licenses implicit relative readings while *some* does not. §4.1 provides an overview of existing research on this contrast, emphasizing arguments to the effect that the distribution of *possible*’s readings provides strong evidence in favor of the two-domains model.

§4.2 argues in direct opposition to two-domains theories of this contrast: whether *possible* occupies an inner or an outer modifier position cannot be predicted on the basis of whether it has a direct or indirect/implicit relative reading. Therefore, the postulation of two modification domains is mostly orthogonal to the question of how *possible* receives its interpretation. Following Larson's (2000a) analysis of the implicit relative reading, though, I do assume a silent clausal complement of *possible* on its indirect reading, and hence that *possible* can be a covert/reduced relative clause under certain conditions. Despite this syntactic variability, though, it remains possible and viable to analyze *possible* and related adjectives as lexically unambiguous but semantically polymorphic: they have schematic types and can compose with propositional or nominal arguments.

With this analysis in hand, §4.3 investigates the puzzling determiner restriction on implicit relative adjectives. While it remains somewhat mysterious, I point out that the very same restriction is active in amount/degree relative clauses, as has been known since Carlson (1977a). Key to establishing this point is the observation that the adjective *necessary* can also admit of implicit relative readings. The parallel with amount relatives is then used as motivation for a degree-based analysis of all implicit relative adjectives, a generalization of Romero's (2013) proposal for modal superlatives. While I implement the syntax and semantics of the silent degree relative in a slightly different way, this analysis complements Romero's. §4.4 presents derivations of the direct and indirect readings in superlatives, universals, and plain definites—a kind of DP that had not been recognized to admit of implicit relative adjectives until the present work. The section and chapter close by revisiting Larson's (2000a) and Cinque's (2010) generalizations about the distribution of direct and implicit relative readings. Point-by-point, I show how the analysis of modal adjectives in this chapter, combined with the theory of DP from Chapter 2, accounts directly for the basic pattern in English. Related issues and observations are identified intermittently throughout the chapter.

1.3.5 Chapter 5: Modified names, subtriggering, and manner implicatures

Chapter 5 applies the theory of DP structure advanced in preceding chapters to three additional issues in the syntax, semantics, and pragmatics of modification: the composition of modified proper names and related constructions (§5.1); the licensing of *any* in positive contexts by modification (subtriggering, §5.2); and an asymmetry in the implicatures of restrictive versus non-restrictive modifiers in downward entailing environments (§5.3).

§5.1 revisits the type mismatch repair mechanism utilized in Chapter 3's analysis of non-restrictive

modification (TMAP). While the principle has somewhat of an ad hoc nature, it affords natural analyses of article-less modified names (*poor John*), and potentially also polydefinite nominals in Greek, expressive adjectives, and attributive quantity modifiers. Any compositional theory of side-issue meaning necessarily invokes a mechanism whose function is equivalent to the one proposed here (e.g. Potts 2005; Morzycki 2008). Thus, as the applications of TMAP extend beyond those of competing mechanisms, I conclude that TMAP is a reasonable operation to assume in a formal theory of compositional interpretation. Following the general discussion of composition, I present a theory of modified and unmodified names accompanied by articles (*an/the exhausted John Smith; der Hans*), proposing an analysis according to which names (can) denote predicates of spatio-temporal individual stages. This analysis treats articles and modifiers in complex names as no different than their occurrences in ordinary DPs, an improvement over competing theories that maintain referential analyses of names and complicate the treatment of all other elements in a modified name. The names-as-stage-predicates analysis fits nicely with the theory of noun meaning from Chapter 2 in that names can denote individuals like nouns can denote kinds, but names can also denote predicates over stages of an individual like nouns can denote predicates over subkinds of a kind. In other words, individuals are to kinds what stage-predicates are to subkinds; the difference between names and common nouns is thus simply that names denote in the individual domain while nouns denote in the kind domain.

§5.2 summarizes some interesting generalizations about the occurrence of *any* in non-negative, non-modal contexts—also known as subtriggering. Interestingly, *any* can be licensed in positive contexts only by a proper subset of nominal modifiers, characterized by Dayal (1998) as relative clauses with *essential* but not accidental readings. The points I make about subtriggering are relatively minor: first, a counterexample with implicit relative *possible* shows that it is not only postnominal modifiers that can subtrigger *any*, contrary to Carlson's (1981) decades-old generalization. Second, Dayal's (1998) essential-accidental analysis can be reinterpreted in terms of rhetorical connections: a relative clause subtrigs *any* only if it establishes a discourse coherence relation with the content of the sentence in which *any* occurs. If correct, this idea provides further evidence that coherence relations can interact with the grammar of a language, as suggested for non-restrictive modification in Chapter 3.

Finally, §5.3 reflects upon an asymmetry in the implicatures of restrictive versus non-restrictive modifiers: as noted by Katzir (2007) and others, relative clauses and other restrictive modifiers trigger manner/quantity implicatures in downward-entailing contexts. Building on discussion in Leffel 2011, I discuss the fact that no such implicatures are observed for non-restrictive adjectives and argue that this fact motivates

a form of Magri's (2011) blindness hypothesis for manner implicatures. Without reaching firm conclusions, I entertain various explanations for the facts described in this section. Potentially the most interesting hypothesis also explains why postnominal modifiers in English do not admit of non-restrictive readings as easily as do prenominal modifiers: if postnominal modifiers are reduced relative clauses, and if clauses are prosodic domains, then default prominence is assigned to postnominal modifiers via a nuclear stress principle; as a result the modifier is focus-marked and thus evokes a set of alternatives whose existence is incompatible with a non-restrictive construal. In contrast, there is in most cases little or no motivation for a reduced relative analysis of prenominal attributive adjectives, and as a result they need not be focus-marked, need not evoke alternatives, and need not be interpreted restrictively.

In sum, I hope to establish in this dissertation that intricate distributional generalizations and structure-sensitive ambiguities in modified nouns do not necessarily imply a highly articulated nominal structure, nor do they imply a general but idiosyncratic ambiguity in adjectives. Instead, by limiting the kinds of grammatical enrichments we allow into our theoretical toolbox, such patterns can be used to inform our understanding of lexical noun meaning, nominal morphosyntax and morpho-semantics, the mechanics of composition, and even natural language ontology.

Chapter 2

The internal syntax and semantics of noun phrases

2.0 Introduction

The investigation of Bolinger contrasts requires a theoretically informed and cross-linguistically applicable theory of nouns and of DP-internal composition. After motivating the need for such a theory, I establish a syntactic parallel between kind- or type-denoting nominals on the one hand, and generically interpreted nominal modifiers on the other. By unifying the observations of Larson (1998,1999) and Vergnaud & Zubizarreta (1992), it becomes evident that certain aspects of a modifier's interpretation depend critically upon semantic properties of the noun it modifies, as well as the structural level at which the modifier is introduced.

The overall goal of the chapter is to develop a concrete set of principles governing the structure and interpretation of noun phrases that is both informed by recent advances in the study of DP, and which provides a principled explanation for the aforementioned parallel between generic meaning in nouns and in modifiers. Capturing this parallel within a formal theory of English grammar, I will argue, affords a natural analysis of the stage-level versus individual-level contrast discussed in Chapter 1. Specifically, I propose in §5 that the i-level/s-level Bolinger contrast can be explained as an interaction between attachment height and the conversion of a noun's type-based denotation into a token-based denotation.

The specific set of principles and assumptions about DP that I develop integrates observations and ideas from a range of recent theories about nouns and modifiers. The resulting theory covers a variety of related issues including the semantics of number- and gender-marking, the cross-linguistic inventory of (hypothesized) expletive and null determiners, and how parametric variation in NP and DP morphosyntax influences the nature of modification structures. The point of building such a theory is to have a precise, concrete, and internally consistent framework in which to formulate analyses of other Bolinger contrasts—the main

objective of the remainder of this dissertation.

Chapter outline: §1 presents some arguments for why studying modification (especially Bolinger contrasts) requires a well-motivated and cross-linguistically applicable theory of nouns and noun phrases. §2 summarizes a range of recent observations/theories/proposals about nouns, determiners, NPs, and DPs. §2 establishes a heretofore unobserved parallel between Vergnaud & Zubizarreta's (1992) NP/DP-type/token correspondence and Larson's two-domains theory of nominal modification; sketches the kind of DP architecture required by these converging observations. §4 stitches together a theory of syntactic and semantic composition of nouns, determiners, inflectional morphology, NPs, and DPs; draws together insights from each of the proposals introduced in §2. §4 also applies the theory to a contrastive paradigm between English/German and French/Italian/Spanish simple noun phrases. Finally, §5 applies the newly developed theory to the individual-level versus stage-level contrast in modifier interpretation, and compares the analysis to alternatives. Positive results and shortcomings are discussed throughout.

2.1 Motivation for a fine-grained theory of nominal reference

The interpretation of a modified noun is a function not only of the modifier's semantics, but also of the noun's semantics, including whatever might be contributed by inflectional or derivational morphology. Perhaps because modifiers are (quasi-definitionally) non-obligatory elements, research on DP often assumes a relatively simplistic syntax and semantics for modifiers. At the same time, research focused on modification often makes simplifying assumptions about the semantics of nouns, e.g. that they uniformly have type $\langle e, t \rangle$. The latter kind of simplifying assumption, though, is not as innocuous as the former: every noun modifier is contained in a noun phrase, whereas not every noun phrase contains a modifier. Thus, while many aspects of DP syntax and semantics can be examined independently of adjectives or relative clauses, the theory of nouns adopted will necessarily affect the predictions of a theory of modification. It is therefore crucial for the present study that we be equipped with a well-motivated perspective on the semantics and syntax of simple, unmodified noun phrases (like *the dog* or *dogs*).

To see why it is important to have a well-grounded theory of nouns when studying modification, consider the following particularly well-known example: Romance common nouns are inflected for gender and number, while English common nouns, only for number. This difference raises the question whether gender marking has any semantic effect. In other words, does an inflected singular count noun in, say, French have a different semantics than an inflected singular count noun in English? There is reason to believe so. Consider

the French noun *chien* ‘dog’. The lexical gender of *chien*—masculine—is morphologically visible on the definite article when *chien* heads an argument nominal. Ungrammaticality results if the determiner does not match the noun in gender.

- (2.1) a. le chien
the.MASC dog
- b. *la chien
the.FEM dog

However, in reference to a female dog, *chien* can also appear in the feminine form *chienne* (as in 2.2b). In this case, the feminine article *la* is used.

- (2.2) a. *le chienne
the.MASC dog.FEM
- b. la chienne
the.FEM dog.FEM

Such contrasts do not occur with nouns that denote inanimate objects; e.g. there is no *la maison/*le maison(ne)* ‘the house’ alternation comparable to *le chien/la chienne*.

Put another way, if gender marking has a non-vacuous effect on semantics, then we might expect the lexical semantics of a noun in French to look different than the lexical semantics of its English translation. Which is just to say that language-specific considerations necessarily come into play in even the most simple of semantic domains: should the extension of *chien* be the same as that of *dog*, or that of *male dog*? What about *chienne*? The answers are not self-evident.

The idea that common nouns uniformly denote sets of individuals (and that that’s the end of the story), while elegant in its simplicity, encounters a wide variety of empirical problems, some more obvious than others. One obvious example is the singular/plural distinction. If number is purely morphosyntactic, then what accounts for the clear difference in meaning between *the dog* and *the dogs*? A wide range of well-motivated theoretical proposals have been advanced about the semantics of number, in particular the interpretation of plural versus singular nouns (e.g. Link 1983; Schwarzschild 1996; Krifka 2003; Dayal 2004; Zweig 2008).

There are many other problems with the simple view, such as the count/mass distinction, for instance. What kind of entity should count as satisfying the predicate *water*? Is it individual molecules, or large seas? Both? This kind of consideration has led some to posit distinct ontological domains for count- versus mass-noun reference (Link (1983); Chierchia (1998); a.o.).

A further problem has to do with subcategorization. Certain nouns appear to have a semantic argument position for complements, while other nouns do not. Compare: *picture of John/*pencil of John, decision*

to leave/**conclusion to leave*. Interestingly, such alternations in deverbal nouns sometimes appear to be determined by the selectional restrictions of the verbal root—as in *decided to leave*/**concluded to leave*. A natural conclusion to draw is thus that the semantics of the verbal root from which a noun is derived can affect a noun’s semantic argument structure (and hence its semantic type).

And what about proper names? It is common to assume that proper names have a semantic type distinct from that of common nouns (since at least Montague 1974). If true, then a strict Montogovian category-type mapping cannot be maintained.

The few issues just mentioned, along with many others, suggest that nouns are semantically a heterogeneous category. Furthermore, the semantic type and denotation of a noun can apparently depend upon its inflectional or derivational morphology, the kind of entity it refers to, the semantics of its root, historical accidents, and many other factors.

Syntactic and semantic heterogeneity are also observed in the domain of noun phrases. A common assumption in contemporary semantics is that names and definite descriptions have semantic type e , while quantified noun phrases have type $\langle\langle e, t \rangle, t\rangle$ (see Heim & Kratzer 1998, a departure from the classical type-category correspondence). The semantic type of a DP can also vary depending upon its grammatical function, e.g. predicate versus argument indefinites (Partee 1987; Partee & Rooth 1983; Chierchia 1998). Some have even argued that definite descriptions in argument positions are type-ambiguous between e and $\langle e, t \rangle$ (Mikkelsen (2002)). In the following section, I will discuss in detail a few analyses of noun and DP syntax/semantics that will be of particular importance in stating my own assumptions about nouns in §5, which will form the starting point for subsequent analyses in this dissertation.

2.2 Some recent advances in DP syntax and semantics

A variety of alternatives to the nouns-as-sets theory have been advanced in recent decades. What many of these theories have in common is that they all relate—or are shown to relate—to Bolinger contrasts, whether the relations are immediately obvious on the surface or not. In the following expositions, I survey some theoretical proposals parts of which are adopted in §3-4’s analysis.

2.2.1 Kind- and individual-reference in common noun denotations

In many languages, a singular count noun can have a so-called kind-reading (*the dog descends from the wolf*) in addition to its ordinary predicate reading (see Carlson 1977b for foundational discussion). On the kind interpretation, a noun such as *dog* refers to the species *Canus Lupus*, or to the reified concept of ‘dogness’.

For example both *dog* and *wolf* have a kind reading in the sentence *The dog descends from the wolf*. Krifka et al. (1995) identify a third interpretation, the “taxonomic” or “subkind” (or “taxonomic subkind”) reading. On this reading, *whale*, for example, acts “as a predicate applying to the *subkind*[s] of the kind *Cetacea* [Latin for *whale*], that is, the blue whale, the sperm whale, the dolphin [*sic*], etc.” (Krifka et al. 1995:74) As illustrated in 2.3, this reading is available for nominals of various sorts, indicative of genuine—and general—polysemy in the lexical meaning of nouns.

- (2.3) a. The terrier is a dog.
 b. This dog—the terrier—makes a great pet.
 c. Every dog—even the terrier—descended from wolves.

To talk about natural language reference to kinds and sub-kinds in a formal language, Krifka et al. (1995) introduce a binary “taxonomic sub-kind” relation **T**, which holds between kinds *k* and *m* iff *m* is a more specific kind than *k* (e.g. *k* is “dog” and *m* is “terrier”). Summarizing, a common noun in a DP can be interpreted in one of three ways (we put off until §5 the issue of how these three readings arise): as a predicate of individuals, as a kind, and as a predicate of subkinds.

As evidence that the sub-kind and predicate readings are truly distinct, Krifka et al. (1995) present the following minimal pair from Chinese, which uses classifiers to categorize nouns. The form of the classifier determines whether the noun *xiong* ‘bear’ is interpreted as an individual- or a sub-kind-predicate.

- | | | | |
|-------|---|-------|--|
| (2.4) | yi zhī xióng one CL bear ‘an individual bear’ | (2.5) | yi zhong xióng one CL bear ‘a kind of bear/a bear species’ |
|-------|---|-------|--|

Dayal (2004) notes that Chinese nouns can have a kind-referring reading as well:

- (2.6) Gou juezhong le
 dog extinct ASP
 ‘Dogs are extinct.’

The conclusion of this brief discussion is that there is a systematic polysemy (*not* ambiguity) in noun denotations, where the three possible interpretations are as indicated above.

2.2.2 Event- and kind-relativity in adjectives and nouns

Larson (1998) argued that nouns have a hidden eventuality variable in their semantic argument structure. This conclusion is motivated by his analysis of ambiguities such as the ones below, primarily 2.7. These ambiguities—paradigm cases of Bolinger contrasts—have remained largely unexplained since at least Vendler (1968).

- | | | | |
|-------|--|-------|--|
| (2.7) | Olga is a beautiful dancer. | (2.8) | a. a quick cup of coffee |
| | <i>Intersective</i> : “beautiful and a dancer” | | (the cup or the drinking can be quick) |
| | <i>Manner</i> : “dances beautifully” | | b. my old friend |
| | | | (my friend or our friendship can be old) |

Larson argued that the source of these ambiguities is nominal in nature. Specifically, he argued that the semantic representation of certain common nouns contains a Davidsonian event argument in addition to its individual argument, and that adjectives like *beautiful* are lexically polymorphic, denoting either properties of individuals or of events. This allows an adjective to be semantically linked with a noun via identification of two different variables, and for Larson this explains the ambiguity in 2.7 (Q is Larson’s notation for a quantifier of unspecified force).

- | | | |
|-------|---|-------------------------------|
| (2.9) | a. <i>Intersective</i> : $Qe[\mathbf{dancing}(\mathbf{olga})(e) \dots \mathbf{beautiful}_C(\mathbf{olga})]$ | (Olga is beautiful) |
| | b. <i>Manner</i> : $Qe[\mathbf{dancing}(\mathbf{olga})(e) \dots \mathbf{beautiful}_C(e)]$ | (Olga’s dancing is beautiful) |

An advantage of such an approach is that it allows one to maintain a uniform, intersective analysis of adjectival modification even in cases that appear to involve non-intersective interpretations: both readings are intersective, the only distinction being in whether the denotation of the noun is intersected with the denotation of (the individual-predicate version of) the adjective, or whether the event-predicate hidden inside the noun is intersected with (the event-predicate version of) the adjective.

This kind of reasoning—that is, analyzing apparently non-intersective adjectives as intersective along an abstract dimension—has been extended by McNally & Boleda (2004) to the interpretation of “relational” or “classificatory” adjectives in constructions such as the following:

- | | |
|--------|---|
| (2.10) | a. John is a technical architect. |
| | b. El Martí és arquitecte tècnic. (Catalan) |
| | the M. is architect technical |

- c. C'est chaleur solaire. (French)
it's heat solar

Relational modifiers have certain distributional properties of intersective adjectives, the most crucial of which is that they appear only postnominally in Romance languages (McNally & Boleda 2004). But relational adjectives are not intersective in the traditional sense: from (2.10a) it does not follow that John is technical (in fact *John is technical* even feels mildly ungrammatical). To reconcile these observations, McNally & Boleda (2004) propose that relational adjectives are intersective in an abstract sense: they denote properties of “kinds” (parallel to *widespread* or *extinct*), and semantically combine with nouns via a composition mode that identifies the kind variable in the representation of the adjective with that of the noun. To ensure that the combinatorics of the analysis work properly, McNally & Boleda propose that all nouns come equipped with an extra, kind argument, and have the following general meaning (which is based upon the idea of the taxonomic sub-kind reading of nouns discussed in §3.2 above, identified by Krifka et al. 1995):

$$(2.11) \quad \text{For all nouns } N, \llbracket N \rrbracket = \lambda x_k [\lambda y_o [R(y)(x) \wedge \mathbf{N}(x)]]$$

And parallel to Larson’s event analysis of adverbial adjectives like *beautiful* in 2.7, they claim that relational adjectives are kind predicates that combine with nouns via the composition rule specified in 2.12—similar to standard approaches to intersective semantics for modification.

(2.12) Let N be a noun and AP a relational adjective phrase. Then

- a. $\llbracket N \rrbracket = \lambda x_k [\lambda y_o [R(y)(x) \wedge \mathbf{N}(x)]]$
- b. $\llbracket AP \rrbracket = \lambda x_k [\mathbf{AP}(x_k)]$
- c. $\llbracket N[AP] \rrbracket = \lambda x_k [\lambda y_o [R(y_o)(x_k) \wedge \mathbf{N}(x_k) \wedge \mathbf{AP}(x_k)]]$

Notice that this composition does not saturate the noun’s kind argument, so that *arciteqtue tècnic* ‘technical architect’ is still of type $\langle k, \langle e, t \rangle \rangle$. McNally & Boleda remedy this situation by proposing that “the kind argument gets saturated by a contextually-determined kind” (189). It follows that a simple sentence with a relational adjective modifying a noun works out to the following formula:

- (2.13) a. El Martí és arquitecte tècnic.
b. $[R(\mathbf{marti})(k) \wedge \mathbf{architect}(k) \wedge \mathbf{technical}(k)]$

The analysis guarantees that *technical* is not the kind of property that, say, you or I could have: it is a

property that holds of kinds. Furthermore, it predicts that relational adjectives should pattern syntactically like intersective adjectives and unlike “operator adjectives” such as *former* or *alleged*. The prediction is partially confirmed by the fact that like intersectives, relational adjectives appear obligatorily in postnominal position in French and Catalan (Catalan, McNally & Boleda 2004:181):

- (2.14) a. un presumepte assassi / *un assassi presumepte
 a presumed assassin (operator adjective)
- b. una malaltia pulmonar
 a disease pulmonary (relational adjective)
- c. un escriptor jove
 a writer young (intersective adjective)

The underlying idea in Larson’s (1998) and McNally & Boleda’s (2004) analyses is that the lexical semantics of the noun can just as easily affect the interpretation of a modified noun as can the lexical semantics of the adjective. This recognition has potentially widespread consequences. For instance, if noun semantics is a parameter of cross-linguistic variation (as argued by Chierchia (1998), Vergnaud & Zubizarreta (1992)), then cross-linguistic differences in the semantics of modified nouns should be expected.

2.2.3 The semantics of number and the unmarked status of PL

Assume that language makes reference to both atomic and plural individuals, also called atoms and (mereological) sums, respectively (see Link 1983; Schwarzschild 1996; Barker 1992 for discussion). As noted by Link (1983), some predicates appear to be capable of applying only to atomic individuals, others only to plural individuals, and still others to either; thus the entailment in (2.15c) goes through on one interpretation.

- (2.15) a. John and Mary died. \Rightarrow John died and Mary died.
- b. John and Mary are a good team. $\not\Rightarrow$ John is a good team and Mary is a good team.
- c. John and Mary lifted a piano. \Rightarrow ? John lifted a piano and Mary lifted a piano.

Considerations like this led Link to introduce a semantic pluralization operator $*$, which takes a set and returns the set of all sums created from the atoms of that set. \oplus is mereological sum formation.

- (2.16) a. $\llbracket P \rrbracket = \{a, b, c, \dots\}$
- b. $\llbracket *P \rrbracket = \{a, b, c, a \oplus b, b \oplus c, a \oplus c, a \oplus b \oplus c, \dots\}$

Application of $*$ to a verb phrase, for instance, can convert the VP’s extension from a set of individuals to the closure of that set under sum formation. And since VPs and Ns often have a common semantic type, the operators in 2.16 could just as well be applied to noun denotations. In fact, traditional accounts of the semantics of plurality (e.g. Link (1983); Schwarzschild (1996)) hold that (setting aside intensionality) singular nouns denote sets of atomic individuals, while plural nouns denote sets of sum-individuals (in some accounts, in addition to atoms).

- (2.17) a. $\llbracket \text{dog} \rrbracket = \{ \text{sparky}, \text{fido}, \text{rags}, \dots \}$
 b. $\llbracket \text{dogs} \rrbracket = \{ \text{sparky}, \text{fido}, \text{rags}, \text{sparky} \oplus \text{fido}, \text{sparky} \oplus \text{fido} \oplus \text{rags}, \dots \}$

This perspective allows one to factor out the semantic contribution of plural morphology: the meaning of $-s$ is just $*$. Many slight variants of this analysis have been proposed. For example, Tonhauser & Kiparsky (2011) argue for an “exclusive” semantics for PL. On this view, the contribution of PL is \otimes instead of $*$, where \otimes is the same as $*$ except that $\llbracket \otimes P \rrbracket$ contains no atoms.


An important feature of any plural-as-sum-formation analysis is that the singular is taken to be the primitive or unmarked form of a noun, and the plural is compositionally constructed from the singular. A recent alternative to the $\text{PL} = *$ analysis, which I will call the “number as agreement” approach (Sauerland 2003; Sauerland et al. 2005), holds that it is the plural that is semantically basic/unmarked, and that singular morphology actually imposes selectional restrictions on the noun. More specifically, Sauerland (2003) and Sauerland et al. (2005) argue that number morphology indicates agreement with a higher head hosting the semantically contentful agreement features sg or pl . This is conceptually motivated by languages like German, French, and many others, where both nouns and articles (and even adjectives) are inflected for number (e.g. *das_{sg} Haus/die_{pl} Häuser* ‘the house/s’). On this theory, the semantics of the features SG and PL are encoded presuppositionally, i.e. number features denote partial identity functions that are defined only when their argument’s denotation has an appropriate mereological structure.¹

- (2.18) a. $\llbracket \text{SG} \rrbracket = id_{\{x \in D \mid \text{atom}(x)\}}$ b. $\llbracket \text{PL} \rrbracket = id_D$

Notice that on this approach, since sums still belong to D_e , it is the plural that is semantically unmarked and the singular that introduces a non-trivial condition, namely the presupposition that the DP refers to an atomic individual. While not essential to the agreement approach, Sauerland et al. (2005) argue that

¹ Since only the singular imposes a non-trivial presupposition on this view, Sauerland (2003) elaborates this analysis in terms of Heim’s (1992) pragmatic constraint *Maximize Presupposition!*.

the number presupposition is introduced above DP, by a functional layer ϕ P hosting “phi-features” such as person, number, and gender.



Because of this syntactic assumption, number morphology applies semantically to expressions of type e , meaning that it is DPs that are singular or plural, not Ns or NPs. In §5 below, I will develop a modified version of this analysis that introduces the semantics of number DP-internally.

2.2.4 DP layers and the type/token distinction in nominal reference

It has been proposed that the syntactic N and D layers of the nominal correspond to qualitatively different kinds of semantic reference. In particular, Vergnaud & Zubizarreta (1992) and others have argued that nouns denote types (kinds in contemporary terms) and that quantification over tokens (individuals) in nominals is contributed by D (or by some silent element within the projection of D, depending on the theory). Vergnaud & Zubizarreta (1992) (henceforth VZ) propose the following principle.

(2.20) Correspondence Law

- 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)

Intuitive motivation for this kind of position can be obtained by examining anaphoric relations licensed by different kinds of DPs. Recall that English bare plural subjects can be interpreted generically, as in (2.21a); quantified subjects, by contrast, generally cannot (cf. (2.21b)). However, 2.22 shows that the plural pronoun *they* can refer back either to the plurality of individual bears from (2.21b) (tokens) or to the kind denoted by the noun within the quantified DP (the type).

- (2.21) a. Bears are dangerous.
b. Several bears were seen at the campground.

- (2.22) a. They looked right at us!
 b. They are common around here. (#‘several bears are common’)

In fact, the situation is even more complex than has been noted in the literature. Particularly, *one*-anaphora in subject and object position is licensed by generic bare plurals, and the former is unambiguously individual-denoting. In other words, *one* refers to an individual bear token, but its antecedent refers to a type/kind.

- (2.23) Bears are common around here...
 a. In fact I just saw one.
 b. One just walked up to me and gave me a scare.

More direct evidence for the N/D-type/token correspondence can be found by examining cross-linguistic variation in type- and token-reference. VZ’s original arguments for the correspondence law 2.20 have to do with so-called inalienable possession constructions in French (exemplified in 2.24), which are absent in English. 2.24 is ambiguous between A(lienable) and I(nalienable) readings, whereas the corresponding English sentence—identical to the gloss—has only the A interpretation.

- (2.24) Les enfants ont levé la main.
 the children have raised the hand
Alienable: ‘The children raised the unique hand.’
Inalienable: ‘Each of the children raised his/her hand.’

VZ propose that this difference between English and French is due to syntactico-semantic differences in NP and DP. On their analysis, the reason that French has inalienable possession constructions and English doesn’t is that French has an expletive (pronounced but semantically vacuous) definite article that in certain constructions allows a full DP to refer to a single type but to have a bound interpretation. While they do not offer an explicit semantic formalization, I believe the idea is that in 2.24, the object DP *la main* ‘hand’ simply inherits the property-type semantics of the NP *main*, so that the compositional meaning of 2.24 ends up something like “Each child participated in a hand-raising event.” Because, on their analysis, English lacks a semantically empty definite article, the inalienable reading is absent. This parametric difference will form an important part of the theory developed in Section 5.

Regarding the licensing of expletive articles in French, VZ state that “the complementation relation between D and NP is morphologically licensed, by the overt person, number, and gender agreement relation

that holds between the determiner and its complement NP (as in the case of Agr and VP).” (VZ, 615) This explanation implies that English, whose definite determiner is morphologically invariant, should not have expletive articles.² And this is exactly what VZ propose: “The definite determiner may function as an expletive from the point of view of denotation in French but not in English. In other words, in English the definite determiner must be absent in a type-denoting expression (635-6).” This parametric difference, in turn, is then used to explain a number of further contrasts between French and English. Two are: that French has but English lacks inalienable possession constructions such as 2.24; that in examples like 2.25, French definite plurals are ambiguous in a way that their English counterparts are not: according to VZ, *les baleines* ‘the whales’ in (2.25a) can denote the kind “whales” but also the set of whale-subkinds, whereas English *the whales* has only the latter interpretation.

- (2.25) a. Les baleines sont en train de disparaître.
 the.PL whales are in process of disappear
 b. The whales are becoming extinct.

Zooming out a bit, how are these interpretive asymmetries represented in the grammar? The answer to this question depends in part upon what the primitive elements are in our semantic ontology. Carlson (1977b) proposed that language makes reference to individuals—concrete or abstract objects in the world—but also to kinds (which he views roughly as indivisible platonic reifications of properties). In a survey article, Carlson (2003) has characterized the NP/DP situation as follows:

- (2.26) a. Reference to individuals is a function of D.
 b. Reference to kinds is a function of N.

The principle in 2.26 is closely related to VZ’s correspondence law 2.20, but is more general since it applies to non-referential nominals as well.

An intuitive application of VZ’s/Carlson’s idea is Longobardi’s (1994) analysis of proper names (and by extension, Postal’s (1969) analysis of pronouns as determiners). Since Jackendoff (1977) it has been observed that in contrast to many other languages, English does not allow proper names to be accompanied by a definite article, e.g.

- (2.27) a. *the John

² This implication exists, of course, only if it is assumed that morphological agreement on an article is both necessary and sufficient for expletive interpretation.

- b. der Hans
the.MASC.sg H.
- c. il mio Gianni
the.MASC.sg my G.
- d. (a) Mari- \emptyset vendég-e- \emptyset
the M.-NOM guest-poss-3sg
“Mary’s guest” (Hungarian, Szabolcsi 1983:89)

with one systematic exception: *modified* proper names can appear with definite and indefinite articles.

- (2.28) a. the Paris that I love b. an exhausted John

Longobardi argues that bare proper names in English are base-generated in N, but (in the absence of a suitable modifier) raise to the higher D position:

- (2.29) $[_{DP} D [_{NP} [_{N} \text{John}]]] \implies [_{DP} [_{D} \text{John}_i] [_{NP} [_{N} t_i]]]$

Longobardi’s (1994) study was mostly syntactic in nature, but we can reconstruct what the semantics of this raising operation would have to be: a name in the N position is interpreted generically, e.g. $[_{N} \text{Paris}]$ would mean “the kind consisting of every instantiation of Paris throughout time,” while $[_{D} \text{Paris}]$ would simply refer to Paris (an individual). Since Paris has had many instantiations in history, the uniqueness presupposition of *the* fails to be satisfied in **the Paris*. This analysis explains the contrastive nature of phrases like *the Paris that I know*, since the intersection of things I know with Paris-instantiations could plausibly be a singleton. See Chapter 5:§2 for further development of this idea.

Returning to the N/D-type/token correspondence, a straightforward hypothesis about the lexical semantics of common nouns and determiners might be

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

The correlation suggested by 2.30 has been argued to be even more fine-grained than just N versus D. Zamparelli (2000) developed a so-called “DP-layers Hypothesis,” which says that even argumental noun phrases come in different sizes (i.e. with different numbers of functional projections), and that the size of a noun phrase determines what kind of semantics it has. The idea of DP-layers is analogous to Rizzi’s (1997) split-C hypothesis for the left periphery, which holds that the clause-type and information structural

properties of clauses are encoded at various structural heights above TP/IP.

Zamparelli's (2000) specific proposal about the interpretation of DP layers is consonant with the above theories of type- versus token-reference in DP: the smallest variety of DP are called KIP (kind phrases) and refer to kinds (a sort of individuals); next largest are PDP (predicative DPs), which denote $\langle e, t \rangle$ properties; and biggest are SDP (strong DPs), which denote either generalized quantifiers ($\langle \langle e, t \rangle, t \rangle$) like *every dog* or individuals (e) like *John*. Cross-linguistic variation in nominal syntax and semantics is then analyzed in terms of which layers are available in what syntactic environments.

In this subsection, I have introduced a number of related proposals correlating the syntactic size of a nominal expression with the ontological category/type of object it denotes. The moral of the story is roughly that the smaller the nominal gets, the more “basic,” “general,” or “context-independent” its interpretation becomes (possibly excluding pronouns). Importantly, the same general trend was observed in the domain of nominal modifiers in Chapter 2—inner modifiers often have “generic” interpretations (or in Bolinger's (1967) terms are “reference-modifying”), while outer modifiers tend to be “episodic” (Bolinger's “referent-modifying”). This parallel constitutes the key conceptual motivation for the theory of DP I present in §5. We return to the issue of genericity in modification extensively in §4.

2.2.5 Kind reference and cross-linguistic considerations

Chierchia (1998) built a cross-linguistically oriented theory of kind-reference upon the NP/DP-type/token correspondence introduced in the previous subsection. The empirical point of departure is the observation that bare plural NPs in English can be interpreted generically while bare plural NP arguments in Romance languages are in most cases simply ungrammatical. The overall semantic theory differs from the proposals of Vergnaud & Zubizarreta (1992), Zamparelli (2000), etc. in various respects, but the relation between the two lines of research should become evident.

An important innovation of Chierchia (1998) is the idea that the denotation of common nouns is subject to cross-linguistic variation. In discussing cross-linguistic differences between the grammatical status of bare nouns in argumental and predicate positions, he states “The denotation of nouns might vary across languages, and this variation might be responsible for the different distributions of bare nominal arguments” (344). The kind of alternations he has in mind are exemplified in 2.31-2.32 for English versus French:

(2.31) a. Doctors are nice.

b. *Médecins sont gentils.
doctors are nice (French)

(2.32) a. *John is doctor.

b. Jean est médecin.
J. is doctor (French)

Chierchia’s theory holds that whether a noun refers to a set of individuals or to a kind can be a matter of cross- and intra-linguistic variation. Crucial to the proposal is a set of universally available type-shifters that convert a noun’s denotation from one type of object to another and back. Cross-linguistic variation in noun reference, then, is derived via a small set of parameter settings, namely [+/-pred] (whether an N(P) can be a predicate) and [+/-arg] (whether an N(P) can be an argument).

Within this parametric approach, Chierchia proposes that English is an NP[+arg,+pred] language. Thus, English nouns can be predicates (and hence refer to sets of individuals) or arguments (and hence have an argumental semantic type, either type e or $\langle\langle e,t \rangle, t\rangle$). When bare NPs are argumental, they have type e and refer to kinds. By contrast, Italian, French, and other Romance languages are NP[-arg,+pred]. In these languages, “an NP cannot be made into an argument without projecting D” (355). In other words, Romance nouns unambiguously refer to sets of individuals, unlike their English counterparts. Within Romance, Chierchia suggests the further distinction that Italian and Spanish (which can have bare NPs in object positions) have null Ds, while French (which always rejects bare arguments) lacks a null D.

Ontologically, for Chierchia kinds are functions from worlds to maximal \oplus -sum individuals, and properties are intensionalized sets (functions from worlds to sets of individuals).³

(2.33) Let P be a property (subset of $D_{\langle e,t \rangle}$), s a situation/world, and d a kind.

- a. $P_s := \{x \in D_e | P(s)(x)\}$ “The extension of P in s ”
- b. $d_s := \sum_{R_s(d)(x)} Atom(x)$ “The mereological sum of all atomic elements of d in s ”

Type-shifting operations play an indispensable role in Chierchia’s framework. The shifters that convert a property into a kind and vice versa, nominalization (\cap) and predicativization (\cup), are defined as follows:

(2.34) Chierchia (1998) type-shifters. Let d be a kind, P be a property, and s be a situation/world

- a. Nominalization: $\cap P := \lambda s [tP_s]$ if $\lambda s [tP_s] \in K$; undefined otherwise.
- b. Predicativization: $\cup d := \lambda x [x \leq d_s]$ (if d_s is defined; false otherwise)

The theory additionally postulates a definite type-shifter ι and an existential type-shifter \exists . The semantics

³ Dayal (1992) and Chierchia (1998) distinguish between semantic objects they call “singular” versus “plural” kinds. Ontologically, Chierchia implements this distinction as one between atomic, non-decomposable “groups”—which model the notion of a singular kind—and functions from worlds to maximal sums—the basic notion of kind introduced in this section. While there are good reasons to introduce this distinction, e.g. the contrast between bare plurals and singular definite kind terms in compatibility with *numerous*, most such reasons are not directly relevant for the semantics of modification. To enhance readability and simplicity, I will therefore ignore this distinction throughout.

of ι is similar to run-of-the-mill lexical entries for the definite article *the*, and is defined as in 2.35.⁴ \exists has a meaning comparable to an indefinite article, defined in 2.36.

(2.35) $\iota X =$ the largest member of X if there is one (else, undefined)

- a. the dogs = $\iota DOGS =$ the largest plurality of dogs
- b. the dog = $\iota DOG =$ the only dog (if there is one)

(2.36) $\exists X = \lambda P[\exists y[X(y) \wedge P(y)]]$ ($\exists : \langle e, t \rangle \Rightarrow \langle \langle e, t \rangle, t \rangle$)

Also crucial to Chierchia’s analysis are a competition principle and an economy constraint against superfluous structure. Both are used to capture cross-linguistic differences in the interpretation of bare plurals and of kind reference in general.

(2.37) **Blocking Principle** (“Type shifting as last resort”): For any type shifting operation τ and any X : $*\tau(X)$ if there is a determiner D s.t. for any set X in its domain, $[[D]](X) = \tau(X)$.

(2.38) **Avoid Structure**: Apply SHIFT at the earliest possible level. (And ranking: $\cap > \{\iota, \exists\}$.)

Here is an example of these principles in action: Chierchia posits a null determiner ∂ for Italian and Spanish bare plural arguments (when grammatical). ∂ has an underspecified semantics represented as a variable SHIFT over type-shifting operations (387). According to the Blocking Principle, ∂ cannot have ι as its value because the existence of definite articles in Italian/Spanish blocks it. Instead, the default value for SHIFT is \cap , and when unavailable (e.g. when NP has no kind correlate), \exists can also be used. In English, SHIFT is available at the NP level—since NPs can be arguments—whereas it is not in Spanish or Italian—since those languages are hypothesized to have only DP arguments. Chierchia argues that the choice of SHIFT over projecting D accounts for why English bare plurals but not definite plurals can refer to kinds.⁵

Cross-linguistic variation in noun semantics suggests cross-linguistic variation in determiner semantics:

“[In] argumental languages, determiners will have to apply to kinds; but this can be obtained in a straightfor-

⁴ Unfortunately, he fails to specify whether t is meant to apply to expressions of type $\langle e, t \rangle$ or of type $\langle s, \langle e, t \rangle \rangle$. In §5 it will be crucial to make a choice on this issue, and I will adopt the latter, intensional type. This means that 2.35 should strictly speaking be read $\iota X_{\langle s, \langle e, t \rangle \rangle} = \dots$ or more perspicuously as $\iota(\lambda w[\lambda x[X_w(x)]])) = \dots$

⁵ Chierchia makes an important but understated point about the analysis of English kind-reference and bare arguments: the analysis involving Avoid Structure “does not extend to singular definite generic *the*, [which] involves singular kinds—whereas plural *the* and \cap involve plural kinds. These two sorts of entities (pluralities and collective singularities) are distinct in our ontology—a distinction we had to countenance on independent grounds. Hence the condition that the meanings be identical is not met when singular generic *the* is involved and Avoid Structure becomes irrelevant. The presence of bare arguments in the language won’t, therefore, bleed uses of singular *the*; but it will bleed generic uses of the plural *the*” (393).

ward way by assuming that determiner meanings have predictable kind-taking variants [and] that languages are free to pick and choose the variant fitting their NP type” (353). The basic idea here is that, where k is the type of a kind, determiners in languages like Italian will take k -type expressions as their first arguments, while determiners in English will take the more traditional $\langle e, t \rangle$ - or $\langle s, \langle e, t \rangle \rangle$ -type expressions. Specifically, Chierchia proposes the general template in (2.39a) for determiners in Italian-type languages:

(2.39) Alternative determiner denotations, (Chierchia 1998:353)

- a. $\text{DET}'(x_k)(P) = \text{DET}(\cup x_k)(P)$
- b. $\text{EVERY}'(\mathbf{dog}_k)(\mathbf{barks}) = \text{EVERY}(\cup \mathbf{dog})(\mathbf{barks})$

So in addition to variation in the grammaticality of bare arguments and predicates, variation in kind reference can be derived as a function of parameter settings including not only [+/-arg] and [+/-pred], but also whether languages have DETs or DET's.

Turning now to some concrete examples, Chierchia analyzes bare plural arguments in English as kind-referring. This yields a straightforward semantics for generic sentences with kind-oriented predicates:

(2.40) a. Dinosaurs are extinct. b. **extinct**($\cap d$)

However, for those instances in which a bare plural is interpreted existentially, as in *Dinosaurs ate my grandma*, an operation called Derived Kind Predication (DKP) is required. Chierchia considers DKP to be a last-resort mechanism which resolves the type/sort mismatch between an individual-oriented predicate and a kind argument. The principle is:

(2.41) *Derived Kind Predication* (Chierchia 1998:364, ex. 34c)

If P applies to objects and k denotes a kind, then $P(k) = \exists x[\cup k(x) \wedge P(x)]$

This principle is essentially an adapted form of the DET' strategy adopted for determiners in (2.39a). DKP provides a straightforward semantics for existential bare plurals:

(2.42) $\llbracket \text{Dinosaurs ate my grandma} \rrbracket = \text{DKP}(\lambda x[\mathbf{ate}(\mathbf{grandma})(x)])(d) = \exists x[\cup d(x) \wedge \mathbf{ate}(\mathbf{grandma})(x)]$
‘Some plurality of dinosaur-instantiations ate my grandmother.’

In Romance languages, which are [-arg], analogues of 2.40 simply cannot be derived because regardless of grammatical number, Romance nouns cannot have type e or $\langle \langle e, t \rangle, t \rangle$. This restriction on Romance

noun denotations has in fact been argued for on independent grounds by McNally (2004), who provides convincing reasons to believe that Spanish bare plural arguments denote $\langle e, t \rangle$ -type properties.

Finally, the prohibition against bare singular arguments in English (and Romance) can be captured as follows: by 2.33, the kind corresponding to dogs in @ is the mereological sum of atomic dog instantiations in @, and is of type e . Therefore (2.43a) is a presupposition failure as there is no unique maximal element in the extension of **dog**, which consists of three atoms. In (2.43b), on the other hand, **dogs** refers to the set of non-atomic dogs, and assuming a lattice-like structure of individuals ordered by \oplus (e.g. Link 1983), the denotation of **dogs** will have a maximal element in @, namely the plural individual **sparky** \oplus **fido** \oplus **rag**.

(2.43) Suppose in @ the set of dogs is {**sparky**, **fido**, **rag**}. Then

- a. $[[dog]]^@ = (\cap \mathbf{dog})(@) = \iota \mathbf{dog}_@ = ???$
- b. $[[dogs]]^@ = (\cap (\mathbf{PL}(\mathbf{dog})))(@) = \iota \mathbf{dogs}_@ = \mathbf{sparky} \oplus \mathbf{fido} \oplus \mathbf{rag}$

2.2.6 Summary of proposals to enrich NP/DP syntax and semantics

Here is a summary of the ideas presented in this section:

Polysemy in noun semantics: Count nouns are three-ways polysemous, potentially having as denotation: an intensional property of individuals (type $\langle s, \langle e, t \rangle \rangle$); a kind (theories diverge as to what type kind terms should have); or a taxonomic/property of sub-kinds reading (type $\langle k, t \rangle$ where k is whatever type kinds end up having).

Hidden parameters in nouns and adjectives: The denotation of a noun can depend upon abstract parameters such as events; similarly for adjectives and kinds.

The semantics of number: The number feature [PL] is semantically basic, and [SG] imposes more stringent semantic requirements.

Category-type correspondence: The syntactic NP/DP distinction is mirrored by the semantic type/token distinction (or kind/individual); this correspondence is manifest differently across languages.

The semantics of nouns: The denotation of a common noun is subject to cross-linguistic variation, as is the ability of determinerless NPs to act as arguments or as predicates.

2.3 The parallel between type-reference in N and generic readings of inner modifiers

We have seen that the structural NP layer is arguably where reference to kinds or types takes place in the semantics of noun phrases, and that conversion to token reference is encoded somewhere above NP within the nominal. The situation was framed by Carlson (2003) as 2.26, repeated here.

- 2.26 a. Reference to individuals is a function of D.
 b. Reference to kinds is a function of N.

I suggest that the existence of this two-tiered system of reference can explain some previously mysterious observations about interpretations of inner modifiers. In particular, I propose the following correspondence, which is more fully spelled-out in the discussion that follows.

(2.44) Generic properties of inner modifiers are the result of composition with a kind-denoting noun.

The following discussion summarizes what is meant by “generic properties,” and spells out more explicitly why 2.44 is a plausible generalization.

Larson (2000b), Larson & Marušič (2004), Umbach (2006), and others have observed that inner modifiers, intuitively speaking, often have something semantically to do with genericity. This can be seen in the following data, mostly repeated from Chapter 1:

- (2.45) a. Olga is a beautiful dancer. (Though she was clumsy tonight)
 “*Generally*, when Olga dances, it is beautiful.” **Manner/Adverbial**
 b. Our Thursday meeting (has been moved to Tuesday this week)
 “Our meeting, which *generally* takes place on Thursday” **Temporal**
 c. Sirius is a visible star. (But it currently can’t be seen)
 “Sirius is a star which is *generally* visible.” **Individual-level**

In broad strokes, Larson’s theory states that inner readings result from a variable in the adjective being bound by an NP-internal generic operator/event quantifier Gen/Γ , while outer readings involve an existential event quantifier \exists . This proposal bears a striking resemblance to NP/DP-type/token correspondences discussed in §3—in fact, if our starting point were *this* correspondence, then it is not at all unexpected that the attachment point of a modifier should be correlated with its semantics in exactly this kind of way. However, Larson’s only theoretical motivation for introducing a generic operator inside NP is Chierchia’s (1995) idea that individual-level predication is generic quantification, combined with the observation that only inner adjectives can have i-level readings.

The variables present in adjectives’ semantic representations are either ordinary individual variables, or else hidden parameters such as event variables (described in §3.2 above). And the structural positions of the operators Γ and \exists are fixed by the syntax of English (hence may be different in other languages). In other

words, this analysis reduces the inner-outer distinction to a matter of quantificational force (and scope).

(2.46) **Larson’s DP-internal operators**

$[_{DP} \exists e [AP_{outer} [\Gamma e [_{NP} AP_{inner} N]] AP_{outer}]]$ (Larson & Marusic 2004)

(Only AP_{inner} is in the scope of generic event quantifier Γe)

Bolinger contrasts are semantically non-uniform, with each contrast encoding a truth-conditionally independent meaning distinction. Furthermore, aside from vague intuition that inner readings are somehow “generic,” there is little independent motivation for the postulation of Γ in NP (or \exists in DP, for that matter). And not all inner readings of modifiers are generic; consider the direct reading of *possible* in *Mary interviewed every possible candidate*—here, *possible candidate* just means ‘someone who is possibly a candidate’. It is difficult to find a sense in which this interpretation is “generic.”

I alternatively propose that it is the interpretation of nouns, and not of adjectives or of silent operators, that is responsible for the generic component in inner readings. More concretely, as stated above in 2.44, I propose that inner readings arise when a modifier composes with a type-denoting noun, and that outer readings arise when a modifier composes with a token-denoting noun. This means that inner modifiers must attach at a low enough point in DP that N has not yet been converted to a token-predicate. And outer modifiers must correspondingly attach above the structural site of this conversion. Hence the core component of Larson’s two-domains theory is preserved, namely that inner adjectives bear a tighter syntactic relationship to nouns than do outer adjectives. But the specific analytical means to achieve this generalization are very different than Larson’s: on Larson’s approach the generic component comes from a silent operator whereas on mine the generic component simply comes from the ordinary processes that compositionally interpret noun phrases.

In the remainder of this Chapter, I will use the insights detailed in §1-3 to build a concrete and internally consistent theory of how noun phrase meanings are built. We proceed in two steps: first, I develop a set of principles governing the interpretation of simple noun phrases consisting of determiners and nouns, and show that the theory can derive the basic distribution of interpretations for a small range of noun phrase types in English, German, French, Italian, and Spanish (though most attention is given to English). In the second part, I formalize the intuition expressed in 2.44 above, namely that generic properties of inner modifiers come from composition with type-denoting nouns. Once the principles governing modification are in place, I apply the system to the individual-level/stage-level Bolinger contrast in English. This is the contrast that

perhaps most clearly illustrates the relevance of genericity in modification, and the basic distribution of readings pre- and postnominally in English is derived.

2.4 Building a theory of composition in DP, Part 1: Nouns and determiners

With so many different theories about the syntax and semantics of DP—most of which do not take into consideration the assumptions or results of related theories—it is difficult to identify just what the right set of assumptions is. In §3, for example, I introduced several proposals correlating the syntactic size of a nominal with the semantic domain of its referent. But if any two of these particular proposals were formalized in the same logical language, they would often constitute an inconsistent theory (“theory” qua set of sentences closed under logical consequence, relative to some class of frames/model-theoretic structures). Just one example is the clash between McNally & Boleda’s claim that nouns have lexical type $\langle e_k, \langle e_o, t \rangle \rangle$ and Chierchia’s claim that nouns have lexical type $\langle s, \langle e, t \rangle \rangle$ and some can be shifted to kind-denoting types (none of which is $\langle e_k, \langle e_o, t \rangle \rangle$ on his analysis). Because of this lack of consistency, I will be careful to see that the specific set of assumptions I adopt is concrete and internally consistent.

As a minimal empirical requirement, a theory of nouns should predict the (un)grammaticality and range of interpretations that argument bare nouns and simple definite descriptions can have in a handful of languages. For example, the theory should guarantee that bare plural subjects in English are ambiguous between a kind and an existential reading, while they are typically ungrammatical in Romance languages.

As a minimal theoretical requirement, the theory must be internally consistent, be informed by recent research on NP/DP, and be explicit enough to be formalized as a theory of Ty_2 (see Chapter 1). Most of the assumptions I adopt come from the proposals explicated in §3, but some are new as well, as is the specific collection of example sentences I will be applying them to. The flow of the section is: first some basic assumptions about the lexical semantics of nouns and number morphology (just enough to compute the meaning of a simple definite description); then I present a contrastive cross-linguistic paradigm describing the available readings for several noun phrase-types. Then, simultaneously with the analysis of the paradigm, I motivate and state further assumptions and principles that are necessary (given the basic assumptions) to account for the paradigm.

2.4.1 Basic assumptions about noun meanings and plural/singular

While the technical choices made in this discussion do have theoretical motivation, their status in the theory is that of basic principles/axioms/premises or “assumptions.”

Prior to any type-shifting or syntactic composition, I assume that uninflected nominal roots denote sub-kind predicates; i.e. the polysemy discussed in §2.1 will be explained by means other than multiple lexical entries. For example, the root *dog* will have type $\langle\langle s, e \rangle, t\rangle$ and will be true of those kinds k which are subkinds of the kind “dog.”

$$(2.47) \quad \llbracket \text{dog} \rrbracket = \lambda k_{\langle s, e \rangle} : \mathbf{kind}(k) \cdot [k \leq (\lambda w [\iota x [* \mathbf{dog}_w(x)])]] : \langle\langle s, e \rangle, t\rangle$$


(where \leq is the taxonomic subkind relation)

For example, the kinds “terrier” and “dalmation” will belong to the extension of *dog* (in models compatible with the actual world), as will the kind “dog” itself. To save space, kind-denoting terms are sometimes abbreviated as follows.

$$(2.48) \quad \mathbf{Notational\ convention:} \text{ The term ‘NOUN : } \langle s, e \rangle \text{’ is shorthand for ‘} \lambda w [\iota x [* \mathbf{noun}_w(x)]] : \langle s, e \rangle \text{’.$$

For example the dog-kind expression ‘ $\lambda w [\iota x [* \mathbf{dog}_w(x)]]$ ’ is abbreviated ‘**DOG**.’

Next, following Sauerland (2003); Zweig (2009); and others, I assume that the plural is semantically unmarked compared to the singular in the sense that singular morphology on a noun requires that the noun apply only to atomic individuals while plural nouns have no mereological restrictions and apply to atoms and \oplus -sums alike. However, departing from the aforementioned proposals and appealing to an insight of Déprez (2005), I propose that grammatical number additionally is what converts a noun’s type- or kind-denotation into a predicate of individuals. Déprez achieves precisely this ontological conversion via grammatical number by postulating a projection NumP between NP and DP. So for her, NPs denote kinds, and a phonetically empty Num head hosting number features then converts an NP to a predicate of individuals, which can then be fed to a determiner. I assume this analytical strategy, but will introduce number features internal to NP instead. Specifically, I assume that all agreement features in the nominal, e.g. person, number, gender, are introduced in the specifier of N. For example, English agreement will have the structure indicated below. But since German nouns inflect additionally for gender, the agreement node—which I call ϕ (for “phi-features”)—will host both a number and a gender specification in German, as indicated below as well.



Given the assumption that nominal roots denote subkind predicates, the semantics of ϕ will take a subkind predicate as its first argument. And since it is ϕ that converts kind-reference to individual-reference, an NP consisting of a number specification and a noun will have the type of an individual-predicate. Integrating the unmarked-plural assumption noted above, the lexical entries for [SG] and [PL] that I posit are given in 2.50. On these meanings, a number-marked noun will denote an intensional property (a term with type $\langle s, \langle e, t \rangle \rangle$)—the kind of meaning that standard theories assume NPs to have. More specifically, plural *-s* simply indicates morphological agreement with the higher [PL] morpheme, which is semantically vacuous (as in Sauerland 2003). The corresponding abstract [SG] morpheme “atomizes” noun denotations; i.e. singular number marking filters out all non-atomic individuals from the denotation of a noun. This operation, the inverse of $*$, I will write as \circ ($\circ P$ is defined as $\{x \in P \mid x \text{ is atomic} \}$). So despite the non-standard semantics for nominal *roots*, this theory does generate familiar denotations for full NPs.

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

The only difference between [SG] and [PL] is that the extension of a singular noun will consist solely of atoms while that extension of a plural noun will be closed under $*$ and therefore contain atoms and sums. This echoes the analysis of Sauerland (2003) cited above. Also: note that the final conjunct $x \sqsubseteq z_w$ in these definitions states that x is a mereological part of z_w (i.e. $x \oplus y = z_w$ for some y). Since kinds are ontologically functions from worlds to maximal \oplus -sums and since z is a kind, the statement $x \sqsubseteq z_w$ has the same semantic effect as a formula asserting x to realize the kind z in w via Carlson’s (1977b) realization relation \mathbf{R} . In other words, $x \sqsubseteq z_w$ is equivalent to $\mathbf{R}_w(z)(x)$.

Determiners will have standard denotations, on which they take a restrictor argument (provided by NP) to form a generalized quantifier denoting the set of properties standing in some relation to the predicate defined by the rest of the sentence. For example, *every* has the following denotation.

$$(2.51) \quad \llbracket \text{every} \rrbracket = \lambda P[\lambda Q[\forall x[P(x) \rightarrow Q(x)]]]$$

As a brief preview of how these moving parts work in action, consider the sentence *the dog barked*, and assume that the VP *barked* has as its meaning the $\langle e, t \rangle$ -type term $\lambda x[\mathbf{barked}(x)]$. Composition inside the referential subject nominal *the dog* proceeds as below, yielding the expression on the right as the meaning for the DP *the dog*.

$$(2.52) \quad \begin{array}{c} \text{DP} \\ \swarrow \quad \searrow \\ \text{D} \quad \text{NP} \\ \text{the} \\ \lambda P[\iota y[P(y)]] \\ \phi \\ \text{[SG]} \\ \lambda P[\lambda w[\circ[\lambda x[\exists z \\ [P(z) \wedge x \sqsubseteq z_w]]]]] \\ \text{N} \\ \text{dog}_{\text{SG}} \\ \lambda k : \mathbf{kind}(k). \\ [k \leq \mathbf{DOG}] \end{array} \quad \llbracket \text{DP} \rrbracket = \iota y[\circ[\lambda x[\exists z[: \mathbf{kind}(z).z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]](y)$$

Because \circ syntactically operates on a lambda-abstract, the resulting expression cannot be β -reduced when fed the argument y , resulting in an expression that is somewhat difficult to parse. However, this term is a equivalent to the easier-to-read (2.53a). To complete the semantics for *the dog barked*, we compose the DP meaning with *barked* via function-argument application. After β -reductions, this results in the expression (2.53b), a notational variant of the standard textbook semantics for the sentence in question modulo the \circ -analysis of the singular (shown for comparison in (2.53c)).

$$(2.53) \quad \begin{array}{l} \text{a.} \quad \llbracket \text{DP} \rrbracket = \iota y[\mathbf{atom}(y) \wedge [\circ[\exists z[: \mathbf{kind}(z).z \leq \mathbf{DOG} \wedge y \sqsubseteq z_w]]]] \\ \text{b.} \quad \lambda x[\mathbf{barked}(x)]((2.53\text{a})) \implies_{\beta} \mathbf{barked}(\iota y[\mathbf{atom}(y) \wedge [\circ[\exists z[: \mathbf{kind}(z).z \leq \mathbf{DOG} \wedge y \sqsubseteq z_w]]]]) \\ \text{c.} \quad \lambda x[\mathbf{barked}(x)](\iota y[\mathbf{dog}(y)]) \implies_{\beta} \mathbf{barked}(\iota y[\mathbf{dog}(y)]) \end{array}$$

To summarize: noun roots denote subkind predicates, and number converts these to predicates of individuals, the suitable argument-type for determiners.

2.4.2 A dataset for theory building

I now use the following empirical generalizations, represented paradigmatically in **Table 1**, to motivate additional principles governing the interpretation of basic noun phrases. These principles are built on top of the basic assumptions laid out in the previous subsection. Thus the theory is designed to account for a sharply delimited dataset, and further applications of the theory will therefore constitute points in its favor.

The chapter closes with Part 2, wherein the theory developed here is applied to the individual-level versus stage-level Bolinger contrast in English.

| arg. type ↓ lang ⇒ | English | German | French | Italian | Spanish |
|-----------------------------|---------|-------------------|--------|---------|---------|
| SG DEF | R/K | R/K | R/K | R/K | R/K |
| PL DEF | R | R/K | R/K | R/K | R/K |
| SG BARE | * | * | * | * | * |
| PL BARE | E/K | E ² /K | * | * | * |

Table 1: Grammaticality and available readings for argument noun phrase types.

Key: R = referential;
K = kind-reading; E = existential;
* = mostly ungrammatical as arg.

We begin by comparing bare plurals with definite plurals.

2.4.3 The interpretation of plurals

Consider the contrast between bare plural (row PL BARE) versus definite plural (row PL DEF) noun phrases with count noun heads. As indicated in **Table 1**, the former are ambiguous in English and German between kind and existential readings; and the latter are unambiguously referential in English and are ambiguous in German. In French, Italian, and Spanish, bare plurals are ungrammatical as subjects and definite plurals have the same set of readings that they do in German. These generalizations are illustrated in 2.54-2.58, expanded from Dayal 2004:397(6).⁶

(2.54) (*The) dogs are common pets.

(2.55) (Die) Pandabären sind vom Aussterben bedroht.
the.PL pandas are to extinction face
“Pandas are facing extinction.” (German)⁷

(2.56) *(Les) pandas sont éteint.
the.PL pandas are extinct
“Pandas are extinct.” (French)

(2.57) *(I) cani sono diffusi.
the.PL dogs are widespread
“Dogs are widespread/common.” (Italian)

(2.58) *(Los) arquitectos construyen las casas.
the.PL architects construct the.PL houses

⁶ Where ‘(*...)’ means ... makes the sentence ungrammatical and ‘*(...)’ indicates ungrammaticality if ... is absent.

⁷ Brugger (1993) states that an example directly parallel to 2.55 with the article is ungrammatical; further investigation required.

“Architects build houses.”

(Spanish)

A potential starting point to account for the variation illustrated in (2.54)-(2.58) is a parametric principle, advocated by Chierchia (1998) and others, according to which some languages allow determinerless NPs to act as arguments of a verb, while others do not.

(2.59) **The NP/DP parameter**

- a. In Germanic, NPs can be arguments.
- b. In Romance, NPs cannot be arguments.

It follows from (2.59) that in English and German, a plural NP can serve as an argument of a verb, from which it follows that plural NPs must be able to receive argumental types. The ungrammaticality of bare plural subjects in Romance also follows from (2.59). (2.59) will of course need to be supplemented with something else if it is to account for the ungrammaticality of bare singular count nouns in Germanic, regardless of whether they are intended to be generic or referential:

- (2.60) a. *(The) panda is facing extinction.
- b. *(The) panda ate bamboo yesterday.
- (2.61) a. *(Der) Pandabär ist vom Aussterben bedroht.
 the panda is to elimination face
 “The panda is facing extinction.”
- b. *(Der) Pandabär hat gestern Bambus gegessen.
 the panda has yesterday bamboo eaten
 “The panda ate bamboo yesterday.” (German)

If 2.59 is correct, then grammatical number, gender, etc., must be introduced within NP, unlike Sauerland’s implementation of DP-agreement. However, the assumptions in 2.49-2.50 state exactly this: agreement (ϕ) is introduced in the specifier of NP and therefore agreement applies to both NPs and DPs.

Here is how the kind reading of English and German bare plurals is attained on the present theory. Chierchia’s (1998) nominalization operator \cap (defined in (2.34a)) converts a plural noun with type $\langle s, \langle e, t \rangle \rangle$ into a suitable argument type: since $PL(\mathbf{dog})$ denotes a function from worlds to sets of \oplus -sums, it follows that $\cap(PL(\mathbf{dog}))$ has type $\langle s, e \rangle$ denotes the function from worlds to *maximal* \oplus -sums specified in (2.62b)—which is precisely the function corresponding to the kind “dog,” as desired. In symbols:

- (2.62) a. $PL(\mathbf{dog}) = \lambda w[*[\lambda x[\exists z[: \mathbf{kind}(z).z \leq \mathbf{DOG} \wedge x \sqsubseteq z_w]]]]$
 b. $\cap(PL(\mathbf{dog})) = \lambda s[\iota y[*[\lambda x[\exists z[: \mathbf{kind}(z).z \leq \mathbf{DOG} \wedge x \sqsubseteq z_s]]]](y)]$

Crucially, this process cannot apply to singular NPs: since $SG(\mathbf{dog})$ denotes a function from worlds to sets of atomic dogs, applying \cap to this expression causes a presupposition failure in those worlds with more than one dog. Again, this is the desired result for the languages under consideration and all languages that disallow bare singular arguments (see Dayal 2004 and Krifka et al. 1995) for discussion of languages that lack this restriction).

Bare plurals in English and German also have existential readings, as indicated in **Table 1**. This interpretation is generated by a single application of Chierchia’s Derived Kind Predication (DKP, see 2.41), which composes kind-terms with object-level predicates. For example we can compose the kind-denoting NP meaning $\cap(PL(\mathbf{dog}))$ with the individual-predicate $\lambda x[\mathbf{barking}(x)]$ to yield

$$\begin{aligned}
 (2.63) \quad \llbracket [_{NP} \text{Dogs}] [_{VP} \text{are barking}] \rrbracket &= \mathbf{DKP}(\lambda x[\mathbf{barking}(x)])(\cap(PL(\mathbf{dog}))) \\
 &= \exists x[\cup(\cap(PL(\mathbf{dog}))) (x) \wedge \mathbf{barking}(x)] \\
 &= \exists x[(PL(\mathbf{dog}))(x)(@) \wedge \mathbf{barking}(x)] \\
 &= \exists x[*\mathbf{dog}_@(x) \wedge \mathbf{barking}(x)]
 \end{aligned}$$

Finally, as shown in 2.55, German patterns like Romance and unlike English in having kind-readings of definite plurals. The kind reading of plural definites in these languages, I propose, arises because the definite article in this construction is expletive, i.e. phonologically pronounced but semantically vacuous. Combining Vergnaud & Zubizarreta’s (1992) argument for expletive articles in Romance with 2.64, definite plurals should have the same interpretive options in Romance and German.

(2.64) **Expletive articles in Germanic**

- a. German has an expletive definite article, $D_{[\emptyset]}$.
- b. English does not have expletive articles.

Suppose expletive articles denote identity functions, reflecting their semantic irrelevance. This allows a kind-denoting plural NP—which cannot be an argument in Romance languages—to pass its meaning directly up to the DP level, at which point it can be an argument. And this derives the kind interpretation of definite plurals, explaining what makes English unique in disallowing this reading.

Fortunately, 2.64 has independent syntactic evidence from both empirical and theoretical considerations: empirically, German optionally allows (many) proper names to appear with an agreeing definite article, e.g. *der/*die Hans*, *die/*der Maria*, while English does not; cf. **the John*, **the Mary*. Assuming that *der Hans* and *Hans* have the same referent, it is natural to conclude that articles co-occurring with proper names are semantically empty and/or should count as instances of D_0 (but see Chapter 5:§1). The theoretical argument in favor of 2.64 is an extension of Vergnaud & Zubizarreta’s (1992) proposal about expletive articles in French: expletive articles are licensed by morphological agreement on D, which is why French has them and English doesn’t. Since German does have determiner agreement, the existence of expletive articles in that language is therefore a natural expectation. And again application of **DKP** in German or Romance will derive the existential reading of definite plurals.

It is actually interesting to reflect for a moment upon the fact that English is the only language in the present sample that does not permit plural definites to receive kind readings (though they can receive taxonomic readings, which we turn to shortly). Interestingly, it is also the only language that does not morphologically distinguish singular and plural definite articles, and the only language that exhibits no form of attributive adjective agreement. This correlation is potentially indicative of a real semantic—and not purely morpho-syntactic—difference between singular and plural forms of the definite article in Romance, German, and other languages with such a distinction. The absence of two forms in English could potentially even be linked to its absence of kind readings of plural definites. At this point, though, the connection must remain rather speculative.

2.4.4 The interpretation of non-plurals

Moving on now to singulars (more aptly called “non-plurals” for reasons that will become clear), the following paradigm summarizes the remainder of the facts from **Table 1** above. In all five languages, bare singular count nouns are ungrammatical as arguments, and definite singulars are uniformly ambiguous between a referential and a kind-denoting reading, the latter of which is illustrated by the kind-level predicates included in the sentences in 2.65.

- (2.65) a. **(The) dodo is extinct.* (English)
 b. **(Il) dodo é estinto.*
 the dodo is extinct (Italian, Chierchia (1998) :342)
 c. **(Der) Pandabär ist vom Aussterben bedroht.*
 the panda is to elimination face (German, Dayal 2004:442)

- d. *(La) baleine est réputée être le plus grand mammifère.
the whale is reputed be the most large mammal (French)
- e. *(El) leopardo es fácil de domesticar.
the cheetah is easy to tame/domesticate (Spanish)

Bare singular arguments in Romance are impossible because of 2.59. But what about English and German, which have no restriction against bare arguments? First of all, singular nouns denote (intensionalized) predicates of atomic individuals, and predicate-type nominals cannot be arguments. But what about applying \cap to a singular noun, as was done for plurals? This too will fail, since only atomic individuals can be in the extension of inflected singular nouns; the creation of a kind by \cap requires a noun with a lattice-like structure imposed by \oplus , since kind extensions are maximal \oplus -sums.

Concerning those cases in which Romance arguments can be determinerless (such as Italian direct object bare plurals; see Chierchia 1998), standard procedure in the literature is to postulate a phonetically empty but semantically contentful definite article that must be licensed by a dominating verb (Vergnaud & Zubizarreta 1992; Chierchia 1998; Dayal 2004). Call such hypothetical articles $D_{/\emptyset/s}$. If $D_{/\emptyset/s}$ do indeed exist, then why can't a bare singular noun compose with $D_{/\emptyset/}$ to form a suitable argument? One possibility is that $D_{/\emptyset/}$ is semantically plural.

(2.66) **Semantics of null articles:** $D_{/\emptyset/}$ is plural.

A principle akin to 2.66 would block insertion of $D_{/\emptyset/}$ into a structure with a singular NP, because such a co-occurrence would result in a feature mismatch and agreement would fail to hold. Despite its stipulative nature, 2.66 makes a certain degree of conceptual sense: plural is semantically unmarked when compared with singular, plural should therefore be the default number specification in some sense, and phonetically empty elements ought to have default properties whenever possible.

The final issue we address here is the interpretation of singular definites. A nominal of the form [_{DP} the [_{NP} ...]] can refer to the unique salient individual satisfying the NP predicate. This interpretation—the referential reading—is probably by far the most common sense of *the dog*, for example, as well as probably of *das Hund*, *le chien*, *il cane*, and *el perro*. In addition, all of the languages under consideration have some mechanism by which a definite singular subject can refer to a kind (or be interpreted generically). The referential reading can be derived without special operators or principles: simply take the property-denoting semantics for [_{NP} [SG] dog], and combine it with the meaning of the definite article defined as in 2.35,

i.e. $\lambda P[\iota x[P(x)]]$. The application of *the* to $[_{NP} [SG] \text{dog}]$ will be defined iff there is a unique salient dog in the context of utterance, as desired.

The kind reading of singular definites is less straightforward. Existing accounts of singular kind reference have introduced some kind of ad hoc mechanism into the formalism for exactly this purpose, whether it is introducing an additional type-shifter (Chierchia 1998), introducing a “generic” definite article homophonous with the ordinary singular article (Carlson 1977b; Krifka 2003), introducing an ontological distinction between singular and plural kinds (Chierchia 1998; Dayal 1992,2004), or some combination of these options. Another option would be to let nouns denote kinds directly, and posit an expletive article in singular kind terms (along the lines of Krifka 2003). Applying this strategy to English, though, conflicts with the principle 2.64 which states that English possesses no such articles.

A particularly interesting and well-motivated analysis of singular generics has been advanced by Dayal (1992,2004). Dayal holds that singular definite kind terms are compositionally derived from the ordinary definite article and a non-standard taxonomic interpretation of the head noun. In her words “the singular definite generic is the regular definite determiner quantifying over a domain of taxonomic entities.” According to the proposal, the kind reading of *the lion* arises when the extension of *lion* is a singleton set consisting of the kind “lion.”

While conceptually attractive, details of Dayal’s implementation are difficult to reconcile with cases of nouns that more clearly have a taxonomic reading. If *the African lion* refers to a subkind of lion, then we can only maintain a standard semantics for *the* if the extension of *lion* consists of lion subkinds—not of the kind “lion” alone. Here I suggest precisely this amendment to Dayal’s approach: as guaranteed by the semantics for nouns given above, common nouns denote subkind predicates. Thus *lion* will have as its extension not $\{\mathbf{LION}\}$, but rather $\{k|k \text{ is a subkind of } \mathbf{LION}\}$. Since every kind is a subkind of itself (\leq is reflexive), since all subkinds of \mathbf{LION} are subkinds of lions, and since $k \leq m$ and $m \leq k$ imply $m = k$, it follows that \mathbf{LION} will be the unique maximal element in the extension of *lion* as defined above. Uniqueness and maximality are exactly the presuppositions of the definite article, and therefore composing a noun with a determiner directly—without mediation via ϕ —results in a DP that refers to the most general kind that the head N applies to. So *the lion* will refer to the $\langle s, e \rangle$ -type kind term \mathbf{LION} , *the dog* to \mathbf{DOG} , and so on.

To formalize this sketch, a couple of new assumptions are required. First, we need a more general version of *the* with a polymorphic type (parallel to Larson’s analysis of *beautiful dancer*; see §3.4 above), which allows it to compose with a predicate of unspecified type.

(2.67) **Generalized polymorphic semantics for *the***

$$\llbracket the \rrbracket = \lambda P_{\langle \sigma, t \rangle} [\iota x_{\sigma} [P(x)]]$$

I allow the definite article to compose with a term denoting an intensional property by first saturating that term’s world-argument with @—the constant denoting the actual world.⁸ This allows *the* to uniformly take a single semantic argument in all of its instantiations.

For ι to be well-defined, a generalized definition of maximality (in terms of which ι is defined) is required that guarantees the kind “dog” to be maximal among the subkinds in the extension 2.47. A generalized notion of maximality guarantees ι to pick out $\mathbf{s} \oplus \mathbf{f} \oplus \mathbf{r}$ from the extension $\llbracket \text{PL}(\mathbf{dog})_{@} \rrbracket = \{\mathbf{s}, \mathbf{f}, \mathbf{r}, \mathbf{s} \oplus \mathbf{r}, \dots, \mathbf{s} \oplus \mathbf{f} \oplus \mathbf{r}\}$, while also picking out $\lambda w[\iota[*\mathbf{dog}_w]]$ from the extension of *dog*’s lexical semantics, i.e. $\llbracket \text{dog} \rrbracket = \{\lambda w[\iota[*\mathbf{terrier}_w]], \lambda w[\iota[*\mathbf{dalmation}_w]], \dots, \lambda w[\iota[*\mathbf{dog}_w]]\}$.

Assuming such a notion of maximality is in place, it is sufficient to introduce one final assumption: that the noun in a singular kind-term does not have number morphology (and so is “singular” by name only).

(2.68) **The semantics of kind-denoting singular definites:** Like proper names, the head noun in a kind-referring singular definite does not have a number feature.

Technically speaking, on this analysis, unmarked nouns in definite kind-terms (and such terms themselves) simply lack number features altogether. The idea of a DP without number is not far-fetched: proper names and perhaps even mass nouns constitute good candidates for DPs without grammatical number—it is in some sense a category mistake to ask whether *John* or *water* is singular or plural since they just cannot receive number morphology (without clear coercion). An apt analogy might be proper names occurring with definite articles, as is common in many of the world’s languages. Furthermore, if we take seriously Carlson’s (1977b) influential idea that nouns can act as “proper names of kinds,” then we should expect nouns acting in this capacity to have the same number features as a genuine proper name, namely none. Omitting number from an NP allows the lexical semantics of the noun to trickle up to the D level and combine directly with *the*, which gives us the right meaning for singular kind terms. Granting 2.68, *the dog* has an interpretation built via the following steps:

$$(2.69) \quad \begin{array}{l} \text{a. } \llbracket \text{dog} \rrbracket = \lambda k_{\langle s, e \rangle} : \mathbf{kind}(k). [k \leq (\lambda w[\iota[*\mathbf{dog}_w(x)])]] : \langle \langle s, e \rangle, t \rangle \\ \text{b. } \llbracket \text{the}_{\langle s, e \rangle} \rrbracket = \lambda P_{\langle \langle s, e \rangle, t \rangle} [\iota k_{\langle s, e \rangle} [P(k)]] \end{array}$$

⁸ This will cause complications in intensional contexts but not in matrix contexts.

$$c. \quad \llbracket \text{the}_{\langle s,e \rangle} \rrbracket(\llbracket \text{dog} \rrbracket) = \iota k[\text{kind}(k).k \leq \mathbf{DOG}]$$

In plain English (2.69c) can be paraphrased “the biggest subkind of **DOG**.” Again by reflexivity of \leq , this amounts to exactly the kind **DOG** as the meaning for the generic reading of *the dog*, as desired.

2.4.5 A theory of nouns, summarized


In the previous subsections, we have developed a set of principles which jointly constitute a more-or-less concrete theory of nouns. These principles are summarized here:

The NP/DP parameter. English/German arguments can have category NP or DP. Spanish/Italian/French arguments are always of category DP.


Noun root semantics. Count noun roots uniformly denote taxonomic sub-kind predicates. Grammatical number marking converts noun roots into predicates of individuals. Thus constituents below NP refer to kinds, and nominal constituents above NP refer to individuals—a precise formalization of the NP/DP-type/token correspondence discussed at length in §3.

The morpho-syntax of inflection and agreement. Feature bundles hosting number and gender (and possibly other) features are syntactically introduced in [Spec,NP]. Features trigger Spec-head agreement on the noun (e.g. plural *-s* in English) and in some languages on the determiner as well (e.g. *das~die* ‘the_{neut.sg~neut.pl}’ in German).

English agreement



German agreement



Semantics of number. Number morphology on N operates on the mereological and ontological structure of N denotations. [PL] closes a noun denotation under \oplus -sum formation. [SG] filters out all non-atomic individuals from the denotation of N. Number morphology may also be absent in singular definite kind terms, in which case N surfaces uninflected for number, and projects its lexical semantics to the NP level.

$$2.50 \quad a. \quad \llbracket \text{SG} \rrbracket = \lambda P_{\langle s,e \rangle,t} [\lambda w [\circ [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z) \wedge x \sqsubseteq z_w]]]]]]$$

$$b. \quad \llbracket \text{PL} \rrbracket = \lambda P_{\langle s,e \rangle,t} [\lambda w [* [\lambda x_e [\exists z_{\langle s,e \rangle} [P(z) \wedge x \sqsubseteq z_w]]]]]]$$

(where \sqsubseteq is the part-of relation on \oplus -sum formation)

Covert type-shifting operations. The operations \cap and DKP can be applied to N/DPs, subject to the constraints described above.

2.5 Building a theory of composition in DP, Part 2: Modification

Having established a parallel between the syntactic NP and DP categories on the one hand, and the semantic categories of type/kind and token/individual on the other, a natural question to ask is whether modifiers like attributive adjectives and relative clauses compose within the NP-level—thus constraining the type of entity a DP refers to—or whether they compose external to the NP level—thus restricting the set of individuals in the extension of the NPs/DPs in which they occur. It is not easy to tease apart these two possibilities, because in most cases the two options will result in identical truth-conditional contributions. Disjoint bodies of literature exist which addresses subparts of this question: for example the syntactic position of nominal modifiers (Abney 1987; Svenonius 2008; Alexiadou et al. 2007; Cinque 2010); the semantic interpretation of nominal modifiers (Siegel 1976; Kamp 1975; Morzycki 2005; Morzycki 2014); and how the syntactic position and semantic interpretation of modifiers interact (Larson 1998,1999; Demonte 2008; Kennedy 1999). As mentioned in §1, much of this research assumes a simplified syntax and semantics for nouns. In other words, these theoretical inquiries have not taken into account important subtle intricacies of noun phrases, especially those summarized in §2.

Larson’s influential two-domains model of (adjectival) modification is designed to explain the nature of Bolinger contrasts—systematic ambiguities in adjectival modification along a variety of dimensions (see Chapter 1 for background). At its core, the two tenets of the two-domains model are that attributive adjectives can occupy one of two distinct syntactic regions of DP, and that the interpretation an adjective receives is in part a function of which syntactic region it occupies. Bolinger contrasts are thus a byproduct of the adjective’s underdetermined syntactic role within DP.

(2.70) The two-domains model of nominal modification

- a. Attributive adjectives can appear in two structural tiers: **inner** and **outer**.
- b. The interpretation of an attributive adjective is determined in part by whether it is in an inner or outer position.

This general idea has seen a number of specific implementations over the years. Subsequent syntactically-oriented variants of Larson’s theory, e.g. those of Svenonius (2008), Alexiadou et al. (2007), and Cinque (2010), have with various degrees of success employed and extended Larson’s basic insights to account for cross-linguistic properties of Bolinger contrasts and related phenomena (see Chapter 2). Larson’s specific

theory based upon the principles in 2.70 introduced silent semantic operators at different structural levels in order to guarantee property (2.70b). As I argued in Chapter 1, there is limited empirical support for this implementation beyond the data it was designed to capture. Surprisingly, there have been no attempts—to my knowledge at least—at solidifying the semantic foundations of the two-domains model.

Here I finish laying out the proposed semantic foundations for the two-domains model as characterized in 2.70. Having presented a general theory of DP-internal composition in simple nominals not containing modifiers, I now integrate modification into the theory. There are three components: syntactic rules for the formation of modified nouns; lexical semantics for (here only intersective) adjectives; and a set of principles and parameter-settings for the syntactic form of English versus Romance modified nouns. In keeping with the noun-centric theme, only a few modifier-specific assumptions are explicitly adopted.

2.5.1 A theory of nominal modification

Syntax. First, following Sadler & Arnold (1994), prenominal attributive modification involves a syntactic operation of head-adjunction. The phrase-structure rule in 2.71 states that (some) modified nouns of the form [Adj N] are syntactically complex heads lacking phrasal structure.

(2.71) **Head-Adjunction for A/N (Sadler & Arnold 1994)**

$$N \Rightarrow A N \quad (\text{for certain values of } A)$$

Thus, for example, simple phrases such as *mere mortal* function grammatically like bare head nouns. Departing from Sadler & Arnold, though, I propose that head-adjunction is one of two ways to syntactically integrate a modifier into a nominal. NPs containing postnominal modifiers or outer prenominal adjectives are formed by phrasal adjunction as specified in 2.72.

(2.72) **Phrasal Adjunction**

- a. [XP Adjunct [XP (Spec) [X' X (Comp)]]] (left-adjunction)
- b. [XP [XP (Spec) [X' X (Comp)]] Adjunct] (right-adjunction)

Phrasal adjunction is the operation that integrates relative clauses and outer adjectives (where $X = N$), and even appositives (in which case $X = D$).

Semantics. I hypothesize that head- versus phrase-adjunction is the syntactic distinction that sets apart inner from outer modifiers. And the interaction between adjunction height and the semantics of inflection

is the most general factor that distinguishes the interpretations of inner modifiers from outer modifiers. To elaborate a bit: because agreement is introduced by a functional node at the top of NP, the difference between head-adjunction and phrasal adjunction is semantically significant. In particular, head-adjoined modifiers—i.e. inner modifiers—compose with kind-denoting nominal roots prior to the noun’s conversion to a token-based denotation. Phrasally adjoined modifiers, on the other hand, compose with nouns *after* the noun’s number is fixed and it is converted into an individual-predicate. Since a single adjective can appear as an inner modifier or as an outer modifier, it is necessary to assume that adjectives have variable types, similar to (but different from in an important way) the analysis of *the* above. In particular adjectives are polysemous, being able to compose with constituents with types $\langle\langle s, e \rangle, t\rangle$ (for N heads) as well as $\langle s, \langle e, t \rangle \rangle$ (for NPs). This can be achieved by assuming the following schematic denotations:

- (2.73) a. $\llbracket \text{adjective} \rrbracket \xrightarrow{\text{inner}} \lambda k_{\langle s, e \rangle} [k \leq \text{ADJECTIVE}]$
 b. $\llbracket \text{adjective} \rrbracket \xrightarrow{\text{outer}} \lambda w [\lambda x [\text{adjective}_w(x)]]$

Recall Larson’s (1998) argument that some adjectives can act as predicates of individuals but also of events, and McNally & Boleda’s (2004) proposal that certain intersective adjectives are unambiguously predicates of kinds. Since in the present system, kinds have type $\langle s, e \rangle$, (2.73a) is precisely a kind predicate. So in effect we are combining and generalizing these two ideas: *visible* (and many other adjectives) can apply to individuals or to kinds; the polysemy part comes from Larson 1998 and the kinds part, from McNally & Boleda (2004).

One way of thinking about the two different interpretations would be that—much like nouns versus NPs—adjectives lexically denote subkind-predicates as in (2.73a), and are subsequently converted to intensionalized individual-predicates as in (2.73b) at the AP level. Then it would not be necessary to posit two distinct lexical entries, but rather (2.73a) would be the interpretation of head-adjoined adjectives while (2.73b) would be the interpretation of phrasally-adjoined adjectives. How such a conversion might take place within AP is unclear, but if correct this parallel with nouns would be quite striking.

Finally, I will utilize two composition modes for the combination of Ty_2 terms: function-argument application (with β -reductions to enhance readability), and intersection (Heim & Kratzer’s (1998) Predicate Modification). Function-argument application can apply to $A : \alpha$ and B if B ’s type is of the form $\langle \alpha, \langle \dots \rangle \rangle$. Intersection can apply whenever A and B have the same type and if both types terminate in t .

Language-specific parameters of variation. Within the framework specified so far, nothing constrains whether a modifier adjoins to the left or to the right of a noun, nor whether a modifier adjoins at the head or phrase level. How can cross-linguistic variation be captured given the assumptions developed so far? Just as there are parameters of variation like head-final versus head-initial, I assume the following difference between English and Romance (the English constraint plausibly also holds for German and other languages with predominant Adj-N word order):

(2.74) **Parameter of variation: directionality of head-adjunction**

- a. In English, head-adjunction structures must have the form $[_X [_Y y] [_X x]]$.
- b. In Romance, head-adjunction can have the form $[_X [_Y y] [_X x]]$ or $[_X [_X x] [_Y y]]$.

2.74 reflects the fact that in English, the default position of adjectival modifiers is pre-noun, while in Romance, the relative order between a noun and an attributive adjective is less constrained (largely by the specific lexical items involved). The restriction against prenominal outer modifiers in Romance languages can be captured similarly, by the following additional dimension of variation:

(2.75) **Parameter of variation: directionality of NP-adjunction**

- a. In Romance, NP-adjunction must have the form $[_{NP} [_{NP} \dots] [_{XP} \dots]]$.
- b. In English, NP-adjunction can have the form $[_{NP} [_{NP} \dots] [_{XP} \dots]]$ or $[_{NP} [_{XP} \dots] [_{NP} \dots]]$.

The restriction against prenominal relative clauses in Romance follows from 2.75, but the same restriction in English does not. Other than the vague notion that “heavy” modifiers tend to prefer postnominal positions in English, I offer no principled explanation for this fact in the current study (though one should be sought). One speculative possibility is that if relatives are head-internal, then a left-adjoined relative clause would necessarily involve raising the head noun to a linearly rightward position, a kind of movement that is often assumed to be disallowed (e.g. $*[_{DP} \text{the} [_{CP} \text{that I saw } t_i] [_{NP} \text{dog}_i]]$).

2.5.2 Individual-level versus stage-level attributive adjectives

With respective theories of nouns (§4) and adjectives (§5.1) sketched, we are now in a position to write a toy grammar and generate some derivations for modified nouns. The grammar will generate X-bar phrase structures and compositionally built Ty_2 terms for each node, and a small inventory of composition rules. The output of a derivation for an expression is a phrase-structure tree each leaf of which is annotated with

a Ty_2 term corresponding to the lexical semantics of the expression. To close the chapter, I apply the theory as developed so far to the individual-level versus stage-level contrast for the simple alternation *the visible stars/the stars visible*. §5.3 reflects briefly on the nature and potential source of the restrictions expressed in 2.74 and 2.75. See in particular Larson & Marušič 2004 and Marušič & Žaucer 2009 for background.

As a testing ground, we consider the individual-level versus stage-level interpretations of *visible* in DPs like *the visible stars* and *the stars visible*. Recall that in both English and Romance, the canonical word order is when we see ambiguity—Adj-N for English, N-Adj for Romance—and that the marked word order is when only one interpretation is available—N-Adj for English, Adj-N for Romance. Thus *visible stars* has both an i-level and an s-level interpretation, while *the stars visible* has only the latter.

- | | | | |
|--------|------------------------------------|--------|------------------------------------|
| (2.76) | The visible stars include Capella. | (2.77) | The stars visible include Capella. |
| | <i>✓i-level, ✓s-level</i> | | <i>✗i-level, ✓s-level</i> |


The initial task is to crank out two meaning-structure pairs for 2.76, and to see whether they conform to our intuitions about what 2.76 can mean. With the same means we will then aim to construct a single meaning and structure for 2.77, and show why an additional meaning-structure pair cannot be generated. This will constitute the theory’s analysis of the stage- versus individual-level Bolinger contrast in English.

The principles developed over the last two sections generate the syntactic structures in (2.78a) and (2.78b) for *the visible stars*. *Visible* in (2.78a) is an inner modifier, and hence should have the i-level reading; in (2.78b) *visible* is an outer modifier and should have the s-level reading.

- (2.78) a. [DP [D the] [NP [ϕ [PL]] [N' [N [A visible] [N stars_[PL]]]]]]
 b. [DP [D the] [NP [AP [A visible] [NP [ϕ [PL]] [N' [N stars_[PL]]]]]]]]

Consider first the inner modification structure (2.78a), annotated with semantic terms for each lexical item:

(2.79)



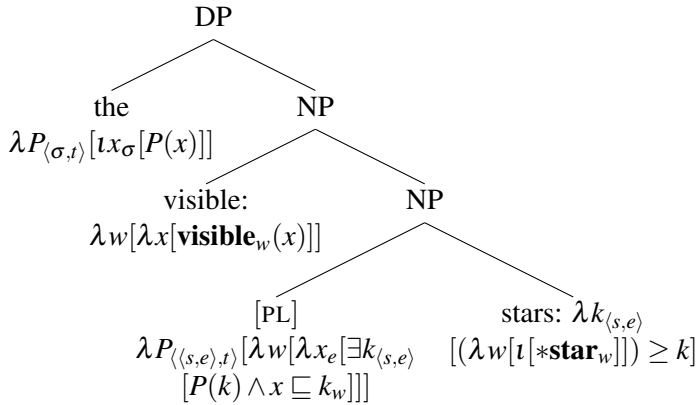
Allowing *visible* and *stars* to compose via generalized predicate modification and all remaining composition via function-argument application, the term in 2.80 is computed as the semantics for the i-level interpretation of *the visible stars*.

(2.80) **Individual-level reading of *the visible stars***

$$\iota x_e [\exists k_{\langle s,e \rangle} [(\lambda w [\iota [*visible_w]]) \geq k \wedge (\lambda w [\iota [*star_w]]) \geq k \wedge x \subseteq k_{@}]]$$

Turning now to (2.78b), the following annotated tree and the lexical entries contain the information necessary to compute the s-level semantics of *the visible stars*. Crucially, observe that the NP $[_{NP} [PL] stars]$ has type $\langle e, t \rangle$. It is therefore impossible for *stars* to compose with *visible*₂, which we used in the i-level derivation, and thus we use *visible*₁ instead.

(2.81)



This structure generates the following term for the s-level reading of *the visible stars*:

(2.82) **Stage-level reading of *the visible stars***

$$\iota x_e [visible_{@}(x) \wedge \exists k_{\langle s,e \rangle} [(\lambda w [\iota [*star_w]]) \geq k \wedge x \subseteq k_{@}]]$$

To compare the terms that our theory generates for the two readings, we embed each of them into the full

sentence *The visible stars include Capella* to facilitate comparison of truth-conditions: the i-level formula is (2.83a) and the s-level formula is (2.83b).


(2.83) The visible stars include Capella.

- a. **include**(**capella**)($\iota x_e [\exists k_{\langle s,e \rangle} [(\lambda w [\iota [*visible]_w]) \geq k \wedge (\lambda w [\iota [*star]_w]) \geq k \wedge x \sqsubseteq k_{@}]]]$)
- b. **include**(**capella**)($\iota x_e [visible_{@}(x) \wedge \exists k_{\langle s,e \rangle} [[(\lambda w [\iota [*star]_w]) \geq k] \wedge x \sqsubseteq k_{@}]]]$)

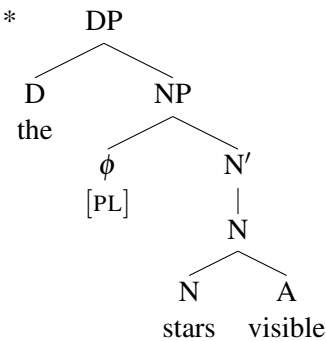
The truth conditions of (2.83a) and (2.83b) are similar but differ in an important way: (2.83a) is true iff Capella is part of the maximal sum-individual x such that for some kind k , (i) k is a subkind of “stars;” (ii) k is a subkind of “visible things;” and (iii) x realizes k in @. In plain English, this means something like “Capella is one of the stars that are characteristically visible.” By contrast, (2.83b) is true iff Capella is part of the maximal sum-individual x such that (i') x is visible in @; (ii') for some kind k , k is a subkind of “stars;” and (iii') x realizes k in @. In plain English, this might be paraphrased “Capella happens to be visible and is one of the objects that are characteristically stars.” The difference between these two sets of truth-conditions, I believe, captures the intuitive difference between the two relevant readings.

The other major consequence the theory should generate is that only the s-level reading is available in *the stars visible*—and the system as defined predicts exactly this. Here is how: as stated above, English allows head-adjunction to the left only. This means that 2.84 is a possible structure while 2.85 is not.

(2.84) **Postnominal visible possible for NP-adjunction only**




(2.85) *



So the theory captures the difference in meaning between prenominal inner and outer *visible*, as well as the ungrammaticality of postnominal inner *visible*. These are the defining empirical characteristics of this Bolinger contrast in English. The principles advanced in §5.1 will also predict the correct distribution of readings in Italian, as described by Cinque (2010).

2.5.3 Why ban rightward head-adjunction (in English)?

The ban against rightward head-adjunction in English is a stipulation. It is an important stipulation, though, because every Bolinger contrast requires it. While I have no definitive explanation for why this restriction should hold, there are a couple of reasonable options to consider. One possibility is that all postnominal adjectives in English (and possibly in all languages) are reduced relative clauses. This is actually quite a common assumption in modern syntactic theory, having been argued for explicitly by Larson & Marušič (2004); Larson (1999); Cinque (2010); Alexiadou et al. (2007); Kayne (1994); Leu (2008), and others. If we adopt it, then it would follow immediately that head-adjunction could only be to the left in English: if all postnominal adjectives are reduced relatives, then they are all phrasal, and phrasal constituents cannot participate in head-adjunction. The reduced relative assumption also forces us to modify our syntax slightly, so that it is CP and not AP that adjoins rightward to NP; as far as I can tell this is harmless aside from the fact that it predicts all postnominal adjectives are predicative. The structure for *the stars visible* under this assumption would resemble 2.86 (some positions in CP are left empty because we have not stated a theory of relative clauses at this point; the only assumption we need for now is that they have $\langle e, t \rangle$ -type semantics).



An alternative explanation for why rightward head-adjunction is impossible in English can be stated in terms of parametric variation. Suppose head adjunction is bi-directional in French (perhaps Romance in general). We could explain this difference as follows: in English, *-s* is a suffix that forms a syntactic constituent with the highest head N only, and the structural position of *-s* always corresponds to its surface position. Therefore if head-adjunction were to occur to the right in English, we would have the ungrammatical word order $*[{}_N N A]-s$ instead of the proper $[{}_N A N]-s$ order (for example, $[{}_N \textit{big dog}]-s$ versus $*[{}_N \textit{dog big}]-s$).

By contrast, recall that in French, both the head noun and attributive adjectives are morphologically marked as plural; in other words, a suffix (typically *-(e)s*) must appear not just on the head noun stem, but also on all elements below the highest N. But suppose that modified nouns contain only a single plural morpheme, which attaches at the highest N, and that languages can differ as to how this morpheme gets spelled out. Then French chooses to spell out plural on all elements below the highest N, while in English plural is a true suffix, combining with the rightmost overt element in N. This would render the relative linear order of N and A irrelevant in French: both [A-s N-s] and [N-s A-s] are grammatical as far as proper morphology goes. But in English, [{A,N}]-s would give us [A N-s] or *[N A-s], the latter of which is ungrammatical. So rightward head-adjunction in English would result in plural marking on the adjective and not the noun, which violates a language-specific morphological/phonotactic property of English. A concrete example: we could have in French [_N *bon chien*]-s⇒[_N *bons chiens*], but also (and more typically) [_N *chien bon*]-s⇒[_N *chiens bons*]. And for English [_N *good dog*]-s⇒[_N *good dogs*], but not [_N *dog good*]-s⇒*[_N *dog goods*].

This second possibility has less existing support in the literature, as far as I know, but seems to make sense conceptually. If concrete predictions could be extracted from this kind of reasoning, it would be an interesting hypothesis to pursue more systematically. To conclude the point, though, there are a number of reasons why it might be that English does not permit rightward head-adjunction.

2.6 Conclusion

This chapter has developed a theory of nouns and noun phrases based upon some recent theoretical advances in the semantics of nouns and noun phrases. I also stated the basics of a new two-domains based theory of modification. The resulting system derives basic properties of the stage-level versus individual-level contrast in English and in Italian. There are a handful of loose ends, but the analysis as it stands is precise enough to be applied to other Bolinger contrasts and modification-related phenomena. And this is the agenda for the remainder of the dissertation.

Chapter 3

Case study: restrictive versus non-restrictive modification

3.0 Introduction

Having introduced some basic architectural assumptions about composition in DP, we now turn to the first major analytical domain of this dissertation: restrictive versus non-restrictive modification. While this distinction is widely referred to in linguistic theory, there is great variation in what exactly the words *restrictive* and *non-restrictive* are intended to mean: some researchers have used the terms as syntactic labels distinguishing between relative clause types, others as semantic labels distinguishing between subsective and non-subsective modifiers, and still others have used the terms in a pragmatic sense, so that a phrase can be restrictive or non-restrictive only relative to a specific utterance in a specific context of use. The issue of precisely what it means for a word or phrase to be restrictive or non-restrictive has rarely been addressed in an explicit way (though see Piñón 2005 and §2.2 for some ideas).

This chapter has two central aims. The first is to develop technically precise definitions for the intuitive terms *restrictive* and *non-restrictive*. In §1.1, I motivate the need for a reexamination of these terms. In §1.2-1.5, I summarize some common conceptions of the distinction within linguistic theory, and argue that no existing definitions for *restrictive/non-restrictive* fully capture the intuitive notions underlying the terminology. §1.6 introduces two additional desiderata for a definition of *(non-)restrictive*, including the most powerful piece of evidence for a discourse-based characterization of non-restrictive modification. In §2, I introduce novel definitions for *restrictive modification* and *non-restrictive modification* in terms of discourse coherence establishment, using insights from existing frameworks for discourse semantics/pragmatics (especially those of Lascarides & Asher (1993,2003); Kehler (2002)) as well as the novel theoretical construct of articulated conversational contexts with active and passive components, motivated by some important

observations of Kripke (2009). In sum, I propose that a nominal modifier is non-restrictive if and only if two conditions hold: (i) that the speaker believes that the modified noun phrase has the same denotation as the unmodified noun; and (ii) that the modifier can be used to deduce an implication that is discourse-related in a particular way to some salient proposition in the sentence or discourse. Restrictive modification is defined (roughly) as the denial of property (i). The definitions are meant to capture in a precise way the intuitions behind working linguists' usage of the terms, i.e. the semantico-pragmatic core underlying all modification structures that have been dubbed "(non-)restrictive" in various research contributions. Because the analysis is built upon a theoretical framework for discourse semantics—specifically a hybrid extension of Asher & Lascarides' and Kehler's theories—I devote §2.1-2.2 to the introduction of discourse semantics/pragmatics as a body of research in linguistic theory, and to the definition of a simple, stripped-down theory of discourse inspired by Asher & Lascarides's (2003) architecture.

The second aim is to apply the newly developed coherence-based characterization of non-restrictiveness to the putative "restrictive/non-restrictive" Bolinger contrast. In §3.1 I show with empirical and conceptual arguments that R/NR is fundamentally different in nature from other Bolinger contrasts, despite its superficially similar appearance. In §3.2 I scrutinize and demonstrate some problems for Larson's (1998)/Larson & Marušič's (2004) and Cinque's (2010) theories of R/NR, irrespective of how (*non-*)*restrictive* is understood. §4 then proceeds to apply the theory of NP/DP from Chapter 2 and the characterizations from §2 (of this chapter) to the syntax/semantics/pragmatics of R/NR. The main conclusion of the new analysis is that R/NR only superficially has the distribution of Bolinger contrasts as described by Cinque (2010), and that restrictive/non-restrictive does not track the inner/outer distinction as related contrasts do. Two novel factors that play crucial roles in the analysis are (i) Larson's (1991) theory of DP shells; and (ii) a novel type mismatch adjustment principle, according to which the composition of a noun (phrase) and a modifier in DP results in an expression that functions syntactically just like a noun (phrase) but is defined only if a particular property holds of the referent of the noun (phrase). As in Chapter 2, variation across binary syntactic parameters is argued to underly different empirical patterns across languages (here English versus Italian). §5 concludes.

3.1 What are restrictive and non-restrictive modification (not)?

The objective of this section is to consider a variety of existing and hypothetical characterizations about what it means for a nominal modifier to be (non-)restrictive, and to show that all of them clash with intuitions

in some way or another about what it means for a modifier to have said properties. In §2, I introduce the discourse-coherence based characterization of non-restrictiveness that will form (part of) the backbone of the analysis presented in §4 of this chapter, as well as some of the discussion in Chapter 6.

3.1.1 Restrictive versus non-restrictive relative clauses

It is useful to begin with some brief comments on relative clauses. In English, two relative clause-types can be distinguished: what I will call integrated relative clauses, (IRCs), more commonly known as restrictive relative clauses, and appositive relative clauses (ARCs), more commonly known as non-restrictive relative clauses. These two modifier types are illustrated in (3.1a) and (3.1b), respectively.¹

- (3.1) a. The book that I bought is on the table. (IRC)
b. The book, which I bought, is on the table. (ARC)

ARCs and non-restrictive adjectives (e.g. *sick* in *my sick mother*) introduce a pragmatically similar kind of information, but the grammatical differences between them are vast (see McCawley 1981; del Gobbo 2003; Schlenker 2013). Non-restrictive adjectives are in fact syntactically more similar to IRCs than they are to ARCs, and IRCs differ from ARCs in several respects: IRCs have a wider range of relative pronoun options, they attach syntactically to common nouns, and have an intersective semantics. ARCs, by contrast, attach syntactically (mostly) to referential nominals, are preceded and followed by a prosodic break, and semantically attribute a property to an individual. While these differences are syntactic, phonological, and semantic in nature, the intuitive notion “restrictive” seems inherently and exclusively semantic in nature. Presumably, restrictive relatives are so called because they “restrict” the denotation of the noun they modify. By the same token, non-restrictive relatives are probably so called because they fail to alter noun denotations—they provide additional information about an independently identifiable referent. In the words of SIL:

(3.2) Definitions from SIL’s *Glossary of Linguistic Terms*

- a. **Restrictive relative clause:** A restrictive relative clause is a relative clause that helps to identify the referent of the word that it modifies.
- b. **Nonrestrictive relative clause:** A nonrestr. relative clause is a relative clause that does not aid in the identification of the referent of its head noun, but only provides information about it.

¹ ARCs are also sometimes called “supplement(al) relative clauses,” but should be distinguished from syntactically less constrained supplement parentheticals such as *I bought it* in *The book—I bought it—is on the table*.

These informal characterizations capture what I believe most researchers have had in mind by “restrictive” and “non-restrictive.” So how can they be made technically precise and generalized to apply to all nominal modifiers? Simplifying matters for the sake of argument, suppose that intersective adjectives, IRCs, PP noun-modifiers (e.g. locatives), and sortal common nouns simply denote (characteristic functions of) sets of individuals. Then whether a modifier is restrictive or not must depend on the actual extension of the modifier and of the noun it modifies in a fixed situation; restrictiveness is therefore a *context-sensitive* property. With these background remarks in mind, we now consider some candidate definitions.

3.1.2 Restriction =_{df} $\llbracket \text{modifier}(\text{noun}) \rrbracket \subset \llbracket \text{noun} \rrbracket$?

An straightforward way to formalize the intuitions in 3.2 is the following: given a noun with denotation $\llbracket \text{noun}_{\langle e,t \rangle} \rrbracket$, a modifier is restrictive if composing the modifier’s denotation with $\llbracket \text{noun}_{\langle e,t \rangle} \rrbracket$ yields a function whose characteristic set is a proper subset of that of $\llbracket \text{noun}_{\langle e,t \rangle} \rrbracket$. One potential point of variation in terminology is whether this proper-subset requirement must hold in all contexts, in some, for all nouns, for some, etc. The following definitions, very much in this spirit, were proposed by Piñón (2005):²

(3.3) Let $[M]$ and H designate functions from situations and objects to truth-values. (Piñón 2005:3)

a. M restrictively modifies H in s iff $\llbracket \lambda x[M_s(x) \wedge H_s(x)] \rrbracket_{\mathcal{M},g} \subset \llbracket H_s \rrbracket_{\mathcal{M},g}$; or

$$\llbracket \lambda x[M_s(x) \wedge H_s(x)] \rrbracket_{\mathcal{M},g} = \emptyset$$

b. M non-restrictively modifies H in s iff $\llbracket \lambda x[M_s(x) \wedge H_s(x)] \rrbracket_{\mathcal{M},g} = \llbracket H_s \rrbracket_{\mathcal{M},g}$

By convention, it is typically modifiers and modifiers only that are referred to as “restrictive” or “non-restrictive.” But if one asks why complements, for example, should be excluded from this distinction, brief inspection provides us with no real reason: complements can have the same intuitive semantic effect on the interpretation of a noun as can an adjective or a relative clause. This point can be illustrated by example. According to commonly used terminology, since attributive adjectives are modifiers but *of*-PPs are (generally) complements, *pretty* in *pretty picture* could be considered “restrictive” while *of Mary* in *picture of Mary* could not. According to common usage, it would be a category mistake to ask whether *of Mary* is restrictive. But 3.4 shows that the procedure of semantic composition is the same in both cases:³ the denotation of the modified noun is simply the intersection of the head noun’s extension with the other phrase’s extension. In

² I leave out two supplementary definitions that Piñón introduces alongside these—they are meant to deal with cases involving, e.g. non-subjective adjectives.

³ There are of course other possible analyses for *of*-PP complements, such as Heim & Kratzer’s (1998) below, but this is not of central importance here.

any context containing more than one picture, both the modifier *pretty* and the complement of *Mary* would “restrict” the denotation of *picture* in exactly the same way (**PM** is predicate modification).⁴

$$\begin{aligned}
 (3.4) \quad a. \quad & \mathbf{PM}(\llbracket \lambda x[\mathbf{pretty}(x)] \rrbracket)(\llbracket \lambda y[\mathbf{picture}(y)] \rrbracket) \\
 & = \llbracket \lambda x[\mathbf{pretty}(x) \wedge \mathbf{picture}(x)] \rrbracket && (\textit{pretty picture}) \\
 b. \quad & \mathbf{PM}(\llbracket \lambda x[\mathbf{of-mary}(x)] \rrbracket)(\llbracket \lambda y[\mathbf{picture}(y)] \rrbracket) \\
 & = \llbracket \lambda x[\mathbf{of-mary}(x) \wedge \mathbf{picture}(x)] \rrbracket && (\textit{picture of Mary})
 \end{aligned}$$

The point of this comparison is just to show that a phrase “restricting” the denotation of a noun seems not to be a sufficient condition for considering that phrase to be restrictive. Conversely, we can also show by example that restricting the denotation of a noun is not a necessary condition for a modifier to be considered “restrictive.” The argument is by contrapositive: given the scenario in 3.5, the modifiers in (3.5a)-(3.5c) are all intuitively restrictive. But the context crucially entails that all students happen to be sick (and displaying flu-like symptoms), and therefore the extensions of *students*, *sick students*, *students with flu-like symptoms*, and *students who have flu-like symptoms* will all be equal (simplifying a bit). In other words, we might say these modifiers are *restrictive* but not *restricting*.

(3.5) **Context:** *A highly contagious flu-bug has caused every single student to become sick, displaying various flu-like symptoms. The school nurse hasn’t seen every kid, but she knows a lot of them are sick. She wants to see every student who is sick or has flu-like symptoms, so she can see how bad the epidemic is. So she gets on the PA and says...*

- a. All sick students should report to the nurse’s office at once.
- b. All students with flu-like symptoms should report to the nurse’s office at once.
- c. All students who have flu-like symptoms should report to the nurse’s office at once.

What 3.5 shows is that a modifier can be considered “restrictive” even if it fails to restrict the denotation of a noun. Taken together with the discussion surrounding 3.4, this demonstrates that the intuitive notion of “restrictiveness” and the set-theoretic notion of “restriction”—combining $\llbracket \mathbf{phrase} \rrbracket$ and $\llbracket \mathbf{noun} \rrbracket$ to yield a proper subset of $\llbracket \mathbf{noun} \rrbracket$ as in Piñón’s (2005) 3.3—are logically independent.

⁴ If *picture* is analyzed as lexically relational, then it should have type $\langle e, \langle e, t \rangle \rangle$ and *of Mary* in this example should be its first argument. However, the fact that *picture* does not require a complement suggests it must be capable of denoting a non-relational property of individuals as well.

A second point of equal importance is raised by the examples in 3.5: the notion of restrictiveness is inextricably linked to speaker intentions.⁵ Consider a slight variant of the scenario, in which everyone in the school *is aware* of the fact that all students are sick. In such a case, none of the sentences in 3.5 would be felicitous at all—the modifiers would either be redundant (and thus violate a Gricean quantity maxim), or lead to the expectation that not all students are sick (which would contradict mutual world-knowledge). But in contrast to the modification constructions in 3.5, an appositive structure with the same semantic content is perfectly felicitous in this second scenario:

(3.6) The students, who are (all) sick with a nasty flu-bug, could easily infect staff members.

I propose that this pattern can be explained by just a few principles: first, that appositives are inherently non-restrictive modifiers while relative clauses and postnominal PPs are inherently restrictive; and second, that a non-restrictive modifier is felicitous only if it is used to contribute an entailment independent of the main assertion of the sentence in which it occurs (the second requirement is elaborated in §1.6). In 3.6, the nurse reminds the school of the state of the students (i.e. reminds everyone that the students are sick) using an appositive; but in (3.5a)-(3.5c), the nurse uses an integrated modifier to specify the subset of students to which she intends to refer. The distinction between these two discourse effects of modification is a theme we return to throughout the chapter.

3.1.3 Heim & Kratzer (1998): Non-restrictiveness \approx Apposition?

In an influential semantics textbook, Heim & Kratzer (1998) introduce the distinction between restrictive and non-restrictive modification. Without stating explicit definitions, they illustrate the difference via example. In discussing the syntax and semantics of noun modification, Heim & Kratzer (1998:63) write that “PPs may appear inside NPs in three distinct semantic roles: as arguments, as restrictive modifiers, or as nonrestrictive modifiers,” and use the following examples to illustrate these three roles:

- (3.7) a. a part of Europe (argument/complement)
 b. a city in Texas (restrictive modifier)
 c. Susan, from Nebraska, (nonrestrictive modifier)

The way that the first two roles are distinguished in Heim & Kratzer 1998 is as follows: complement-taking nouns are relational and have type $\langle e, \langle e, t \rangle \rangle$, so that a complement PP saturates the noun’s first argument

⁵ Thanks to A. Szabolcsi for emphasizing the importance of this point

position. Modifying PPs compose with nouns (both type $\langle e, t \rangle$) via Predicate Modification (intersection):

- (3.8) a. **FA**($[[\lambda x[\lambda y[\mathbf{part}_x(y)]]]])([[\mathbf{europe}]] = [[\lambda y[\mathbf{part}_{\mathbf{europe}}(y)]]]$
 b. **PM**($[[\lambda x[\mathbf{city}(x)]]]])([[\lambda y[\mathbf{in}(\mathbf{texas})(y)]]]) = [[\lambda y[\mathbf{city}(y) \wedge \mathbf{in}(\mathbf{texas})(y)]]]$

The third role of a PP in DP (as a nonrestrictive modifier) they characterize in very different terms:

“The basic intuition that most authors have expressed about the semantics of nonrestrictive modification is that nonrestrictive modifiers are not semantically composed at all with the phrases they modify. Rather, they have the status of separate sentences which serve to make side-remarks of some kind. For example, the meaning of [(3.9a)] is not unlike that of [(3.9b)].

- (3.9) a. *It is surprising that Susan, from Nebraska, finds it cold in here.*
 b. *It is surprising that Susan finds it cold in here. Note that she is from Nebraska*

This makes it reasonable to assume that at the level at which our semantic rules apply, the nonrestrictive modifier isn't part of the structure at all, so the question of how its denotation should be composed with that of its modifier doesn't arise in the first place” (p64).

This view is interesting in a number of ways. First, it raises difficult syntactic questions, such as whether appositives are adjoined to DPs, or whether the structural relationship between an appositive and the nominal it modifies is something more exotic like multi-dominance (see McCawley 1981; Schlenker 2013 for some discussion). Second, Heim & Kratzer's (1998) characterization raises a deep architectural question: by what means can a sentence-medial phrase be interpreted independently of a phrase it is contained in? This situation has been described as scopelessness by Potts (2005), who developed an articulated theory of appositives in the spirit of Heim & Kratzer's (1998) general suggestions. Finally, and most pertinently for this dissertation, the trichotomy in 3.7, if correct, raises the question of which category adjectives like *sick* and *harmful* in 3.10 should fall into:

- (3.10) a. I take care of my sick mother.
 b. Cigarettes contain harmful carcinogens.

Semantically, the contribution of these adjectives appears to pattern just like what Heim & Kratzer call “non-restrictive modifiers” like *from Nebraska* in (3.9a): the same kind of paraphrasing shown in (3.9b) can be constructed for both examples in 3.10.

- 3.10' a. I take care of my mother. My mother is sick.
 b. Cigarettes contain carcinogens. Carcinogens are harmful.

Syntactically, though, *sick* and *harmful* in 3.10 appear to be garden variety attributive adjectives, which in the default case pattern semantically with restrictive modifiers like *in Texas* in (3.7b)—consider the (near) synonymy of the pairs in 3.11. This similarity makes Heim & Kratzer's (1998) "grammatically non-integrated" theory of non-restrictive modifiers implausible for the adjectives in 3.10.

- (3.11) a. a student from Japan / a Japanese student
 b. a city in Texas / a Texan city
 c. a person who is famous / a famous person

The puzzle is thus that in various contexts, attributive adjectives (in English) appear to display a dual behavior: they appear syntactically unremarkable, but share semantic traits with appositives. It is uses of adjectives like *sick* and *harmful* in 3.10 that I will refer to as "non-restrictive attributive adjectives" (NRAs). It is important to recognize that (non-)restrictiveness cannot be a lexical property, since the very same adjectives in 3.10 can be used in contexts in which they pattern syntactically *and* semantically unlike appositives, and thus must be "restrictive."

- (3.12) a. I take care of my sick client.
 a.' #I take care of my client. My client is sick. ('#' meaning ≠(3.12a))
 b. Cigarettes contain harmful substances.
 b.' #Cigarettes contain substances. Substances are harmful. ('#' meaning ≠(3.12b))

An even more remarkable property of (certain) NRAs is that unlike the appositive in *Susan, from Nebraska*, which pretty clearly predicates Nebraskan-ness of Susan, it is not always obvious to determine exactly what entailment/implication they contribute; the identity of the entailment seems to be highly variable. (3.10b) constitutes a particularly interesting case. Note that the object bare plural *harmful carcinogens* has an existential reading; i.e. (3.10b) entails something like "there is some collection of harmful carcinogens that cigarettes generally contain." Despite this existential interpretation, *harmful* appears to attribute a property to the kind "carcinogens," evidenced by the second part of (3.10b)'s paraphrase (3.10b)': *Carcinogens are harmful*. But in (3.10a) sickness is predicated of my mother, not of the kind "mother." This issue is discussed further in §2.6 and in §5.

The discussion thus far can be summarized in three sentences: whatever it means to be “non-restrictive,” all appositives and certain uses of attributive adjectives appear to have this property. It is not at all clear how the semantic contributions of non-restrictive modifiers are integrated into the truth-conditions of the sentences they occur in. And non-restrictiveness does not appear to be a syntactically characterizable property, since the syntax of appositives and of attributive adjectives are distinct.

3.1.4 Umbach (2006): Restriction \approx Focus?

Umbach (2006) observes that non-restrictive adjectives in German cannot receive intonational focus. The adjective *bunte* ‘colorful’ is most naturally interpreted non-restrictively in 3.13. Narrow focus on *bunte* in (3.13b) leads to the unnatural expectation that there are or could be colorless flowers.

- (3.13) a. In Annas Garten sind bunte Blumen, aber kein Gemüse und keine Bäume.
in A.’s garden are colorful flowers but no vegetables and no trees
‘In Anna’s garden there are colorful flowers, but no vegetables and no trees.’
- b. #In Annas Garten sind BUNTE Blumen (...aber keine farblosen Blumen).
in A.’s garden are COLORFUL_F flowers but no colorless flowers
‘In Anna’s garden there are colorful flowers (...but no colorless flowers)’ Umbach (2006:1)

The same pattern holds of English:

- (3.14) a. In Anna’s garden there are colorful flowers.
- b. #In Anna’s garden there are [COLORFUL]_F flowers.

Umbach attributes the contrast in 3.13/3.14 to the fact that focus evokes alternatives in the sense of Rooth (1985,1992), which is presumably an obligatory grammatical process. That there could be non-colorful flowers contradicts world-knowledge, so (3.13b) is odd. Because focus generally provides new or contrasting information, Umbach argues in light of examples like (3.13b) that NR adjectives are “backgrounded,” though she emphasizes that focus (qua new information) and de-accented material (qua old information) are independent from the restrictive/non-restrictive opposition.

While Umbach’s (2006) analysis is not fully formalized, I understand her central thesis to be that non-restrictive modifiers simply do not factor into the focus-background partition of a sentence; they are in some sense part of a different “dimension” of a sentence’s meaning (she analogizes non-restrictive modifiers with the expressive adjectives of Potts 2003,2005,2007). In other words, the contribution of a non-restrictive modifier is semantically independent from truth-conditional meaning.

Given this independence, Umbach (2006:3) proposes that “In the case of indefinite NPs...the restrictive interpretation leads to the intersection of adjective and noun denotation, while on the non-restrictive interpretation the modifier applies to the kind denoted by the noun. In the case of definite NPs on a non-restrictive interpretation the modifier may also apply to the referent.” She also however proposes (in my understanding of the paper) that in referential noun phrases, a non-restrictive modifier is ambiguous between attributing a property to a kind or to an individual. The one formalization she does provide illustrates the kind-related interpretation of non-restrictive adjectives (where \cap is Chierchia’s (1998) nominalization type shifter, introduced in Chapter 2):

- (3.15) a. Ein klein-e-r Pekinese bellt.
 a.MASC small-INFL-MASC Pekinese barks
 “A small Pekinese is barking.” $\implies \exists x[\text{pekinese}(x) \wedge \text{small}(\cap \text{pekinese}) \wedge \text{bark}(x)]$
- b. Der klein-e Pekinese bellt.
 the.MASC small-INFL Pekinese barks
 “The small Pekinese is barking.” $\implies \exists!x[\text{pekinese}(x) \wedge \text{small}(\cap \text{pekinese}) \wedge \text{bark}(x)]$

While Umbach does not provide a logical translation of the individual-related reading, it would have to be something like the following (for the indefinite version (3.15a)).

$$(3.16) \quad \exists x[\text{pekinese}(x) \wedge \text{small}(x) \wedge \text{bark}(x)]$$

In other words, the only difference between the two readings of a non-restrictive modifier on Umbach’s account is whether the modifier attributes a property to a kind or to an individual.

While conceptually attractive, a crucial piece of the puzzle is missing from this analysis: *how* is it that an adjective can compose with a noun with the result that (i) a property is attributed to the kind corresponding to the noun; and (ii) the noun does not contain a modifier when it composes with the determiner? Given traditional assumptions within semantic theory, this situation is just not technically possible. However, in the theory of Potts (2005), formulating such an analysis is possible. And this is more or less the proposal of Morzycki (2008), to which we now turn.

3.1.5 Non-restrictive modifiers as conventional implicature triggers?

Based on an analogy with appositives, Morzycki (2008), Solt (2009), and others have argued that non-restrictive modifiers trigger Conventional Implicatures (CIs) in the sense of Potts (2005). CIs are secondary, side-issue entailments whose projective properties are similar to—but subtly different from—those of pre-

suppositions.⁶

In Potts' framework, at-issue entailments and CIs are on separate "dimensions" of meaning, and so additional non-standard composition rules must be introduced into the grammar. For example, in Potts's system there is both an ordinary function application rule, and a special CI function application rule. The latter is used to combine constituents one of which is a CI trigger (e.g. an appositive relative clause).

Morzycki (2008) extended Potts' architecture to analyze non-restrictive adjectival and adverbial modification. Morzycki argued that the contribution of non-restrictive adjectives is a conventional implicature. The main argument is that the content of a non-restrictive adjective appears to be a "non-at-issue" entailment, much like the contribution of an appositive or an expressive adjective, categories widely considered to be CI triggers. And the pattern of inference from non-restrictive adjectives does appear to parallel what is observed with expressive attributive adjectives like *damn*, which are inherently incapable of restricting noun denotations. An utterance of 3.17 will give rise to both of the inferences in (3.17a)-(3.17b), just like an utterance of 3.18 will entail both (3.18a) and (3.18b) (on the non-restrictive reading of *unsuitable*).

(3.17) The damn Republicans are at it again.

- a. The Republicans are at it again.
- b. The speaker has a negative attitude toward Republicans.

(3.18) Every unsuitable word was deleted.

- a. Every word was deleted.
- b. The words were unsuitable.

To account for the contribution of non-restrictive adjectives, Morzycki introduces the composition rule Expressive Predicate Modification into the grammar, which approximates a Potts-style CI-version of Heim & Kratzer's (1998) Predicate Modification. Combining an adjective α with a noun β will generate a CI (to the left of \bullet), and the denotation will be the same as β 's (to the right of \bullet). The bullet ' \bullet ' separates the at-issue and CI dimensions. The stipulation on relative order is what allows Morzycki to derive the fact that NR adjectives must occur preminally (in English).

⁶ See Potts 2005:Ch2 for a summary of arguments that CIs should be distinguished from presuppositions, but also Schlenker 2007 for arguments that apparent CI triggers can be viewed as presuppositional items that take obligatorily maximal scope (for heterogeneous reasons).

$$(3.19) \quad \alpha(\mathbf{sup}(\beta)) : t^c \bullet \beta : \langle e^a, t^a \rangle \quad \text{where the relative order of } \alpha \text{ and } \beta \text{ is as indicated.}$$

$$\begin{array}{c} \diagup \quad \diagdown \\ \alpha : \langle e^a, t^a \rangle \quad \beta : \langle e^a, t^a \rangle \end{array}$$

The expression $\mathbf{sup}(\beta)$ in 3.19 denotes the sum of (contextually restricted) elements in the extension of β , and $\alpha(\mathbf{sup}(\beta))$ says that this sum has the property α . When applied to *unsuitable* and *words* in 3.18, *unsuitable words* will have the same denotation as *words*, but will carry with it the CI that (roughly) all contextually relevant words are unsuitable—very close in spirit to Umbach’s (2006) analysis in terms of kind reference. This CI is then passed up to the root level via Potts’s (2005) Parsetree Interpretation procedure, thereby accounting for its projective properties.

On the restrictive reading, *unsuitable* and *word* would combine via ordinary Predicate Modification since both linear order and the quasi-universal inference are stipulated in the composition rule, no further assumptions are required. The approach also extends to a related contrast in adverbs, provided that a generalized polymorphic version of the rule is introduced (see the discussion surrounding examples 3.71-3.72).

While Morzycki’s (2008) analysis fares well with respect to the basic contrasts it set out to analyze, it cannot derive the correct interpretations for cases like *my sick mother* or *a small pekinese*, in which the adjective attributes a property to an individual, and not to a kind or a plurality. Furthermore, adopting a CI-based approach requires us to first adopt Potts’s (2005) multi-dimensional semantic framework. It also assumes that linear order requirements may be built into composition rules, an idea that seems to clash with the fact that composition rules are at their core just metalanguage-defined functions from meaning pairs to meanings.⁷ And if Cinque’s (2010) claim is true that postnominal adjectives can be interpreted non-restrictively in Italian (see 3.69 below), then the linear order requirement is undesirable in the first place. One could postulate N-movement plus pre-movement semantic composition to remedy this situation, but we won’t explore this or other potential technical fixes here.

3.1.6 Further considerations: relevance and predicational variability

Here I introduce two properties that I argue a suitable definition of “non-restrictive” should capture. To date, no analyses of non-restrictive modification have acknowledged—much less accounted for—these properties.

Both will form crucial parts of the discourse-coherence-based definition of non-restrictiveness in §2.

⁷ It should be noted that this remark does not apply to certain categorial grammars that define syntactic and semantic composition in tandem (e.g. Jacobson 1999; Barker & Shan 2008). In such a setting, building linear order into composition rules is syntactically motivated since there are not separate phrase-structure rules in the grammar.

3.1.6.1 The inherently discourse-related component of non-restrictive modification

The first property has to do with the discourse function of non-restrictive adjectives, a topic that has not been systematically addressed in print. Recall the well-worn example 3.18 (repeated here as 3.20), on the non-restrictive reading of *unsuitable*. As noted by Larson & Marušič (2004), Morzycki (2008), and others, this sentence has two main implications, namely (3.20a) and (3.20b). What has *not* been observed is that 3.20 also has a third, more subtle implication, paraphrased in (3.20c).

- (3.20) Every unsuitable word was deleted.
- a. Every word was deleted.
 - b. The words were unsuitable.
 - c. Every word was deleted because the words were unsuitable.

The third implication (3.20c) is roughly that the primary assertion is *explained* by the secondary implication (3.20b). Importantly, no such implication necessarily exists on the restrictive reading (though it can)—i.e., it could be the case that some subset of words were accidentally deleted, and it just turns out that all the unsuitable ones were part of this subset (to be elaborated in §2.3). Notice also that the same kind of implication can arise with an appositive modifier: 3.21 has essentially the same three implications as does 3.18.

- (3.21) All the words, which were unsuitable, were deleted.
- a. All the words were deleted.
 - b. The words were unsuitable.
 - c. The words were deleted because they were unsuitable.

The conclusion is thus that the discourse function of a non-restrictive adjective is quite similar to the discourse function of an appositive (exactly how deep this similarity is is a matter of ongoing research). To use an informal but intuitive term, non-restrictive adjectives are required to contribute content that is “relevant” to the discourses they occur in.

(3.22) **Relevance requirement for non-restrictive adjectives:**

A non-restrictive adjective is felicitous only if it provides a relevant link between the implication that it introduces and another implication of the sentence or discourse in which it occurs.

Something resembling the third implication of 3.18 appears to be present in *all* uses of non-restrictive adjectives, although the implication is not always explanatory in nature. Consider the following diverse examples:

(3.23) *Explanation why (colorful cheetahs \rightsquigarrow beautiful; aggressive cheetahs \rightsquigarrow dangerous)*

- a. The savanna is a beautiful place to visit—many colorful cheetahs live there.
- a.' #The savanna is a dangerous place to visit—many colorful cheetahs live there.
- b. The savanna is a dangerous place to visit—many aggressive cheetahs live there.
- b.' #The savanna is a beautiful place to visit—many aggressive cheetahs live there.

(3.24) *As a result (sickness requires care; being tall does not)*

- a. I have to go take care of my sick wife.
- b. \approx My wife is sick, so/therefore I have to go take care of her.
- a.' #I have to go take care of my tall wife.
- b.' # \approx My wife is tall, so/therefore I have to go take care of her.

(3.25) *Surprisal (expanded from Isabelle Charnavel's p.c. (3.25a))*

- a. Can you believe it—the black president gave a racist speech!
- b. \approx The president's blackness makes it surprising that he would give a racist speech.
- a.' #Can you believe it—the tall president gave a racist speech!
- b.' # \approx The president's height makes it surprising that he would give a racist speech.

While there is variation in exactly how the first two implications are linked (whether it is via explanation, surprisal, etc.), what seems to be common among non-restrictive adjectives is that they introduce some kind of rhetorical link between the contribution of the modifier and some other implication salient in the discourse (in the examples here, within the same sentence). Even though non-restrictive adjectives are truth-conditionally “redundant” by definition, their content must be relevant to the surrounding discourse in some intuitive sense. The /-primed examples in 3.23-3.25 show that infelicity results when such a relevance requirement is not satisfied.

A theory of modification should explain why it is that non-restrictive modifiers are apparently always discourse-related in a way that restrictive modifiers need not be. This property is especially remarkable for non-restrictive adjectives, because they appear superficially to be run-of-the-mill integrated modifiers, lacking the unique syntax and intonation associated with appositives. I will in fact propose in §2.4 that it is

this discourse-related property that constitutes the core of what it means for a modifier to be non-restrictive.

3.1.6.2 Variability in non-restrictive adjectival predication

The second property (foreshadowed in discussion above) has to do with variability in what NRAs attribute properties to. As noted by Umbach (2006), there appear to be cases in which non-restrictive adjectives attribute properties to a kind, but also cases in which they predicate properties to an individual. While Umbach does not attempt to explain how the grammar generates these two possibilities, an adequate general theory of non-restrictive modification must.

One possibility is that the contribution of a NRA is determined by the shape of nominal it occurs in. For example, in the possessive description *my sick mother*, sickness is predicated of my mother; in the universally quantified *every harmful toxin*, harmfulness is predicated of the kind “toxin;” and in another universal DP, *every unsuitable word*, unsuitability is predicated of some contextually salient plurality of words (which could be conceived of as a subkind of “words”).

(3.26) Variability in non-restrictive adjectival predication

- a. *my sick mother* ⇒ “my mother is sick” (individual-predicate)
- b. *every harmful toxin* ⇒ “toxins are harmful” (kind-predicate)
- c. *every unsuitable word* ⇒ “the words are/were unsuitable” (sum/subkind-predicate)

Two grammatical factors potentially relevant for the pragmatics of non-restrictive adjectives appear to be definiteness/referential(ness), and whether the noun is singular or plural. Intuitively, for example, a universally quantified DP (*every unsuitable word*) with a non-restrictive adjective (*unsuitable*) should not be capable of predicating a property of a single individual (say, one of the words).

However, if Umbach’s (2006) assumption is correct that *the small Pekinese* is genuinely ambiguous (between “Pekinese are small” and “a particular one is small”), then the structure of the nominal cannot be the only factor in determining what kind of predication non-restrictive adjectives contribute. But even the individual-predication reading is not trivial to generate in a theoretically natural way. Consider again *my sick mother*. The adjective *sick* presumably forms a constituent with *mother* to the exclusion of *my*. So how is it then that sickness can be predicated of *my mother*, which is not a constituent at all? Furthermore, assuming such a predication is possible, how could *my sick mother* have the semantic type of a DP (here probably *e*) if this very structure generates the entailment **sick**($\iota x[\mathbf{mother}_{me}(x)]$)? Put another way, how is it that the

structure in 3.27 can have semantic type e while also generating the entailment that my mother is sick?

(3.27)

```

graph TD
    DP --> D[D]
    DP --> NP1[NP]
    D --- my[my]
    NP1 --> AP[AP]
    NP1 --> NP2[NP]
    AP --- sick[sick]
    NP2 --- mother[mother]

```

\implies *how??*

- DP has semantic type e (or $\langle\langle e, t \rangle, t \rangle$)
- $\llbracket \text{DP} \rrbracket$ “entails” **sick**($\iota x[\mathbf{mother}_{\mathbf{me}}(x)]$), and satisfies the presupposition that I have exactly one mother.

There is also variability in *how* a non-restrictive modifier introduces information. Contrast (3.28a) with (3.28b), for example.

- (3.28)
- a. I have to go take care of my sick mother.
 - b. Cigarettes contain several harmful carcinogens.

The “non-restrictiveness” of the adjectives in (3.28a) and (3.28b) are distinct in two heretofore unobserved respects. To begin with, in (3.28a), it is not possible to determine whether the adjective *sick* restricts the denotation of *mother*, i.e. whether *sick mother* denotes a proper subset of *mother*. (3.28a) could be uttered in a situation in which the set of individuals satisfying *mother* contains more than one element, some of which do not have the property of being sick (in which case *my* would do the restricting). It might also be uttered when there is only one relevant mother in the domain of discourse. In the former case, though, standard assumptions would not guarantee that (3.28a) presupposes I have exactly one mother—it would merely presuppose that I have exactly one sick mother. Crucially, this is compatible with my having two mothers. On the other hand, in all worlds compatible with what we know, carcinogens have the general property of being harmful. Therefore it is impossible given the way the world is for *harmful* to restrict the extension set of the noun *carcinogens*. In a sense, one cannot even recognize that a non-restrictive reading of *harmful* in (3.28b) is intended if one lacks the requisite knowledge that carcinogens are by essence harmful.

The informational contribution of *sick* and *harmful* in (3.28a) and (3.28b) are qualitatively different, as well. *Sick* contributes to the presuppositional content of (3.28a): it licenses the implication that a *particular* mother—namely the speaker’s mother—is sick. But the function of *harmful* in (3.28b) appears to be to *remind*, as opposed to *inform*, the addressee of the fact that carcinogens are in general harmful. This cannot be part of the presuppositional content of (3.28b) in any obvious sense, because *every* does not have the same kind of presupposition that *my* or *the* have.

3.1.7 Wrapping up §1

To conclude, §1 has introduced a number of candidate characterizations of *(non-)restrictiveness*, and all of them were shown to clash with intuitions in at least some respect. What is missing from existing theories is an analysis of discourse function—a factor that I have argued is a defining feature of what it means for a modifier to be “non-restrictive.”

3.2 Defining *(non-)restrictive* in terms of discourse coherence establishment

The primary thesis of this of this section is that

- (3.29) A modifier is *non-restrictive* if and only if
- (i) it can be deleted without altering the truth-conditions of the sentence in which it occurs (after potential differences in presupposition are factored out); and
 - (ii) it introduces a rhetorical link between the implication the modifier generates and another discourse unit in the active context.

In order to state this definition precisely within a formal system, we will need to define some theoretical vocabulary to talk about discourse-level phenomena with (including crucially “active context”). To this end, I begin the section by sketching a formal framework for discourse semantics. The framework is a stripped-down, simplified, and non-representational version of Lascarides & Asher (1993,2003) Segmented DRT, a dynamic theory built upon Kamp’s (1981); Kamp & Reyle’s (1993) Discourse Representation Theory. With the system suitably defined, I present in §2.4 definitions for *restrictive modification* and *non-restrictive modification*, as well as a couple of other related concepts that will prove to be useful in subsequent sections and chapters of this dissertation. I discuss some advantages of the new coherence-based theory over the approaches outlined in §1, and show that the the coherence-based theory successfully captures the core meaning component common among syntactically heterogeneous phrases that have been referred to as “non-restrictive” by various researchers.

3.2.1 Crash course on theoretical discourse semantics

In contemporary generative linguistic theory, the grammar of a language is typically partitioned into discrete tiers of structure: phonological, morphological, syntactic, and semantic. A far lesser studied component of language structure is conversational or discourse structure—that is, the rules, principles, and regularities

about how sentences are ordered and pragmatically linked to one another to yield a well-formed, felicitous, or “coherent” discourse.

3.2.1.1 Motivation for a formal model of discourse coherence

Discourse coherence relations, variously referred to as “rhetorical relations,” “discourse relations,” “coherence relations,” etc., are a family of theoretical constructs designed to provide principled, predictive theories of the semantic and pragmatic relationships between elements of a multi-speech-act text, in as precise a fashion as possible (see e.g. Hobbs 1985 for foundational discussion).⁸ To take a simple example, the contrast between (3.30a) and (3.30b) can be explained in terms of discourse coherence as follows: an abstract pragmatic relation RESULT holds between the two clauses of (3.30a)—we can infer this inductively from a combination of lexical semantic knowledge (RESULT imposes temporal constraints on its relata) and world knowledge (we know that sweating is usually the result of heat). Absent unusual context, the clauses of (3.30b) by contrast are intuitively unrelated.⁹ Put another way, there is no coherence relation that can associate the clauses of (3.30b) given the current discourse state. As a result, (3.30b) is felt to be a bizarre sequence of sentences. Such abstract relations, referred to in SMALL CAPS, are viewed as object-language symbols in formal languages for discourse semantics.

- (3.30) a. It’s hot outside. I’m sweating. (RESULT)
b. #It’s hot outside. My dog’s name is Rags.

A diverse collection of particles (*too, even*), coordinators (*but, because*), and other natural language expressions have been argued by Asher (1993); Lascarides & Asher (1993,2003); Asher et al. (2001); Hardt & Romero (2004), and others to overtly realize coherence relations. Such expressions are sometimes referred to as “coherence markers.” Omission of coherence markers can highly degrade the felicity of certain texts:

- (3.31) *Context:* Who came to the party?
a. Most of my friends showed up. My girlfriend did, [too/as well].
b. #Most of my friends showed up. My girlfriend did.

⁸ The study of coherence relations in linguistic theory is still in its infancy, although coherence relations have been studied from the perspective of artificial intelligence and automated reasoning for decades (see, e.g. Hobbs 1985,1982,1990; Thomason & Hobbs 1997; Scha & Polanyi 1988; Polanyi 1988, and many others).

⁹ That is, unrelated absent supporting contextual information—which is not available when inspecting isolated sentences on a computer screen or paper.

Additive particles such as *too* mark a PARALLEL relation between two clauses, which holds when the clauses syntactically or semantically “match” in some intuitive way (exactly what the matching condition is varies from theory to theory). Likewise *however* and *but* mark CONTRAST relations between clauses/sentences. Many lexical items are compatible with a wider range of relations, the best known such example being *and* (see especially Txurruka 2003). In the following three examples, *and* expresses NARRATION, RESULT, and VIOLATED EXPECTATION, respectively.

- (3.32) a. John became tired and he collapsed on the couch. (NARRATION)
 b. John slipped and he broke his leg. (RESULT)
 c. John is 120 years old and he’s still self-sufficient! (VIOLATED EXPECTATION)

Coherence relations frequently remain linguistically unmarked altogether, as in the case of sequenced declarative sentences (e.g. (3.30a)). The following sentences contrast minimally with their counterparts in 3.32, but the respective coherence relations can hold in 3.33 because juxtaposition is compatible with a wider range of relations than is *and*. Thus, the temporal order of the events described in (3.33a) can vary, depending on whether EXPLANATION or RESULT links the two clauses (cf. (3.32a)).

- (3.33) a. John became tired. He collapsed on the couch.
 b. John slipped. He broke his leg.
 c. John is 120 years old. He’s still self-sufficient!

The semantics and pragmatics of coherence relations are difficult to define precisely. Enumerating all possible coherence relations is probably a futile effort, though there have been a surprising number of attempts. Nevertheless, there are interesting generalizations that can be made about *classes* of coherence relations. One can distinguish, for example, between those relations that require structural similarity between their two arguments (*relata*), and those that don’t (Kehler 2002).

3.2.1.2 A hierarchical model of discourse: Lascarides & Asher 1993,2003, et seq.

The most influential generative theory of discourse semantics is the representational theory of Asher & Lascarides (AL), which has been altered, refined, extended, etc. in a series of papers and books spanning two decades (see Asher & Lascarides 2013,2003 for a comprehensive set of references). Their theory is a system of dynamic semantics, built upon an extension of the language of Discourse Representation Theory (DRT, Kamp 1981; Kamp & Reyle 1993), which they call Segmented DRT (SDRT). In what follows, I will

try to present the core concepts behind AL's theory, while avoiding the formal details of SDRT as much as possible (as they are tangential to present goals).

The central idea behind AL's theory is that discourses have hierarchical structure in much the same way that syntactic phrases do—a view also assumed, but implemented in a different way, by Büring (2001,2003); Roberts (1996), and others. Thus the analysis of a discourse consists of the incremental construction of a tree structure whose terminals correspond in some way to sentences, utterances, or entailments of the discourse (called the minimal discourse units, or MDUs), and in which each vertex connecting two nodes represents the coherence relation that links them. The shape of a discourse tree, then, is determined by two factors: the order in which the sentences are uttered, and the *kind* of coherence relations that link the sentences of the discourse.

The most important distinction among kinds of coherence relations is between those that are *coordinating* versus those that are *subordinating*. Subordinating relations such as EXEMPLIFICATION and (potentially) EXPLANATION create dominance relations among MDUs in a text, similar to how phrase structure creates dominance relations among nodes in a syntactic tree. When S_1 stands in some subordinating relation to S_2 (which we abbreviate $R_S(S_1)(S_2)$), then S_2 typically qualifies or comments upon the content of S_1 ; put another way, S_1 provides a topic for S_2 . Subordinating relations are what create hierarchy in a discourse tree—if S_1 subordinates S_2 , then S_1 will dominate S_2 in the discourse tree (again, see Roberts 1996 and Büring 2003 for alternative conceptions of discourse trees that also postulate topic-comment-induced hierarchy). Coordinating relations such as NARRATIVE and PARALLEL link pairs of sentences that constitute a narrative, are syntactically/semantically parallel or contrasting statements, or that both comment upon some previous sentence that subordinates them both. The two relata of a coordinating relation are sisters in a tree, and the corresponding notation $R_C(S_1)(S_2)$ is used.

Determining whether a given relation is subordinating or coordinating is not a straightforward task. There are a number of properties that seem common to relations of each type, e.g. that “with coordination, there is a temporal progression of the events presented, whereas with subordination, this progression is broken” (Asher & Vieu 2005:596). We will not be concerned here with examining specific properties of individual coherence relations, but see Asher & Vieu 2005 for extensive research on the distinction between discourse coordination versus subordination. An important assumption of this approach is that coordination and subordination are semantic properties of coherence relations and thus are static: an individual coherence relation is either always coordinating or always subordinating.

The above concepts can be illustrated with Lascarides & Asher's (1993,2003) well-worn example about John's wild night:


(3.34) **John's wild night**

- a. John had a great evening last night.
- b. He had a great meal.
- c. He ate salmon.
- d. He devoured lots of cheese.
- e. [_{e₁} He then left the restaurant] and [_{e₂} won a dancing competition].
- f. #The appetizer was calamari.

Intuitively, this discourse becomes deviant down at point f. because e. ends the description of John's meal, and f. tries to restart that description. In 3.34, b. is related to a. via ELABORATION, c. to b. also via ELABORATION, and d. to c. via NARRATION. Crucially, e. is not linked to its preceding sentence d., but rather is a narrative continuation of b. The reason why f. is odd when appended to a.-e., then, is because it attempts to continue the narrative from c. and d., but e. has already closed off that narration, making it inaccessible for subsequent attachment. More precisely, the deviance of f. is due to the fact that NARRATION is a coordinating relation, and hence its relata must be sisters in the tree, while ELABORATION is a subordinating one, so its first argument must dominate its second. In 3.34, c. and d. are dominated by b., but e. is situated at a higher structural level of the text; as a result, certain attachment possibilities for f. are no longer available.


Graphically, the hierarchical relations among a.-e. in 3.34 can be represented as in 3.35. The so-called right frontier constraint of Scha & Polanyi (1988) requires new sentences to attach to rightmost constituents in a discourse tree (which in 3.35 are a., e., and e₂). That d. is not on the right frontier of the tree constructed prior to f.'s integration is AL's explanation for why f. is unable to establish a coherence relation with d.

- (3.35) **Discourse structure tree for a.+b.+c.+d.+e.+...**
 ⇒ f. cannot be integrated because d. is not on the right frontier



But suppose the order of e. and f. in 3.34 were reversed. Then f. would be capable of attaching to d., thus resulting in a well-formed tree that models a coherent discourse. In other words, reversal of e. and f.'s order in 3.34 results in a considerably more natural text, and a well-formed tree can then be constructed:

(3.36) **Discourse structure tree for a.+b.+c.+d.+f.+e.**



In AL's framework, integrating a new utterance into an existing discourse tree is a multi-step process. The basic strategy is that one first assumes that the incoming MDU is rhetorically related to some MDU in the existing tree by some coherence relation, and then uses some kind of knowledge base along with the existing segments of the discourse to infer which coherence relation must have been intended, and which MDU it should be linked to. Whenever the right frontier contains more than one node, attachment site for a newly parsed MDU is underdetermined, so that discourses can be structurally ambiguous. The following is a highly simplified summary of AL's SDRT update procedure, which defines how new sentences are integrated into a discourse tree. Let $\alpha_1, \dots, \alpha_n$ be a coherent sequence of speech-acts, to which speech-act β is to be appended, and let Δ_n be the discourse tree built from $\alpha_1, \dots, \alpha_n$.

(3.37) **General strategy for updating a discourse tree**

- a. Break down β into MDUs β_1, \dots, β_k (assume $k = 1$ for simplicity, i.e. β is a simple declarative clause with no presuppositions);
- b. assume that for some accessible α (most commonly α_n), some coherence relation holds between α and β_1 ; and
- c. use this assumption, along with the knowledge base and the information from $\alpha_1, \dots, \alpha_n$ to prove $R(\alpha)(\beta_1)$ for some specific coherence relation R .
- d. If R is subordinating, then α_n should dominate β_1 ; if R is coordinating, then α_n and β_1 are sisters under a newly inserted placeholder node.

The construction of the tree in 3.35 for 3.34 (minus (3.34f)) follows this process. This can be confirmed by continuing the steps in 3.38 to ultimately result in 3.35.

(3.38) **Building a discourse tree for a.+b.+c.+d.+e.**

- a. Interpret a. and b., and assume $R(\llbracket a. \rrbracket)(\llbracket b. \rrbracket)$ for some R . Solve for R .
- b. We know that dinner is in the evenings, and nothing contradicts the assumption that a. and b. describe events on the same evening.
- c. Since a dinner is part of an evening, we thus conclude that $R = \text{ELABORATION}$, establish

ELABORATION($\llbracket a. \rrbracket$)($\llbracket b. \rrbracket$), and draw $\llbracket a. \rrbracket$

|

$\llbracket b. \rrbracket$

- d. Interpret c. Then we have either $R(\llbracket a. \rrbracket)(\llbracket c. \rrbracket)$ or $R(\llbracket b. \rrbracket)(\llbracket c. \rrbracket)$ for some R .¹⁰ Assume the former case first. Then, ... Now suppose the latter case. Then,...

A fully formalized theory with this kind of architecture would contain constraints on attachment sites, knowledge base axioms, etc. We skip the details here, though, because the only essential parts of SDRT for us are the following:

(3.39) **Summary of key concepts for hierarchical discourse semantics**

- a. A discourse tree consists of minimal discourse units, which are the basic building blocks of a discourse (depending on the theory, they could be sentences, propositions, clauses, etc.).
- b. Every MDU in a given discourse is rhetorically related to some other MDU in that discourse, i.e. is linked via some coherence relation.
- c. Discourse hierarchy is introduced when a subordinating coherence relation links two MDUs.
- d. When a new MDU is introduced, only some of the existing MDUs are accessible for attachment, i.e. there are structural constraints on which portions of a discourse can receive new comments (in most theories, only those on the right frontier).

Zooming out a bit, the idea of reified coherence relations has potentially far-reaching consequences for the theory of grammar in general and for the semantics-pragmatics interface in particular. How many coherence relations are there? How can we tell which relation holds and when? Do all sentences have to be linked by a coherence relation? If coherence relations are operative at level of discourse, can they also operate at lower levels of structure, e.g. between embedded clauses? While most of these questions are far beyond the scope of this dissertation, it will be important for our purposes to address the questions of what kinds of objects coherence relations are and what kinds of objects they can link together (i.e. what the MDUs are).

¹⁰ Both a. and b. are accessible since they are both on the right frontier; if they were linked via a coordinating relation, then only b. would be accessible.

3.2.2 What are the minimal discourse units? An argument from sentence-internal coherence

Coherence-based models of conversation have had success in capturing semantico-pragmatic constraints on multi-speech act texts. However, such theories typically assume coherence relations to be operable at the sentential level, and hence their architecture is not well-suited to analyze discourse coherence internal to mono-clausal utterances (which is required in order to formalize the present analysis of non-restrictive modification). There is little discussion found in the literature explicitly addressing the question of what kinds of objects coherence relations operate on. Lascarides & Asher (1993) states that “Discourse relations are the glue that bind simple constituents together to make more complex constituents” (263-264), but this statement is less than fully informative due to the equivocal nature of *constituent*. There is recent evidence, however, that coherence relations operate in a way that allows them to be embedded and quantified by semantic operators that uncontroversially belong to the grammar proper. In other words, there is evidence that coherence relations can scopally interact with sentence-internal operators, and therefore that MDUs must be smaller than entire sentences or utterances. Here I give two examples.

First, Keshet (2012) has argued that coherence relations can be quantified, and hence are interpreted under the scope of other operators. He cites cases such as the following:

- (3.40) a. Everyone who ate shellfish got sick. (RESULT)
b. Whenever Mary eats shellfish, she gets sick. (RESULT)
- (3.41) a. Every time I get an A one day, I get a B the next day. (PARALLEL)
b. Every time I get an A one day, my GPA goes up the next day. (PARALLEL)

Such examples seem to show that for each instantiation of *time*, a coherence relation is asserted or presupposed to hold. In other words, coherence relations can “embed under quantifiers.” This suggests that they are not only a discourse-level phenomenon, and more interestingly that they can scopally interact with overt quantificational expressions.

Second, Rohde et al. (2011) have recently argued that in the context of “implicit causation verbs” such as *detest*, integrated/restrictive relative clauses can “provide an explanation of the eventuality described in the matrix clause” (Rohde et al. 2011:340). In other words, Rohde et al. claim that a matrix clause can stand in a coherence relation to a clause that is embedded within it. A representative example of this phenomenon is 3.42, which seems to suggest that “[i]n addition to serving the usual function of restricting the reference

of the noun phrase (NP) it modifies (*the coworkers*), in typical contexts the relative clause in 3.42 also generates the inference [that] the coworkers’ arrogance and rudeness constitute reasons why John detests them” (Rohde et al. 2011:340). Notice that the latter inference Rohde et al. identify is essentially identical to the third, explanatory implication from non-restrictive adjectives, introduced in §1.6 above.¹¹

(3.42) John detests the co-workers [_{RC} who are arrogant and rude].

≈“John detests a subset of his co-workers because they are arrogant and rude.”

This kind of observation points to the conclusion that coherence relations can be established between non-matrix clauses. Together, Keshet’s and Rohde et al.’s observations suggest that discourse coherence relations can and should be represented in the grammar of a language—they can be bound and quantified, and are operative at a subsentential level of structure. The question remains what *kind* of objects coherence relations apply to. Are minimal discourse units words or phrases or clauses of a language? Representations thereof? Or something else altogether? As the choice seems to be at least somewhat arbitrary, I propose that minimal discourse units be viewed as propositions, which in the semantics assumed in this dissertation are Ty_2 expressions of type $\langle s, t \rangle$. So for a single monoclausal sentence, we will have MDUs (and hence terminals in the discourse tree) for the main assertion but also for other entailments the sentence might have (e.g. from presupposition triggers or from appositives). Then an example like Keshet’s (3.40a) (=Everyone who ate shellfish got sick) might be represented as in 3.43.¹² If this kind of representation is correct, then sentences cannot be MDUs, and we must use something smaller or more abstract.

(3.43) $\forall x[\text{eat}(\text{shellfish})(x) \rightarrow [\text{RESULT}(\text{eat}(\text{shellfish})(x))(\text{become}(\text{sick})(x))]]$

If we allow coherence relations to be a part of the truth-conditional representation of a sentence, then it seems natural that certain constructions inherently make reference to them. This is a common application of discourse-oriented work on semantics: grammaticized/reified coherence relations have been used as explanatory mechanisms in the analysis of ellipsis (Hobbs & Kehler 1997; Kehler 2000; Hardt & Romero 2004; Asher et al. 2001), anaphora resolution (Asher 1993; Asher & Lascarides 2003; Kehler 1993), par-

¹¹ Rohde et al. (2011) do not discuss how such examples could be integrated into a general theory of discourse semantics, as their primary goal was to determine whether coherence establishment aids in on-line language processing (very interesting results).

¹² How exactly this kind of representation is derived compositionally is not addressed here. Also, it is not obvious how to integrate something like 3.43 into a discourse tree, since the variable x occurs bound in the scope of the RESULT relation, and therefore RESULT does not link genuine MDUs. We will avoid such cases of variable binding throughout, as they are tangential to the primary goal of analyzing non-restrictive modification. In a more comprehensive theory of discourse structure, such issues would be important and interesting to explore.

entheticals (Asher 2000), and presupposition (Asher & Lascarides 1998), for instance. And this is exactly the kind of strategy I propose should be used for the analysis of non-restrictive modification. But before presenting the details of this proposal, we define a simple theory of discourse within which the analysis can be precisely stated.

3.2.3 Defining a simple version of discourse semantics

Explicitly defining “context,” “conversation,” and related notions is essential for a formal theory of discourse semantics. In this section, I motivate and propose the most basic definitions necessary for the purposes of this dissertation (extending what is introduced here could be interesting for comparison to Asher & Lascarides’ and Kehler’s respective frameworks, but this task is beyond the scope of this dissertation).

The notion of “context” that I define is new, and therefore requires some motivation. Such motivation can be provided by reflecting upon an important distinction among presupposition triggers. Presupposition triggers are (roughly) lexical items or constructions that demand some condition be satisfied for their use to be felicitous; e.g. *John quit smoking* is a felicitous statement only if John used to smoke. However, certain presupposition triggers appear to require an overt antecedent; e.g. *#John smokes too* is infelicitous out of the blue despite the fact that it is common knowledge in virtually any situation that many people, and hence someone besides John, smoke(s). *Too* apparently requires its presupposition to be satisfied *overtly*, as in *Bill smokes and John smokes too*. No parallel requirement holds for *quit* and many other lexical triggers. This distinction and related considerations have led some to argue that some presuppositions are anaphoric in the sense that if a clause has presupposition *p*, then some superordinate discourse unit must entail *p* (see Heim 1992; Van der Sandt 1992; Zeevat 1992; Saeboe 1996; Chierchia 2009). Kripke (2009) has used exactly this distinction to argue that there are two notions of “context” relevant for the analysis of presupposition satisfaction/projection (hence for discourse semantics in general):

- (3.44) a. **Active Context:** “[M]aterial that has been explicitly mentioned in the conversation, or is on people’s minds and is known to be on people’s minds, or is highly salient in some way[.]”
- b. **Passive Context:** “[B]ackground information available to the speakers that is not taken as relevant or on their minds.” (Kripke 2009:374)

Kripke’s idea is that some presupposition triggers require an antecedent in the active context—these are the “anaphoric presupposition triggers” (like *too*)—and that others don’t—these are the non-anaphoric triggers.

The difference between *too* and the aspectual verb *quit*, for example, can be naturally analyzed in this fashion. However, a question Kripke does not address in detail is just what kinds of objects active and passive contexts are. His only expressed idea regarding this issue is that the active context should have some kind of complex internal structure that the passive context lacks (373):

“The active context could include a set of questions or topics as well as assertions. [It] might be a complex sort of entity, but it will be the kind of thing that makes uses of *again* and *too* appropriate.”

I propose that Kripke’s notion of an active context should be modeled as an incrementally constructed discourse tree, in which nodes correspond to asserted or presupposed propositions (the MDUs) derived from utterances of a discourse, and edges correspond to instances of discourse relations that link the MDUs to one another. Then, we can state the restriction on *too* semi-formally as follows: an utterance $too(P(x))$ is felicitous only if there is some accessible MDU of the form $P(y)$ for some $y \neq x$ in the discourse tree constructed prior to the introduction of $too(P(x))$. This sketch crucially uses the concepts of discourse hierarchy and accessibility, and seems to capture Kripke’s intuition about the role of an active context. Also, it squares nicely with Van der Sandt’s (1992) (and others’) idea of presupposition as anaphora—in some sense *too* is anaphoric to $P(y)$ in this situation.

Kripke’s notion of passive context can be modeled in the classic way—as a set of propositions called the “context set” (see Stalnaker 1979,2002, and for related notions Stalnaker 1998). The context set is also incrementally updated with each new utterance of a conversation, although it lacks the internal structure that active contexts have. A context set is simply the intersection of all propositions mutually known—and mutually known to be common belief (Stalnaker 2002)—by all conversational participants.

For all of the following definitions, we assume: a set of interlocutors or agents $A = \{a_1, \dots, a_n\}$; a conversation $\sigma_0, \dots, \sigma_n$, which is a sequence of dynamically updated states, a single conversational state σ (which is the kind of object that gets updated as new sentences are uttered); and a non-empty set W (a set of unordered states conceived of as the set of possible worlds/situations). The set of possible sentence meanings—propositions—is defined as $\wp(W)$. Stalnaker’s (1979) notion of common ground, relativized to a set of agents and a spatio-temporal state, can then be defined as the set of propositions that all conversational agents assent to and believe to be mutual belief (Stalnaker 2002).

(3.45) The *common ground* CG at σ , written CG_σ , is the set of all $\varphi \in \wp(W)$ such that

- (i) every $a \in A$ believes (or would assent to) φ (written $B_a(\varphi)$); and
- (ii) if $\varphi \in CG_\sigma$, then $B_a(\varphi) \in CG_\sigma$ and $B_a(B_b(\varphi)) \in CG_\sigma$ for all $a, b \in A$.

The context set can then be defined as the intersection of the common ground, as usual.

(3.46) The *context set* C at σ , written C_σ , is $\bigcap CG_\sigma = \{w \in W \mid \forall p \in CG_\sigma : w \in p\}$

I propose to model Kripke's (2009) concept of a passive context with the context set as defined here. And the discourse trees introduced above will serve in this theory as Kripke's (2009) internally-structured active contexts (trees are mathematically defined shortly).

- (3.47)
- a. The *passive context* at σ is C_σ .
 - b. The *active context* at σ is the discourse tree $T_{\sigma-1}$ that has been constructed prior to the utterance from which σ is introduced.

This allows us to state the restriction on *too* and other anaphoric presupposition triggers precisely, where the satisfaction of a presupposition is modeled via a link to a dominating MDU via some coherence relation. Non-anaphoric presuppositions have a weaker requirement:

(3.48) **Discourse-based hypothesis about presupposition classes**

- a. **Anaphoric presuppositions** are those that must be rhetorically related (by a coherence relation) to an open node (MDU) in the active context, i.e. the current discourse tree.
- b. **Non-anaphoric presuppositions** are those that need not be rhetorically related to an MDU, but still must be entailed by the context set.

This view squares with the idea that *too* introduces a PARALLEL coherence relation (e.g. Asher & Lascarides's (2003)), and also with Kripke's (2009) proposal that there are two notions of context, one with internal structure and one without. However, an important feature of hypothesis 3.48 is that not all presuppositions need to be rhetorically related to some prior MDU, contrary to Asher & Lascarides's (1998) theory of presuppositional SDRT. Because I have introduced both passive and active contexts, non-anaphoric presuppositions are characterized as those that need not be discourse-related, but whose content is merely entailed by the context set (or can be consistently added to it; see Stalnaker 2002; Schlenker 2013 for discussion).

The remaining ingredients for an adequate framework for discourse semantics are (i) a recursive pro-

cedure for incrementally updating contexts (i.e. for adding leaves and edges to a discourse tree, and for reducing the size of the context set); (ii) a procedure for breaking an utterance into MDUs; and (iii) specific constraints, for every individual coherence relation R , on the conditions under which R can link a node in the discourse tree to an incoming MDU.

Concerning the treatment of specific coherence relations (point (iii)), I follow the assumptions of Lascarides & Asher (1993,2003) (without all the associated formalism). Properties of specific relations are explained on a when-relevant basis (the interested reader is referred specifically to Asher & Lascarides 2003:Appendix A, in which all of Asher & Lascarides' proposed coherence relations are enumerated and exemplified).

Concerning (ii), I assume that a sentence α contributes a separate MDU for each entailment, each (anaphoric) presupposition, and each additional side-issue entailment that is grammatically generated from α (e.g. those from appositives, parentheticals, etc.; roughly, for each conventional implicature).¹³ For example, an utterance of 3.49 would require us to integrate the three (simplified) MDUs in (3.49a)-(3.49c) into the discourse tree we had just before the utterance of 3.49. Constraints on the order in which these MDUs are integrated certainly exist, but are tangential to the main phenomena under investigation.¹⁴

(3.49) John quit smoking, which is unfortunate.

- a. MDU₁: $\lambda_w[\neg\text{smoke}_w(\text{john})]$ (assertion)
- b. MDU₂: $\lambda_w[\text{PAST}[\text{smoke}_w(\text{john})]]$ (presupposition)
- c. MDU₃: $\lambda_w[\text{unfortunate}(\lambda_w[\neg\text{smoke}_w(\text{john})])]$ (from appositive)

I simply assume that the semantic component of the grammar is capable of breaking sentences up into appropriate MDUs as in 3.49. This is a substantial assumption, but delving into complicating issues like cross-sentential anaphora would constitute a significant sidetrack.

For ingredient (i), I offer the following informal definition of discourses, along with the update procedure sketched in 3.51 below.

¹³ It is an interesting question to ask whether *scalar* implicatures should be included in this category. The currently popular grammatical approach to scalar implicature (e.g. Chierchia et al. 2011) predicts that a sentence's strengthened truth-conditions should count as the main MDU, i.e. the scalar implicature will not constitute a separate MDU but rather will be part of the sentence's assertion. For example, Chierchia et al.'s (2011) theory predicts that *John ate some of the beans* (typically) has the truth-conditions "John ate some but not all of the beans." Alternative theories of scalar implicature could posit separate MDUs for the implicature. While interesting, systematic investigation is beyond the scope of this dissertation.

¹⁴ While I provide no general theory of such constraints, particular orderings will sometimes be required in subsequent analyses; I will point out situations in which such considerations make a difference.

(3.50) **Definition.** A discourse at stage n is a structure $D_n = \langle \langle C_n, T_n \rangle, \alpha_1, \dots, \alpha_n \rangle$, where

- a. C_n is the passive context at n (the context set);
- b. T_n is the active context at n (the discourse tree); and
- c. $\alpha_1, \dots, \alpha_n$ is a sequence of utterances (those that make up the discourse).

The final object to be defined is the discourse tree. We define *the discourse tree at stage n* , written T_n , in terms of how to update T_{n-1} , by recursion on n . The notation $T_n + \alpha$ refers to the discourse tree from stage n updated by integrating the MDU α into T_n via some coherence relation. Here and throughout, *accessible at n* is intended to mean “on the right frontier of the tree at stage n .” For simplicity I assume that trees are set-theoretically nothing more than sets of Ty_2 formulas exactly one of which has the form $\alpha : \langle s, t \rangle$, and the rest of which have the form $R(\alpha)(\beta)$, where R is a coherence relation symbol (type $\langle \langle s, t \rangle, \langle \langle s, t \rangle, t \rangle \rangle$) and $\alpha, \beta : \langle s, t \rangle$. Despite this simplicity, we can define *dominance*, *precedence*, and related notions in such a way that we are guaranteed a decidable set of right-frontier nodes at any given stage, which is the only essential structural property for discourse update (see Partee et al. 1990:439-450 or Barker & Pullum 1990 for examples of how this could be done in a mathematically rigorous way).

(3.51) **Definition.** A discourse tree T_n updated with MDU α in a discourse D_n is

- a. **Stage 0:** The discourse-initial tree T_0 is empty.
- b. **Stage 1:** If α is discourse-initial, then $T_1 = T_0 + \alpha = \{\alpha\}$ (α is the start symbol).
- c. **Stage n :** If α is not discourse-initial, then $T_n = T_{n-1} + \alpha$, the structure just like T_{n-1} except that for some $\alpha_i \in \text{LEAVES}(T_{n-1})$ accessible at n , we have $R(\alpha_i)(\alpha) \in T_n$ or $R(\alpha)(\alpha_i) \in T_n$.

In case (3.51c), I make the simplifying assumption that each MDU other than the discourse-initial one must attach to exactly one accessible MDU by exactly one coherence relation R . This constraint greatly delimits the class of structures that count as discourse trees, while also enhancing the readability of discourse trees by increasing their resemblance to syntactic phrase markers.

We can illustrate the tree update procedure in 3.51 with Lascarides & Asher’s (1993) example 3.34 from above: a.= *John had a great evening last night*; b.= *he had a great meal*; c.= *he ate salmon*; d.= *he devoured lots of cheese*; e.= [_{e1} *He then left the restaurant*] and [_{e2} *won a dancing competition*]; f.= *The appetizer was calamari*. The start symbol (i.e. the first MDU; a. in 3.52) is always uniquely identifiable because it is the only non-relational statement in a tree.

- (3.52)
- a. $T_0 = \emptyset$
 - b. $T_1 = T_0 + \mathbf{h.g.e.}(\mathbf{john}) = \{\mathbf{h.g.e.}(\mathbf{john})\}$
 - c. $T_2 = T_1 + \mathbf{h.g.m.}(\mathbf{john}) = \{\mathbf{h.g.e.}(\mathbf{john}),$
 $\text{ELABORATION}(\mathbf{h.g.m.}(\mathbf{john}))(\mathbf{h.g.e.}(\mathbf{john}))\}$
 - d. $T_3 = T_2 + \mathbf{ate}(\mathbf{salmon})(\mathbf{john})$
 $= T_2 \cup \{\text{ELABORATION}(\mathbf{h.g.m.}(\mathbf{john}))(\mathbf{ate}(\mathbf{salmon})(\mathbf{john}))\}$
 - e. $T_4 = T_3 + \mathbf{devoured}(\mathbf{cheese})(\mathbf{john})$
 $= T_3 \cup \{\text{PARALLEL}(\mathbf{ate}(\mathbf{salmon})(\mathbf{john}))(\mathbf{devoured}(\mathbf{cheese})(\mathbf{john}))\}$
 - f. $T_5 = T_4 + \mathbf{left}(\iota x[\mathbf{restaurant}(x)])(\mathbf{john})$
 $= T_4 \cup \{\text{NARRATION}(\mathbf{h.g.m.}(\mathbf{john}))(\mathbf{left}(\iota x[\mathbf{restaurant}(x)])(\mathbf{john}))\}$
 - g. $T_6 = T_5 + \exists x[\mathbf{competition}(x) \wedge \mathbf{won}(x)(\mathbf{john})]$
 $= T_5 \cup \{\text{NARRATION}(\mathbf{left}(\iota x[\mathbf{restaurant}(x)])(\mathbf{john}))$
 $(\exists x[\mathbf{competition}(x) \wedge \mathbf{won}(x)(\mathbf{john})])\}$

This derivation alone only gets us a set of relational statements and the root node **h.g.e.(john)**. How to graphically represent this set? One intermediate step that helps to build a structural representation is the following principle, adapted from Asher & Vieu's (2005) *Continuing Discourse Patterns* (CDP) constraint. Assume that all discourse trees are closed under CDP.

(3.53) **Continuing Discourse Patterns (Asher & Vieu's 2005:595 CDP, adapted)**

$$R_S(\alpha)(\beta), R_C(\beta)(\gamma) \implies R_S(\alpha)(\gamma)$$

What CDP does, graphically anyway, is turns structures of the form — into structures of the form \wedge . As an example, CDP guarantees that c. and d. above are both daughters of b. Without CDP, we would only know that b. dominates c. and that d. is sister to c.; CDP adds the information that b. dominates d., as well. This is important mainly for readability purposes, but it also captures the conceptual intuition that if two events are parallel, and one is a sub-event of a larger event, then the other must also be a sub-event of the larger event.

With CDP in place, here is a partial procedure for converting a set of relational statements into a tree-like structure. While incomplete, it is good enough for our purposes here (formulation of a mathematically precise version is research in progress):

(3.54) Let T_n be a discourse tree built from $\alpha_1, \dots, \alpha_n$. To build a graphical representation of T_n ,

- a. α_1 is the root; and
- b. for each α_i, α_k with $i < k \leq n$,
 - (i) if $R_S(\alpha_i)(\alpha_k) \in T_n$ or $R_S(\alpha_k)(\alpha_i) \in T_n$, then α_k is daughter of α_i via R_S ;
 - (ii) if $R_C(\alpha_i)(\alpha_k) \in T_n$ or $R_C(\alpha_k)(\alpha_i) \in T_n$, then α_k is sister to α_i via R_C . And,
- c. whenever $R_S(\alpha_i)(\alpha_k)$ and $R_C(\alpha_k)(\alpha_j)$, then by CDP we have $R_S(\alpha_i)(\alpha_j)$; in this case, α_k and α_j will be sisters, and are both dominated by α_i .

To summarize, a discourse tree T_n generated from a sequence of sentences $\alpha_1, \dots, \alpha_n$ contains as leaves all of the MDUs grammatically derivable from $\alpha_1, \dots, \alpha_n$ (minus the non-anaphoric presuppositions). And each edge of T is individually annotated with the coherence relation that links the pair of nodes. While far from a precise technical definition of trees and related notions, this formalism is enough to introduce the analysis of non-restrictive modification, to which we now turn.

3.2.4 Defining *restrictive* and *non-restrictive* in terms of discourse coherence

All ingredients needed to define the key concepts of this chapter are now in place. The relevance constraint on non-restrictive adjectives, introduced in §1.6 above (ex. 3.22), can be reformulated more precisely in terms of coherence relations as follows:

(3.55) **Relevance requirement for non-restrictive adjectives (in terms of coherence):**

An attributive adjective α uttered at stage n of a coherent discourse is non-restrictive only if the implication generated by α (call it **imp**(α)) is linked to some accessible MDU in T_{n-1} via some coherence relation R .

This constraint alone is not enough to characterize non-restrictiveness. Recall Rohde et al.'s (2011) example *John detests the co-workers who are arrogant and rude*. The relative clause here is restrictive, i.e. this sentence says something about a proper subset of John's co-workers. But nevertheless, there appears to be a relation of EXPLANATION that holds between the relative clause and the matrix clause; i.e. the presupposition (3.56a) and the assertion (3.56b) are linked via the EXPLANATION relation as in (3.56c).¹⁵

(3.56) John detests the co-workers who are arrogant and rude.

- a. $\exists!X[*\mathbf{co-workers}(X) \wedge *\mathbf{arrogant}(X) \wedge *\mathbf{rude}(X)]$ (presupposition)

¹⁵ In Chapter 5§2, I will argue that this perspective makes clear sense of a heretofore puzzling empirical pattern about subtriggering observed by Dayal (1995).

- b. $\text{detests}(X)(\text{john})$ (assertion)
- c. $\text{EXPLANATION}(\exists!X[*\text{co-workers}(X) \wedge *\text{arrogant}(X) \wedge *\text{rude}(X)])(\text{detests}(X)(\text{john}))$

This analysis of Rohde et al.’s 3.42 shows that while establishing a coherence relation may be a necessary condition for a modifier to be considered non-restrictive, it cannot be a sufficient condition. The relative clause *who are arrogant and rude* is intuitively “restrictive,” but still establishes coherence with the assertion. So what is the missing piece? The natural answer would be that non-restrictive modifiers cannot “restrict” noun denotations in the intuitive sense; i.e. composition of a non-restrictive modifier with a noun should always yield a phrase whose extension is identical to that of the unmodified noun’s. Something like this property is important for our analysis—I will call it the property of being “non-restricting.”¹⁶ Define $X[a/b]$ to be the expression just like X except with a token of a replaced by one of b . ε is the empty string.

(3.57) **Definition: non-restricting**

An occurrence α_n of an attributive adjective α uttered in the phrase $[\text{DP} \dots \alpha_n \dots \text{N} \dots]$ or $[\text{DP} \dots \text{N} \dots \alpha_n \dots]$ at index i is non-restricting if and only if $\llbracket \text{DP} \rrbracket^i = \llbracket \text{DP}[\alpha_n/\varepsilon] \rrbracket^i$.

This auxiliary notion of non-restrictingness deals with cases like the nurse-and-sick-students case from 3.5 above: the modifiers in 3.5 are non-restricting according to 3.57, but this does not guarantee them to be non-restrictive, as desired. Rather, what we concluded in §1.2 was that non-restrictiveness additionally requires that the speaker *intends* to use the modifier non-restrictingly. We can thus define the second component of non-restrictiveness in terms of 3.57. The definition then has two parts: the (qualified) semantic/set-theoretic constraint (3.58a), and the discourse-level constraint (3.58b).

(3.58) **Definition: non-restrictive**

An occurrence α_n of a modifier α in the phrase $[\text{DP} \dots \alpha_n \dots \text{N} \dots]$ or $[\text{DP} \dots \text{N} \dots \alpha_n \dots]$ uttered at stage σ of a coherent discourse is *non-restrictive* with respect to index i iff

- a. the speaker believes that $\llbracket \text{DP} \rrbracket^i = \llbracket \text{DP}[\alpha_n/\varepsilon] \rrbracket^i$; and (non-restricting intention)
- b. for some accessible MDU β in $T_{\sigma-1}$ and some coherence relation R , either
- $R(\text{imp}(\alpha_n))(\beta) \in T_\sigma$ or $R(\beta)(\text{imp}(\alpha_n)) \in T_\sigma$. (relevance constraint)

This definition integrates a number of insights from §1, including sensitivity to speaker intentions (§1.2); reference to (lack of) set-theoretic restriction (§1.2); and sensitivity to discourse function (§1.6). It also

¹⁶ Thanks to Philippe Schlenker for suggesting this terminology.

presupposes that non-restrictive adjectives are uniquely associated with an additional implication not always present in restrictive modification; see §1.6 above. An important tacit constraint in this definition is that the implication triggered by the non-restrictive adjective would not exist if the adjective were not present. In other words, component (3.58b) states that (i) non-restrictive adjectives allow one to infer a coherence relation between two implications of a text; and (ii) that coherence relation could not be inferred without the adjective, because one of the two relata (namely $\mathbf{imp}(\alpha_n)$) comes from the presence of the adjective.

A corresponding definition for *restrictive* could take a number of forms. The choice depends mostly upon how one wishes to classify non-subjective modifiers like *alleged*, *former*, etc. Are they restrictive, non-restrictive, or neither? The answer seems somewhat arbitrary, but my persuasions align with the third option: it is a category mistake to ask whether *former* is restrictive. Given this constraint, *restrictive* cannot be defined simply as not non-restrictive. Instead, a certain relationship should hold between the modifier and the *noun* that it modifies, not the entire DP in which it occurs. Again, we introduce speaker relativity into the definition, which distinguishes 3.59 from Piñón's (2005) definition in 3.3 above.

(3.59) **Definition. restrictive:**

A phrase α that modifies a noun \mathbf{N} is *restrictive* with respect to index i if and only if the speaker believes that $\chi(\llbracket[\alpha\mathbf{N}]\rrbracket^i) \subset \chi(\llbracket[\mathbf{N}]\rrbracket^i)$.¹⁷

Defining *restrictive* in terms of \mathbf{N} instead of DP has the added advantage that it actually predicts (correctly) that all DP-adjoined appositives cannot be restrictive.

3.2.5 Testing the definitions on some basic cases

To evaluate these definitions, with a focus specifically on 3.58, let us first re-examine some basic data. The modifiers in 3.10 are most naturally interpreted non-restrictively in the intuitive sense of the word. Does the definition in 3.58 count them as non-restrictive, as desired?

- 3.10 a. I take care of my sick mother.
 b. Cigarettes contain harmful carcinogens.

Accommodating some typical context in which these sentences might be used, let us begin with *my sick mother*. Does *sick* satisfy (3.58a)? Yes, if an individual utters *my sick mother* when s/he has only one mother, then s/he would believe *my mother* and *my sick mother* to denote the same person. What about (3.58b)? Yes,

¹⁷ Where f 's co-domain is Bool, $\chi(f)$ is the characteristic set of f (the set of elements in f 's domain mapped to true).

in the most natural interpretation of (3.10a), the speaker indicates that his/her mother’s sickness *explains* why s/he takes care of her; alternatively, caring for the mother is a *result* of her sickness. In either case, a coherence relation—either EXPLANATION or RESULT—will link the presupposition of the sentence (*I have a sick mother*) with its assertion (*I take care of her*), via a rhetorical statement paraphraseable as *My mother is sick, and as a RESULT, I take care of her* (see §4:3.92 for a slightly different analysis of this case, including motivation).

- (3.60) a. $\exists!x[\mathbf{sick}(x) \wedge \mathbf{mother}_{me}(x)]$ (presupposition)
 b. $\mathbf{take-care}(x)(me)$ (assertion)
 c. EXPLANATION($\exists!x[\mathbf{sick}(x) \wedge \mathbf{mother}_{me}(x)]$)($\mathbf{take-care}(x)(me)$) (rhetorical link)

Moving on to example (3.10b), does *harmful* in *harmful carcinogens* satisfy constraint (3.58a)? Well, in the kind of scenario in which (3.10b) might be uttered, e.g. an infomercial for a smoking-cessation product, a speaker will typically understand that all carcinogens are harmful, and so uses *harmful* to emphasize this point. So, yes, in such a case, the speaker would recognize that *carcinogen* applies to the same set of objects as does *harmful carcinogen*. What about constraint (3.58b)? This part is less obvious. The implication from *harmful* in *Cigarettes contain harmful carcinogens* is something like “carcinogens are harmful.” Consider again (3.10b) in the context of a smoking cessation ad or parental advice. I think it is fair to say that any context in which (3.10b) would be used is a context in which it has already been stated (or will soon after be stated), explicitly or via implicature, that smoking is bad or that one shouldn’t smoke. If correct, then there is an obvious rhetorical link between *carcinogens are harmful*, *cigarettes contain carcinogens*, and *cigarettes are bad/harmful*, possibly via Asher & Lascarides’s (2003) CONSEQUENCE relation (whose most common use is to relate the antecedent and consequent of a conditional).

It is already evident that “computing” a coherence relation is not an exact science, but I believe the intuition is fairly clear: there are systematic rhetorical/coherence relations that hold between certain pairs of sentences in a discourse, and there are also certain pairs of sentences in any given discourse that are *not* rhetorically linked in this way. A given theory of coherence relation inventory and axioms for each relation can make concrete predictions about what kinds of discourses are well-formed and which aren’t. Component (3.58b) states that the content of a non-restrictive adjective must find some accessible proposition with which it can be rhetorically related; if this fails, then the result will be infelicity (as in *#my tall mother*). Thus the content of a non-restrictive adjective in this framework is much like that of an anaphoric presupposition

trigger (see 3.48 above).

It is worth reflecting briefly on the question of exactly why it is that non-restricting modifiers that lack the relevance property are infelicitous. I think that the answer is both simple and interesting. Modifiers are often characterized as “non-obligatory” elements, at least from a syntactic point of view (see Dowty 2003). So if they are not syntactically obligatory, then in order to be useful, they must be *semantically* obligatory, i.e. their absence/presence should correlate with truth-conditional differences. But by definition a non-restricting modifier is one that does not alter truth-conditional meaning. So an non-restricting modifier is syntactically *and* semantically non-obligatory. The next logical step would be to ask if it is *pragmatically* obligatory, i.e. would the implicatures of a sentence differ depending only upon the presence versus absence of the modifier? If not, then the modifier would fail to be useful from any perspective, and hence should violate the Gricean maxim of Manner. If so—i.e. if a non-restricting modifier is responsible for the generation of some implicature—then we should expect for it to be felicitous: there is a *reason* to utter it. This is my hypothesis for why all non-restrictive modifiers are required to be pragmatically relevant—they don’t contribute in a useful way truth-conditionally, so they are required to do something else useful.

An immediate prediction of this quasi-functional analysis is that the coherence inference from relative clauses (as in Rohde et al.’s (2011) 3.42) should be cancelable, while the coherence inference from non-restrictive adjectives should not. As Rohde et al. (2011) observed, the first statement appears to be true: the first sentence of (3.61a)’s implicature can be felicitously cancelled.

- (3.61) a. John detests the co-workers who are arrogant and rude. Of course their personality has nothing to do with it, they just happen to be the only ones he’s ever met.
- b. #I take care of my sick mother, though her sickness has nothing to do with why.

The felicity of (3.61a) suggests that coherence inferences are conversational implicatures, which typically have the defining property of being defeasible/cancelable. But at the same, the coherence inference in (3.61b) apparently *cannot* be canceled: absent contextual information, the continuation in (3.61b) feels inconsistent with the first clause. This apparent puzzle can be understood, I think, by constructing a scenario in which the EXPLANATION inference is canceled by a distinct contextually introduced relation that in some sense overrides an explanatory inference. Consider 3.62: in this case, VIOLATED EXPECTATION links the two statements, instead of (or perhaps in addition to) the EXPLANATION relation in (3.61b). Hence 3.62 is felicitous because *some* coherence relation can be established between its presupposition and its assertion.

- (3.62) *Context:* A devastating zombie apocalypse is upon us. If anyone is known to be sick, they must be immediately quarantined or killed. But I am just too loyal a guy, so...
I take care of my sick mother.

The point this example makes is that in non-restrictive modification, specific coherence inferences can be cancelled, but *only if* replaced by some other coherence inference. So the contrast between (3.61b) and 3.62 boils down to the notion that 3.62 has a coherent link—albeit not the usual one holding for *my sick mother*—while (3.61b) lacks one.

The coherence-based analysis applies equally well to appositives, which are, again, hallmark cases of non-restrictive modifiers.

- (3.63) a. I take care of my mother, who is sick. b. My mother, who I take care of, is sick.

Interestingly, the felicity of both examples suggests that an appositive can contribute either of the two relations for whatever relation it is that holds between the implications of 3.63 (probably EXPLANATION).

3.3 Revisiting theories of R/NR and the two-domains model

This section reexamines the “restrictive/non-restrictive” Bolinger contrast introduced in Chapter 1 (henceforth R/NR), in light of the observations made in §1 of this chapter as well as the analysis of the terms (*non-*)*restrictive* from §2. After reviewing the basic distributional patterns in English and in Italian (as reported in the literature), I evaluate some predictions of Larson’s and Cinque’s two-domains based theories of R/NR. (*Non-*)*restrictive* is not defined in either analysis, but I show that regardless of exactly how it is defined, both theories fail to account for some body of critical data. Having motivated the need for a new analysis, I then show how basic composition principles, combined with the framework developed in Chapter 2, can (partially) explain when and why attributive adjectives can receive non-restrictive readings (where *non-restrictive* is defined as in 3.58). I also provide an analysis of *how* it is that non-restrictive modifiers generate an appositive-like implication (called **imp**(α) above), which is arguably the most technically puzzling property of such modifiers and also a topic that few theories have explicitly addressed (with Umbach 2006 and Morzycki 2008 being the exceptions; see §1 above). Together, these two components shed light on the general distribution of non-restrictive attributive adjectives cross-linguistically.

3.3.1 Some reasons to suspect that R/NR is of a different nature than Bolinger contrasts

What is remarkable about the well-worn minimal pair in 3.64 is that the relative position of adjective and noun apparently determines whether a non-restrictive interpretation of the adjective is possible. The pair in 3.65, modified from Larson & Marušič 2004, shows that this situation can obtain in DPs with various kinds of determiners, so it is not just peculiar to universals.

- (3.64) a. Every unsuitable word was deleted. (Larson & Marušič 2004:275)
✓*Restrictive*: “Every word that was unsuitable was deleted.”
✓*Nonrestrictive*: “Every word was deleted; they were unsuitable.”
- b. Every word unsuitable was deleted.
✓*Restrictive*, ✗*Nonrestrictive*
- (3.65) a. Most/several/all/∅ blessed people were healed. (✓*Restrictive*, ✓*Non-restrictive*)
b. Most/several/all/∅ people blessed were healed. (✓*Restrictive*, ✗*Non-restrictive*)

This pattern is characteristic of many other Bolinger contrasts in English, more narrowly of the family of contrasts presented in Cinque (2010:Ch1). But at the same time, restrictive/non-restrictive intuitively differs from related Bolinger contrasts in an important and fundamental way: intersective/subsective adjectival modifiers are by default “restrictive,” but have the potential to be “non-restrictive” provided the context and surrounding lexical items are appropriate. Most other Bolinger contrasts involve specific closed classes of lexical items, and the opposing readings tend to be idiosyncratic; for example modal/direct applies to a limited set of modal adjectives (*possible*, *conceivable*, etc.); relative/absolute applies to superlatives only; intersective/adverbial applies (roughly) to only those adjectives that have adverbial derivatives, etc. In languages like French, the different meanings observed in inner versus outer modifiers can be even more radical. Unlike other Bolinger contrasts, it makes sense to ask of any (use of an) attributive adjective—of nearly any modifier at all, in fact—whether it is restrictive or non-restrictive. So while the pattern in 3.64-3.65 shows a clear opposition reminiscent of Bolinger contrasts, there is reason to suspect that the contrast has less to do with lexical semantics or DP-internal semantics, than it does with more general characteristics of the English syntax-semantics mapping.

Given the characterization of (non-)restrictiveness from §2, we are now in a position to ask whether the contrast exemplified by 3.64 is indeed a subspecies of the Bolinger contrast. Even if the answer turns out to

be negative—which I will argue extensively that it does—the R/NR phenomenon still has interesting implications for the architecture of DP, and is still relevant for the inner/outer bifurcation in nominal modifiers. I will also argue that the implications triggered by non-restrictive adjectives in definite/referential versus non-definite/non-referential DPs arise via distinct mechanisms, a possibility alluded to in Chapter 1.

Historically, and for good reason, the kinds of facts illustrated by 3.64 have been taken to constitute strong evidence in favor of the two-domains model of modification. According to Larson (1998,1999); Larson & Marušič (2004); Bernstein (1993); Cinque (1994,2010,2003); Alexiadou et al. (2007), and others, non-restrictive readings of attributive adjectives are possible for inner modifiers only. If this idea is correct, then we are led to a number of clear predictions:

(3.66) **Predictions of the inner-outer theory of non-restrictive modification**

- a. Postnominal adjectives in Romance should be able to receive both non-restrictive and restrictive readings, while prenominal variants should be unambiguously non-restrictive (the mirror-image pattern of English).
- b. Ordering restrictions on iterated prenominal adjectives should exist in English, i.e. there should be no structures of the form [AP₁ AP₂ N] in which AP₂ is restrictive and AP₁ is non-restrictive.
- c. Corresponding restrictions should hold for iterated postnominal adjectives in Romance.
- d. (For Cinque's (2010,2003) version of the two-domains model) there should be no situation in which AP₁ has *any* kind of inner reading while AP₂ is restrictive (in English).

The accuracy of these predictions can only be assessed relative to a definition of (*non-*)*restrictive*, and such predictions have not been systematically explored in the literature. I take up this exploratory task in §4.2.

Initial evidence that the source of the contrast in 3.64 is something more general than the source of other Bolinger contrasts is that the former can be generalized from pre- versus postnominal adjectives to pre- versus postnominal modifiers *simpliciter* (excluding appositives): the restrictive paraphrase of (3.64a) is unambiguous, and differs from (3.64b) only in the presence of an overt relative clause structure surrounding *unsuitable*. In other words, the relative clause in 3.67 patterns just like the postnominal adjective in (3.64b).

(3.67) Every word that was unsuitable was deleted.

✓*Restrictive*, ✗*Nonrestrictive*

Such a generalization cannot always be made for related data: consider the fact that while postnominal *visible* in *the stars visible* is unambiguously stage-level, it is less clear what interpretive possibilities there are for the full relative clause *that are visible* as in *the stars that are visible*. In fact, authors have expressed conflicting judgments about exactly the question of whether *the stars that are visible* has the same (singleton) set of readings as *the stars visible* does. Thus, it could be a general property of the English postnominal *position*, and not just of postnominal *adjectives*, that disallows non-restrictive interpretations.

There are also apparent differences in syntactic distribution between R/NR and other Bolinger contrasts, subtle as they may be. Larson in fact used R/NR to illustrate the pre- versus postnominal contrast in adjective interpretation *only*; for R/NR, he made no claims about prenominal ordering restrictions within iterated AP structures, as he did for modal/direct *possible* and i-level/s-level. Nevertheless, certain predictions about the distribution of non-restrictive adjectives can be deduced from the DP architecture advocated in Larson 1998,1999 (see §4.2).

Cinque 2010, on the other hand, explicitly argues that R/NR fits into the mirror-image schema precisely like the other contrasts he considers. This empirical claim, however, can be called into question by further probing into the data presented in Cinque 2010:Ch1. Specifically, Cinque (2010:18-19) claimed that in English, if two adjectives co-occur prenominal, only the closest one to the noun can be read non-restrictively—exactly the pattern Cinque’s version of the two domains model predicts. Unfortunately, this empirical claim is difficult to evaluate because of two complications: the presence of a focused superlative (*most*) in all examples (superlatives have a syntax and semantics very different from ordinary, positive-form adjectives; see Chapter 4 and references therein), and the lack of any explicit semantic criterion distinguishing “restrictive” from “non-restrictive” interpretation.

- (3.68) *restrictive > nonrestrictive > N > restrictive* (Cinque 2010:19, annotation
and judgments preserved)
- a. his MOST UNSUITABLE unsuitable acts
 - b. *his unsuitable MOST UNSUITABLE acts
 - c. his unsuitable acts MOST UNSUITABLE

Cinque also argues that precisely the expected mirror-image pattern holds in Italian: postnominal adjectives are claimed to be generally ambiguous between restrictive and non-restrictive, and prenominal adjectives are claimed to be “unambiguously non-restrictive.” This is illustrated in 3.69, which constitutes the empirical evidence cited for this generalization about Italian. Italian examples corresponding to 3.68 are also presented

(Cinque 2010:21,ex.54), but they are confounded in the same way as 3.68 was.

- (3.69) a. Le lezioni noiose di Ferri se le ricordano tutti (Italian, Cinque 2010:8)
the classes boring of F. remember all
✓*Restrictive*: ‘Everybody remembers just those classes by Ferri that were boring.’
✓*Nonrestrictive*: ‘Everybody remembers Ferri’s classes, all of which were boring.’
- b. Le noiose lezioni di Ferri se le ricordano tutti
✗*Restrictive*, ✓*Nonrestrictive*

While Cinque’s intuition may be clear, there is an equally clear sense in which his claim contradicts other clear empirical facts about Italian. If prenominal adjectives are “unambiguously non-restrictive,” then this would mean that, e.g. prenominal *invisibili* ‘invisible’ in *invisibili stelle* ‘non-visible stars’ (see 3.52 in Chapter 3) would be non-restrictive, despite the fact that *invisibili stelle* has a “restricted” denotation when compared with *stelle*. This terminological inconsistency yet again highlights the fact that the terms “restrictive” and “non-restrictive” are used in wildly inconsistent ways among linguists.¹⁸

Abstracting away from what terminology ought to be used to label the alternation in 3.64-3.69, Cinque’s (2010) following observation shows that R/NR has at least one quirk in its distribution that sets it apart from other Bolinger contrasts. In other words, the situation in Italian is not so simple:

“For reasons that are not entirely clear, omission of the genitive PP *di Ferri* renders the nonrestrictive reading in [(3.69a)] virtually impossible, and [(3.69b)], for which a restrictive reading is unavailable, virtually ungrammatical” (Cinque 2010:8,fn.4).

The significance of this point is amplified by an exactly parallel observation reported for Spanish by Morzycki (2008:3,fn2): omission of the PP *de María* disambiguates (3.70b), allowing only the restrictive reading as shown in (3.70c).¹⁹

¹⁸ More direct evidence of this inconsistency can be seen by the following disagreement between Alexiadou et al. (2007) and Cinque (2010): “One of the reviewers...[says] ‘it is semantically not possible for an adjective to be non-restrictive and non-intersective at the same time’.” to which Cinque replies “it is certainly true that an adjective can be nonrestrictive and nonintersective at the same time” (119:fn18), citing the following example (illustrative for Italian and for the English translation):

- (i) Ieri abbiamo avuto modo di ammirare una meravigliosa ballerina
‘Yesterday, we had the opportunity to admire a beautiful dancer.’

In the same footnote, Cinque also comments that “Alexiadou et al. (2007) claim that the restrictive/non-restrictive and the intersective/non-intersective distinctions are independent from one another...But this appears dubious[.]”

¹⁹ The judgment in (3.70c) fixes what I understand to be a minor but critical typo in the paper’s description.

- (3.70) a. los sofisticados amigos de María
 the sophisticated friends of Maria
 ✗*Restrictive*: ‘those of Maria’s friends who are sophisticated’
 ✓*Nonrestrictive*: ‘Maria’s friends, all of whom happen to be sophisticated’
- b. los amigos sofisticados de María
 ✓*Restrictive*, ✓*Nonrestrictive* (Spanish, Morzycki 2008:3)
- c. los amigos sofisticados
 ✓*Restrictive*: ‘the friends that are sophisticated’
 ✗*Nonrestrictive*: ‘the friends, who are sophisticated’ (attributed to V. Demonte)

Such examples raise further doubt as to whether R/NR does indeed pattern exactly like other Bolinger contrasts. Is there something crucial that links the presence of a postnominal PP with the interpretation of an attributive adjective?

There are yet additional reasons to think that R/NR differs in some fundamental way from other Bolinger contrasts. For one, the pragmatic relevance constraint introduced in §1.6 appears to have no obvious correlate for any other pre-/postnominal contrast in adjective interpretation. Furthermore, a phenomenon corresponding to the adjectival R/NR contrast occurs in the verbal domain as well.²⁰ It is well-known that there are interpretive asymmetries between pre- and postverbal adverbs (e.g. see Ernst 2002,2007; Cinque 1999), and one such asymmetry mimics the R/NR contrast in adjectival interpretation. In 3.71, preverbal *slowly* is compatible with something like a non-restrictive reading relative to the event predicate denoted by the verb phrase (roughly the set of (ship-)sinking events). In postverbal position, the most prominent reading of *slowly* is not just slowly *simpliciter*, but slowly *relative to the typical ship-sinking* 3.72. The truth-conditional difference can be coaxed out with the following scenario: suppose ships typically take 24 hours to sink, and that our particular ship took only 18. 18 hours is a pretty long time to be in a sinking ship, so in some sense 3.71 feels true (or at least has a reading on which it’s true). But at the same time, since 18 hours is quite short *for a boat-sinking*, 3.72 seems false (or has a reading on which it’s false).

(3.71) The ship slowly sank.

(3.72) The ship sank slowly.

Morzycki (2008) uses the following pair to illustrate the contrast: only (3.73a) seems to be compatible with the appositive paraphrase, which he takes to be indicative of a non-restrictive interpretation.

²⁰ The adverbial contrast was observed in a different setting by Shaer (2000), but argued to exemplify a parallel to 3.64 by Morzycki (2008). See also Ernst 2002,2007 for discussion of related phenomena.

- (3.73) a. The Titanic(’s) rapidly sinking caused great loss of life.
Restrictive: “The Titanic’s sinking being rapid caused great loss of life.”
Nonrestrictive: “The Titanic’s sinking, which was rapid, caused great loss of life.”
- b. The Titanic(’s) sinking rapidly caused great loss of life.
Restrictive, #Nonrestrictive

Just like within the nominal domain, the putative non-restrictive readings of preverbal adverbs generate an additional implication, namely that the VP has the property denoted by the Adv, just like *harmful carcinogens* implies that the N *carcinogens* has the property denoted by the Adj *harmful*. This can be seen in 3.74-3.75: 3.74 implies that taking part in insider trading (the VP-event predicate) is illegal (the property denoted by the adverb), while 3.75 merely implies that John’s particular action was illegal—it is compatible with there being legal ways to take part in insider trading.

- (3.74) John illegally took part in insider trading. (3.75) John took part in insider trading illegally.

The parallel between the adverbial pattern in 3.71-3.75 and the adjectival pattern in 3.64-3.67 is unexpected if the R/NR alternation in adjectival interpretation is of the same species as s-level/i-level, direct/implicit relative, and other Bolinger contrasts, none of which have clear verbal correlates. This suggests that the R/NR alternation may be logically independent from specific facts about the nominal, contrary to previous analyses. If correct, then alternative explanations must be sought.

3.3.2 Two-domains approaches to the data, and where they go wrong

Some approaches to the notion of non-restrictiveness were sketched in §1. Here I briefly recap from Chapter 1:§1 Cinque’s two-domains based theory of R/NR—the structure-sensitive contrast widely assumed to track restrictiveness. But first, some general remarks are in order concerning how the data should be best described in terms of a two-domains model.

Because Larson/Cinque frame the discussion of R/NR in terms of related Bolinger contrasts, it is tempting to describe the situation as follows: “non-restrictive adjectives” are inner/direct modifiers, and “restrictive adjectives” are outer/indirect modifiers. But stating the position in this way is actually inconsistent with general traits of Bolinger contrasts. Here’s why: individual-level and stage-level readings are both “restrictive” in the intuitive sense of the word, for example, even though the former are regarded as inner and the latter outer modifiers (see Alexiadou et al. 2007 for some complementary remarks). What’s more, an adject-

tive can be both non-restrictive and individual-level (as in *kind* in *my kind mother*), but also non-restrictive and stage-level (as in *sick* in *my sick mother*). So the restrictive/non-restrictive distinction just cannot be an instance of the more general inner-outer bifurcation. Larson’s point as I understand it is simply that non-restrictive readings of adjectives are possible—but crucially not guaranteed—only in certain syntactic positions, and that modifiers not in such a designated position can never be read non-restrictively. As noted in the discussion of 3.68 above, the position that Cinque (2010,2003) advocates is considerably stronger.

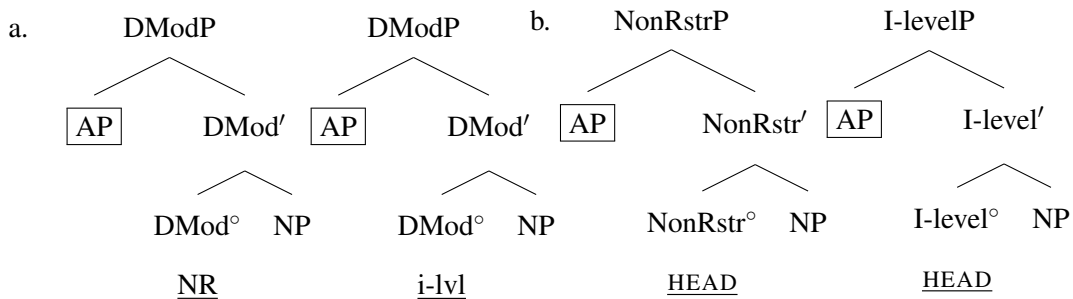
Since many inner readings of adjectives are restrictive readings, a more accurate characterization of the data might be that for any adjective α , if α in $[_{DP} \dots \alpha \dots N \dots]$ or $[_{DP} \dots N \dots \alpha \dots]$ is ambiguous between restrictive and non-restrictive interpretations, then α is non-restrictive if and only if α occupies an inner position (and restrictive iff α is in an outer position). Relativization to ambiguous noun phrases has clear motivation for other Bolinger contrasts, most notably i-level/s-level: some adjectives are i- or s-level in virtue of their lexical semantics, and so will maintain this property regardless of syntactic position (without coercion). It is only with adjectives that can be interpreted more easily in either sense, e.g. *visible*, that the i-level/s-level contrast can be clearly observed. Nevertheless, if restriction to ambiguous nominals is undesirable, an alternative form of the two-domains description could be that for any adjective α , if α in $[_{DP} \dots \alpha \dots N \dots]$ or $[_{DP} \dots N \dots \alpha \dots]$ has a non-restrictive interpretation, then α occupies an inner position.

Recall that Larson analyzed the semantics of the inner-outer distinction in terms of two silent event quantifiers with fixed positions in DP (Γ /Gen for inner, \exists for outer). Put succinctly, “if we think of genericity as a matter of being bound by a covert generic operator binding events/situations, we can view the positional facts...as matters of scope.” (Larson 1999:Lecture1) Such an analysis would appear to work well for non-restrictive adjectives that attribute properties to kinds (e.g. *harmful toxins*), given that certain varieties of kind-reference in language are traditionally analyzed in terms of generic quantification (Carlson 1977b; Krifka et al. 1995; Dayal 2004; a.o.). However, many non-restrictive modification structures have nothing to do with kind reference; the clearest cases being with referential DPs. For example *sick* and *Texan* attribute properties to individuals in *my sick mother* and *the Texan president*, not to kinds. Further, *sick* is a lexically stage-level predicate, so if it were to be bound by one of the operators, it would have to be the outer \exists , and *not* the inner Γ . This latter example is a straightforward counterexample to the idea that non-restrictive adjectives must always be bound by Γ ; by extension, it is also an argument that non-restrictive modifiers need not always be located in the inner modification domain. In other words, if R/NR is to be analyzed in terms of existential versus generic quantification (i.e. in terms of \exists versus Γ), then the case of *sick* leads to

contradiction: *sick* cannot be in the the scope of Γ because it is not generic, nor can it be structurally above Γ , because in that case it should have an outer—and hence restrictive—interpretation.

Moving on to a different approach to the data, we now assess the analysis of R/NR within Cinque’s (2010,2003) cartographic version of the two-domains model. Recall from Chapter 1 that this theory postulates abstract functional heads as introducers of direct/inner adjectives into the nominal structure. Adjectival modifiers occupy the specifier positions of such projections. Cinque does not provide a category label (or a semantics) for these functional heads, so it is unclear whether all direct modifiers should be introduced by heads of the same category, or whether each individual inner reading is associated with a specific functional category. These two options generate divergent structures for different inner readings. The former is illustrated in (3.76a); the latter, in 3.76b (here NR and i-lvl are the silent functional heads that I assume Cinque assumes are responsible for non-restrictive and i-level readings, respectively). The latter has a certain resemblance to Morzycki’s (2005) theory of mediated modification; and the former, to Rubin’s (2002) theory of modification via the silent functional projection ModP.

(3.76) **Two ways of understanding Cinque’s analysis of inner adjectives**



Cinque’s treatment of outer adjectives, on the other hand, is that they are reduced relative clauses, which in English can appear either prenominally or postnominally in surface order. Note that the NPs in 3.77 below must contain all direct APs, i.e. no D-ModP can appear above the CP_{rel} s.²¹

(3.77) **Cinque’s analysis of outer adjectives**



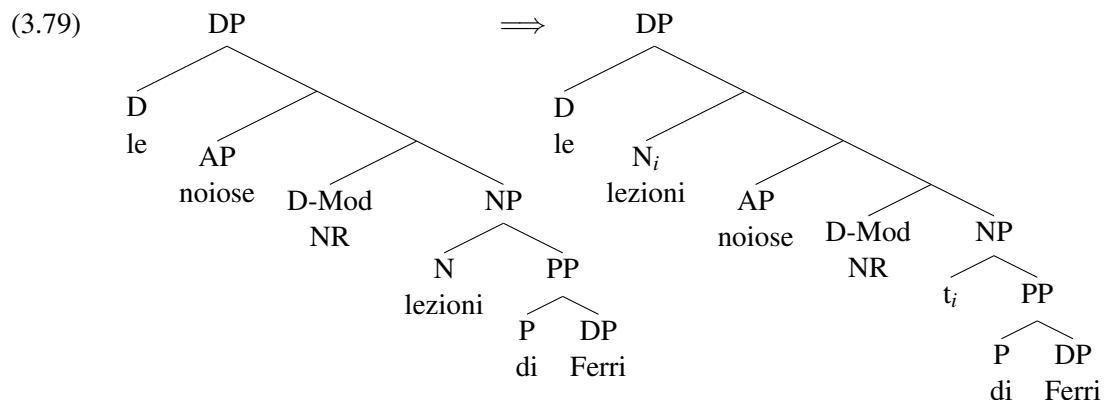
²¹ I suppress the finer-grained (and mysterious) points of his syntactic analysis, such as ordering restrictions between participial and “bare” reduced relative APs, and the assumption that relative clauses are always “merged” prenominally, with postnominal order being derived by noun movement.

So Cinque's theory of R/NR in English is then that non-restrictive adjectives are direct modifiers, and therefore appear in the lower modification domain, in the specifiers of functional projections. These projections appear only preominally (in their base positions, at least), and are lower than all reduced relative APs. This essentially derives the correct results for English.

Concerning Italian, Cinque's (2010) movement analysis does not avoid overgeneration for cases similar to 3.69 above, which I will now show. Recall that Cinque (2010:18-21) reports the mirror-image ordering of readings for Italian: a non-restrictive adjective can appear preominally, and also postnominally provided that no indirect modifier intervenes between it and the noun. This is accounted for in Cinque's (2010) theory by the series of snowball-style movements summarized in Chapter 1 above. While the details are not crucial here, suffice it to say that the basic pattern is captured, assuming there is a reason for these movements to occur in the first place. However, recall the point that I raised about the crucial Italian data in 3.69: when a non-restrictive postnominal adjective is followed by a noun's complement PP, deletion of that complement PP blocks the non-restrictive reading. So the alternation that Cinque describes is:


- (3.78) a. Le lezioni noiose di Ferri se le ricordano tutti
 the classes boring of F. remember all
✓Restrictive, ✓Nonrestrictive
- b. Le lezioni noiose se le ricordano tutti
✓Restrictive, ✗Nonrestrictive

And even assuming the snowball movements, there is no way to block the non-restrictive interpretation (= direct modifier status) in Cinque's theory. To see this, consider first the following derivation for (3.78a), the variant with the PP complement. It involves raising N over all direct modification adjectives.²²



²² Cinque (2010) does not give an explicit syntax for complementation, but the one I have constructed here is intended to be as standard/uncontroversial as possible.

Now consider the same derivation *without* the PP *di Ferri* ‘of Ferri’ (i.e. (3.78b)). Nothing is structurally different about it, so it is safe to assume that the movements Cinque postulates should apply in this case as well. If anything, the movements should be shorter because NP contains less internal structure than in the previous derivation.



In other words, Cinque’s (2010) system generates a structure that assigns a non-restrictive reading to *noiose* ‘boring’ in *le lezioni noiose* ‘the lectures boring’, which contradicts the intuition expressed on page 49 (about non-restrictive readings being blocked by the omission of the genitive PP). The only structure that should be possible given this intuition is one in which *noiose* ‘boring’ is a reduced relative clause adjoined above NP, and is therefore predicted to be semantically “restrictive.”

Having shown that the analysis of Cinque (2010,2003) encounters an empirical difficulty, I will now show that his description of the original data, scrutinized via explicit definitions of the notions involved, does not accurately describe the distribution of readings. Or, more precisely, the contrast he describes does not correspond to any of the notions reviewed in §2.

The schemata below represent Cinque’s (2010,2003) generalizations about restrictive/non-restrictive readings of attributive adjectives in English


(3.81) **Distribution of R/NR adjectives in English (from Cinque 2010,2003)**

- a. *Either reading prenominally, only restrictive postnominally (see 3.64)*
 - (i) [DP D AP_{R/NR} N]
 - (ii) [DP D N AP_{R/*NR}]
- b. *Restrictive precedes nonrestrictive prenominally (see 3.68)*
 - (iii) *[DP D AP_{NR} AP_R N]
 - (iv) [DP D AP_R AP_{NR} N]
 - (v) [DP D AP_{NR} N AP_R]

Patterns (i-ii) state that prenominal adjectives can be read restrictively or non-restrictively, while postnominal adjectives are always restrictive. Patterns (iii)-(iv) jointly state that when a noun is modified by two prenominal adjectives, only the inner one can be NR. Pattern (v) says that a prenominal adjective can be non-restrictive while a co-occurring postnominal adjective is restrictive.

To assess Cinque’s description of the facts, we will need to settle on the crucial definitions. It is unlikely that by *(non-)restrictive* Cinque had in mind the coherence-based notions from §2; more likely is that he intended *non-restrictive* to be synonymous with the notion of *non-restricting* defined in 3.57 above, i.e. that a modifier is non-restrictive if its combination with a noun denotes the same set as the unmodified noun does. Under this conception, it is actually straightforward to show that the proposed descriptive generalizations in (3.81b) are either inaccurate or else are actually due to factors completely orthogonal to “non-restrictiveness.”

Consider first (iii)-(iv). These generalizations are inaccurate if “non-restrictive” is identified with “non-restricting” (which is my best guess at what was actually intended): nothing prevents a situation in which for a structure of the form $[A_1 A_2 N]$, A_2 restrictively modifies N but A_1 non-restrictively modifies $[A_2 N]$. Here is an example of such a derivation, with toy denotations:




Given world knowledge, all nuclear weapons are deadly, so every nuclear weapon is a deadly nuclear weapon and vice versa. Hence *deadly* can non-restrictively modify *nuclear weapon*. But not every weapon is a nuclear weapon, so *nuclear* does not non-restrictively modify *weapon*.

In other words, if Cinque’s (2010) intended characterization is that an adjective A is non-restrictive with respect to a noun N iff $\llbracket [\dots A \dots N] \rrbracket = \llbracket [\dots \dots N] \rrbracket$, then (iii)-(iv) just follow immediately from the nature of bottom-up semantic composition. To see why at a more general level, suppose we have a structure $[A_1 [A_2 N]]$ in which A_1 and A_2 are intersective or subsective adjectives and N is a common noun whose extension is $\{a,b,c\}$. If A_2 “restricts” N , then the denotation of $[A_2 N]$ is a proper subset of $\{a,b,c\}$. Since A_1 is


intersective or subsective, it cannot add elements back into the denotation of $[A_2 N]$, and hence it just cannot be the case that $\llbracket [A_1 [A_2 N]] \rrbracket = \llbracket N \rrbracket$.

(3.83) $\{a, b\}$ or $\{a\}$ or $\{b\}$ or \emptyset , but not $\{a, b, c\}$



By the same reasoning, there is nothing that prevents A_2 from combining with N non-restrictively (provided the context is right) and then A_1 combining with $[A_2 N]$ restrictively. And there should be no problem with both A_2 and A_1 being non-restrictive or restrictive (in fact, an apparent tendency for multiple adjectives to be uniformly restrictive or non-restrictive was pointed out by Morzycki (2008)).

Point (v) says that a prenominal adjective can be non-restrictive while a postnominal adjective is restrictive. This follows from the plausible assumption that at least in many cases prenominal adjectives form constituents with nouns to the exclusion of postnominal modifiers (as in 3.84a), a claim argued for convincingly by Sadler & Arnold (1994). Nothing obviously rules out 3.84b as a possible syntactic structure, but in this case AP_1 could be NR only if AP_2 were, and as noted repeatedly throughout the chapter, postnominal modifiers in English lack non-restrictive readings.



Points (i)-(ii)—the generalizations about possible interpretations of single attributive adjectives—yet again prove the most puzzling. Why is it that in English a postnominal adjective cannot be interpreted non-restrictively while a prenominal adjective can? In the following section I sketch an approach that requires some non-standard syntactic assumptions as one possible answer to this question.

To conclude discussion of the cartographic approach to R/NR, the theory cannot derive the alternation in 3.69, in which the absence of a noun complement PP appears to block the non-restrictive reading of a postnominal adjective in Italian. And the distribution of R/NR readings—under the “proper-subset” defini-

tion of *non-restrictive*—is actually just a consequence of the semantic architecture, namely how bottom-up composition works. The theory does guarantee the prenominal-only distribution for non-restrictive adjectives in English (assuming an appropriate semantics for the NR functional head is provided; see Leffel 2011 for an idea), but does so in virtue of introducing a functional projection that exclusively hosts non-restrictive adjectives, and assigning it a fixed position within DP.

3.4 A new syntax-semantics for non-restrictive adjectives

The goal of the present analysis is to capture the syntax and semantics of non-restrictive adjectival modification, while whenever possible avoiding the introduction of otherwise unmotivated grammatical enrichments. I will argue here for two independently motivated enrichments, which have less severe overall implications for the grammar than those of Larson’s/Cinque’s theories. The two enrichments I adopt are Larson’s (1991) syntactic theory of DP shells, and a novel type-clash repair mechanism. Both are motivated by considerations independent from non-restrictive adjectives.

This analysis focuses on the syntax and semantics of non-restrictive adjectives; I assume that their contribution to discourse coherence is simply a byproduct of an interaction between competing constraints: if an adjective restricts, then it has truth-conditional import (and hence is “useful”). If an adjective fails to restrict, then either it introduces an implication that establishes a coherent link with some other implication (and hence is a useful and felicitous non-restrictive adjective); or else it is redundant from any perspective (and hence is non-restricting but infelicitous, in virtue of violating some Gricean maxim of Manner).

A clear obstacle stands in the way of a completely straightforward compositional analysis of non-restrictive adjectival modification: as discussed in §2.4 above, non-restrictive adjectives often appear to have semantic scope over DPs containing them. For example, the nominal *my sick mother* refers to my mother, and *sick* predicates a property of my mother.

3.27

```

graph TD
    DP --> D[D]
    DP --> NP1[NP]
    D --- my[my]
    NP1 --> AP[AP]
    NP1 --> NP2[NP]
    AP --- sick[sick]
    NP2 --- mother[mother]
    
```

\implies *how??*

- DP has semantic type e (or $\langle\langle e, t \rangle, t \rangle$)
- \llbracket DP \rrbracket “entails” $\mathbf{sick}(tx[\mathbf{mother}_{me}(x)])$, and satisfies the presupp. that I have exactly one mother.

As noted by Morzycki (2008), the situation is similar in *every unsuitable word*, which lacks the kind of presupposition that possessives and definite descriptions are associated with: this nominal truth-conditionally

means the same as *every word*, and *unsuitable* attributes a property to the collection of all such words. Similarly for apparently kind-related modifiers such as *harmful* in *every harmful toxin*: how can *harmful* and *toxin* compose in such a way that $\mathbf{harmful}(\cap(\text{PL}(\text{toxin})))$ (“toxins as a kind are harmful”) is entailed, while *harmful* remains essentially semantically transparent in all composition above this level? The lexical content of the elements in DP do appear to be relevant, e.g. an *every*-headed DP containing a non-restrictive adjective should not be capable of predicating a property of a single individual, whereas a possessive should.

The strategy I will use to overcome this problem is to be serious about the analogy between non-restrictive adjectives and appositives. Note that the problem described in 3.27 also arises for appositives: in *John, who I saw*, we have an $\langle e, t \rangle$ -type modifier, *who I saw*, and an *e* type DP, *John*. Given ordinary semantic assumptions, this should result in an expression of type *t*. But in reality, we find that while an implication of type *t* is generated from this structure, the syntactic result is still a DP, and this DP has the same type and reference as does the unmodified version.

It is worth mentioning that Potts (2002) has offered a solution to a related problem, which builds the heavy lifting into the lexical semantics of *which*. He proposes that appositives modifying full clauses (e.g. *Ali is smart, which she knows*) denote sets of nominalized propositions, and that *which* converts such sets into partial identity functions on the set of all individuals. This function will be defined iff the clause that hosts the appositive has the property denoted by the appositive.

$$(3.85) \quad \mathbf{which} = \lambda f \in D_{\langle e, t \rangle} [\lambda x^p \in D_e : Q(x) \text{ is true } [x^p]]$$

While this analysis was constructed for clausal-scope appositives, the analysis could easily be extended to DP-scope appositives. This would derive the appositive implication while maintaining the denotation of the unmodified DP, but there is a sense in which it is slightly unsatisfying: the lexical semantics of *which* would have to be quite different when *which* occurs in a restrictive relative clause. Thus, on an extended version of Potts’s (2002) theory, *which* would have to be lexically ambiguous.

I propose to maintain a uniform, unambiguous analysis of *which*, and derive the appositive implication in a more general way. As in Heim & Kratzer’s (1998) presentation, *which* turns a *t*-type expression into a property-denoting expression via abstraction over the gap in the relative clause. Concerning how the appositive implication is introduced, I posit the following principle, which can be understood as a certain way of encoding the intuition that appositive content is in a way independent from truth-conditional content:

(3.86) **Type mismatch adjustment principle (TMAP):**

Let $\alpha : \langle \sigma, t \rangle$, $\beta : \sigma$, and assume XP is not a clausal syntactic category.

Then $\llbracket [_{XP} \alpha \beta] \rrbracket = \lambda x_{\sigma} [: \alpha(\beta)] . [\alpha]$.

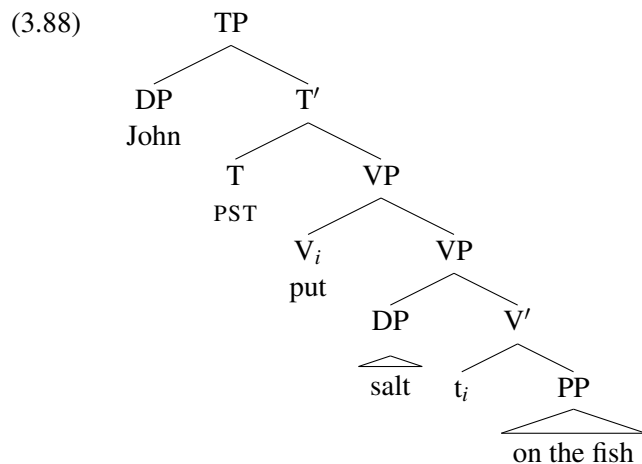
This is a general principle that applies to appositives as well as certain non-appositive relative structures that have proved difficult to analyze with standard semantic tools. For example, syntactically integrated relatives in Mandarin Chinese can apply to proper names (del Gobbo 2003). With **TMAP**, this is analyzed as follows: the constituent *xihuan yinyue de Zhangsan* ‘music-liking Zhangsan’ in 3.87 is defined iff Zhangsan likes music, and if defined refers to Zhangsan.

(3.87) Xihuan yinyue de Zhangsan changchang qu yinyuehui ma?

like music REL Z. often go concert Q
lit. “Does music-liking Zhangsan often go to concerts?”


(del Gobbo 2003:143)

I will show now that the introduction of **TMAP** also solves the predication puzzle about non-restrictive adjectives—that is, provided one more crucial syntactic assumption. The crucial assumption is that some DPs have an extended “shell” structure, as proposed by Larson (1991). This shell structure is parallel to the highly influential VP-shell analysis of Larson (1988). According to the latter, ditransitive verbs originate in a low position, in the complement of which the indirect object is introduced; the verb subsequently raises to its higher surface position.



Basically, the theory of DP shells as explicated in Larson 1991 argues that certain “triadic determiners” such as *every...except...* have a structure that mirrors that of verbs which take three arguments, such as *put*. In this analogy, the argument in the exceptive phrase corresponds to the indirect object *the fish* in 3.88, and the NP argument corresponds to the direct object *salt*. The structure also extends directly to certain kinds of

comparative structures with *more*, as shown here.



While this analysis works for (putative) complex determiners, it would appear to introduce structure superfluous to the analysis of ordinary nominals of the form $[_{DP} D [_{NP} \dots]]$. If we assume that the grammar is capable of generating nominals of both of the following forms, the one in (3.90a) would appear—by considerations of simplicity—to be superior.

- (3.90) a. $[_{DP} D [_{NP} \dots]]$ b. $[_{DP} D_i [_{DP} t_i [_{NP} \dots]]]$

But in the case of *modified* DPs, the shell analysis has the potential to offer a wider range of analytical possibilities than offered by simpler or more conventional alternatives. Recall for example the problem we encountered in the analysis of *my sick mother*—the adjective *sick* semantically says something about my mother, but under standard syntactic assumptions *my mother* is nowhere a constituent in this nominal. Applying Larson’s (1991) hypothesis, though, we have one piece of the puzzle: provided we assume semantic reconstruction is possible for elements of category D, then (3.91a)’s D in its base-generated position forms a constituent with NP, which can then be composed with *sick* to form DP_2 . It follows that DP_2 has semantic type t , as shown in (3.91b).

- (3.91) $[_{DP_1} [_{D_i} my] [_{DP_2} [_{AP} sick] [_{D'} t_i [_{NP} mother]]]]$
- a. $[[DP_2]] = \mathbf{sick}(tx[\mathbf{mother}_{me}(x)]) : t$
- b. $[[D']] = tx[\mathbf{mother}_{me}(x)] : e$


What we actually need is for *sick* to compose with *my mother* to form a semantic unit which refers to *my mother* and which introduces the implication that my mother is sick. The structure in 3.91 gives us the latter, but not the former. The former, however, follows automatically from **TMAP** (3.86): here α is the AP *sick*, σ is e , and β is the D' *my mother*:

- (3.92) a. $\llbracket D' \rrbracket = \iota x[\mathbf{mother}_{me}(x)]$
 b. $\llbracket AP \rrbracket = \lambda x[\mathbf{sick}(x)]$
 c. $\mathbf{TMAP}(\llbracket D' \rrbracket)(\llbracket AP \rrbracket) = \mathbf{TMAP}(\iota x[\mathbf{mother}_{me}(x)])(\lambda x[\mathbf{sick}(x)])$
 $= \underline{[: \mathbf{sick}(\iota x[\mathbf{mother}_{me}(x)])]}.[\iota x[\mathbf{mother}_{me}(x)]]}$

This is the desired result: *my sick mother* refers to my mother, and presupposes that I have a mother (the presupposition of *my*), and that my mother is sick. The fact that the implication from *sick* is encoded as a presupposition is not of particular importance here. There are almost certainly pragmatic/discourse-level differences between this implication and bona fide, lexically triggered presuppositions, but such differences are not our focus here: the important thing is that the truth-conditional and non-assertional components of the meaning of *my sick mother* are adequately separated.

The analysis in 3.92 has interesting consequences for postnominal modification. On the DP-shell analysis, the difference between pre- and postnominal outer modification is a matter of the relative positions of the NP and the modifier, as shown below.

(3.93) **Outer modification structures with DP shells**



Now consider what happens if we attempt to derive a non-restrictive interpretation of *my mother who is sick*, where the relative clause is integrated, not appositive. If D in 3.94 semantically reconstructs, we get the incorrect meaning that DP refers to the unique entity that stands in a possession relation to me and that is sick (via a derivation parallel to that of 3.92). But if D fails to reconstruct, then the noun and modifier must combine via intersection and the determiner takes the modified noun as argument. This derives an intersective semantics for the relative clause, and crucially nowhere in the derivation is **TMAP** invoked. In other words, DP shells plus **TMAP** jointly predict that there should be no non-restrictive construal of integrated postnominal modifiers, including adjectives or relative clauses like *who is sick* in 3.94.²³

²³ The treatment of relational nouns, in this case *mother*, is sloppy in this example. Strictly speaking we need to spell out the mechanism by which **my** saturates the possessor parameter on **mother**. This could be done in a number of ways, none of which would crucially change the analysis here.

(3.94) $[[DP_1 [D_i \text{ my}] [DP_2 [NP \text{ mother}] [D' t_i [CP \text{ who is sick}]]]]]$

a. $[[DP_2]] = \lambda x[\mathbf{mother}(x) \wedge \mathbf{sick}(x)]$

b. $[[DP_1]] = \iota x[\mathbf{mother}_{me}(x) \wedge \mathbf{sick}(x)]$

We now have a principled explanation for why *who is sick* in the DP *my mother who is sick* cannot be read non-restrictively. Of course, we still need an explanation for why this DP is either infelicitous or else has the implicature that I have more than one mother. The explanation is presumably related to the relevance constraint and the additional implication triggered by non-restrictive adjectives.

There is another, more surprising prediction of the analysis in 3.94: in referential nominals at least, it is *outer* pronominal modifiers that can give rise to non-restrictive readings. This entails that R/NR is logically independent from the inner-outer structural distinction among attributive adjectives. This is a novel perspective, which departs significantly from the analyses of Cinque (2010) and Larson (1998); Larson & Marušič (2004), and others mentioned above. The situation is different for non-referential nominals, though. If we assume that DP-shells are projected only in the presence of outer modification (also three-place determiners, etc.), then we can maintain a simple head-adjunction structure for inner modifiers. Consider the case of *many harmful toxins*, as in *cigarettes contain many harmful toxins*, which would have the structure in 3.95 according to the syntax of Chapter 3.

(3.95) $[[DP [D \text{ many}] [NP \phi [N' [N_1 [A \text{ harmful}] [N_2 \text{ toxins}]]]]]]]$

In this case, **TMAP** does the legwork, with some help from the kind-referring semantics for nominal roots. According to Chapter 2, the noun *toxins* denotes a predicate of toxin subkinds, and has type $\langle\langle s, e \rangle, t\rangle$. And *harmful* too can be a kind-level predicate given the assumptions of Chapter 2:§5, also with semantic type $\langle\langle s, e \rangle, t\rangle$. This allows for an intersective interpretation of *harmful*. But what we actually need is something else, namely for the Ty_2 constant **toxin** to denote a kind (type $\langle s, e \rangle$). Suppose for the sake of argument that this is possible. Then the correct result would fall out via **TMAP**:

(3.96) $[[N_1]] = \mathbf{TMAP}([A])([N_2]) = \mathbf{TMAP}(\lambda_{x\langle s, e \rangle}[\mathbf{harmful}(x)])(\lambda_w[\iota[*\mathbf{toxin}_w]])$
 $= [:\mathbf{harmful}(\lambda_w[\iota[*\mathbf{toxin}_w]])].[\lambda_w[\iota[*\mathbf{toxin}_w]]]$

The referent of the constituent *harmful toxins* is just the kind “toxins,” but **TMAP** additionally generates, as a side-effect of the composition, the implication that toxins as a kind are harmful. This appears to be just

what we want for kind-related non-restrictive adjectives. A similar analysis could be constructed for *every unsuitable word* and related examples; the only difference would be minor adjustments in the composition required by the singular number on the noun.


To account for the patterns in Italian, I appeal to two of the binary parameters introduced in Chapter 2:

(3.97) **Parametric difference between English and Italian DPs**

- a. In English, head-adjunction is to the left only.
- b. In Italian, phrasal adjunction is to the right only.

The statements in 3.97 entail that English and Italian have three possible modification structures for non-appositives, but that the three structures differ in a systematic way. For English we have:

(3.98) **Possible modification structures for English**



In 3.98, a. represents prenominal outer adjectives; b. represents postnominal (outer) modifiers; and c. represents prenominal inner adjectives. Structure d. is ruled out because it violates (3.97a). Note that a non-restrictive reading of a modifier is possible in structure a. for definites and possessives, and in structure d. for all other DPs. Thus, as conjectured above, the R/NR alternation does *not* have the same inner/outer distribution as other Bolinger contrasts. In English at least, the contrast is really only between pre- and postnominal modifiers.

In Italian, we have exactly the mirror image of possibilities: prenominal outer adjectives (a.) are out via (3.97b); but postnominal inner adjectives are acceptable since (3.97a) applies to English only.

(3.99) **Possible modification structures for Italian**

- ✗prenom, outer
- ✓postnom, outer
- ✓postnom, inner
- ✓prenom, inner

Unfortunately, the analysis of *my sick mother* in 3.68 does not extend directly to similar examples in Italian. 3.68 appealed directly to the existence DP-shells, which I have proposed are not available in Italian prenom-

inal modification structures. But excluding definites/possessives, the present analysis can be further spelled out so as to capture Cinque's description of the basic pattern in Italian.

3.4.1 Concluding remarks on R/NR and its relation to the two-domains model

In this section I have argued that existing two-domains theories of modification in DP do not fully account for the distribution of non-restrictive adjectives in English or in Italian (as reported in the literature). I have presented a new theory of non-restrictive adjectival modification that has the following features: the implication from non-restrictive modifiers is captured via a type-mismatch repair principle **TMAP**, a placeholder operation that composes appositive meaning; this implication must establish a coherence relation with some other implication in the sentence or in the discourse; and a constrained postulation of DP shells in the syntax guarantees that pronominal adjectives in English can be interpreted non-restrictively in referential nominals. This approach allowed us to formulate analyses of the data that were problematic for Cinque's (2010) cartographic analysis of non-restrictive modification.

3.5 Discussion and conclusion

In this chapter I have argued that non-restrictiveness is inherently related to discourse role. In particular, what it means for a modifier to be non-restrictive is that it does not affect truth conditions (modulo presupposition), and also that it establishes an implication that is discourse-related to some other salient proposition in the sentence/text. This analysis has been shown to provide a fresh perspective on the R/NR Bolinger contrast, and has straightforward applications to the semantics/pragmatics of appositive modifiers and a Chinese relativization construction. See Chapter 5:§1 for an extension of this analysis to modified proper names.

3.5.1 Psycholinguistic connections

The idea that non-restrictive modifiers characteristically introduce coherence relations—and therefore recruit pragmatic reasoning—receives indirect support from a recent neurolinguistic study on the on-line processing of simple modified and unmodified noun phrases like *the (fat) chicken*. Granting certain assumptions about the functional role of the anterior temporal lobes in language comprehension, the results of Leffel et al. (2014) suggest that the processing of non-restrictive modification may involve pragmatic computations not present in restrictive composition (see also Bemis & Pytkäinen 2011 for foundational research on locating linguistic combinatory operations in the brain using magnetoencephalography (MEG)).

The more general idea that nominal modifiers should be—or are generally expected to be—informative

or just generally useful makes a great deal of conceptual and intuitive sense: according to the Gricean maxim of quantity, it is a convention of language use that one should provide as much information as is necessary to achieve the communicative goals at hand, and no more than is necessary. Thus, the use of a modifier that is both truth-conditionally and pragmatically irrelevant constitutes a violation of quantity in virtue of providing more information than is necessary. Such uses can be seen as violations of manner, since an equivalent message could be conveyed with a simpler utterance (see Chapter 5:§3).

Beyond commonsense conceptual motivation, there are empirical and behavioral reasons to think that language users disprefer or find infelicitous unjustified non-restrictive modifiers. The study of on-line semantic processing—particularly within the visual world paradigm (see Huettig et al. 2011 for a survey)—has amassed strong evidence that adjectival modifiers have a default restrictive (or “contrastive”) interpretation in simple noun phrases like *a/the blue pen*. In a research program initiated by the seminal studies of Tanenhaus et al. (1995) and Sedivy et al. (1999), psycholinguists have learned, via tracking eye-movements to objects in a visual display, that subjects identify the referent of a modified noun phrase like *the blue pen* more quickly when there are multiple pens in the display than they do in the presence of a single pen. This finding is naturally interpreted as demonstrating a general expectation that modifiers will be restrictive. With respect to the kinds of stimuli used by Sedivy et al. (1999), one conclusion can be summarized as follows (where the linear order of the noun *N* and the modifier *A* need not be as indicated).

- (3.100) **The contrastive function of modifiers:** The semantic processing system defeasibly infers from a phrase [*the* [*A N*]] that within the discourse context,
- a. exactly one object satisfies both *A* and *N*; and (**uniqueness presupposition of *the***)
 - b. at least one object satisfies *N* but not *A*. (**contrastive/restrictive expectation**)

In a similar paradigm, Engelhardt et al. (2006) showed that subjects exhibit behavioral signs of “confusion” when presented with noun phrases containing referentially irrelevant modifiers.

These and a host of related findings support the hypothesis that nominal modifiers are expected to convey information that is relevant for current conversational purposes. In visual world experiments, the purpose of a modifier is usually to aid in referent identification, a truth-conditionally relevant process. But as argued throughout this chapter, truth-conditional import is not the only factor that can justify the use of a modifier: it is often enough for a modifier to provide or highlight a rhetorical link between implications of a discourse.

3.5.2 A constrained theory of coherence relations for non-restrictive modification?

If non-restrictive modifiers can be justified rhetorically, a natural question to ask is exactly *what range* of rhetorical links can justify their use. One way to approach the question is to think of discourse coherence relations as theoretical proxies for “justifications.” For example it is EXPLANATION that justifies *sick* in *I take care of my sick mother*, but VIOLATED EXPECTATION that justifies *elderly* in *My elderly mother is still quite active*.

Identifying the range of coherence relations that can justify non-restrictive modifiers would be informative both for the theory of modification and for the study of linguistic information-packaging more generally. With different linguistic phenomena in mind, precisely this analytical strategy was deployed by Kehler (2002) and in prior work by Kehler and colleagues, who suggested that coherence relations fall into three broad classes: cause-effect (e.g. RESULT), resemblance (e.g. PARALLEL), and contiguity (e.g. NARRATION). Kehler (2002) argued at length that processes like ellipsis and pronoun resolution are sensitive to the kind of coherence relation that is present. For example he proposes that gapping cannot be licensed by cause-effect relations, and that the grammatical mechanism that resolves verb-phrase ellipsis varies according to whether the antecedent and target clauses are linked by a resemblance relation or a cause-effect relation. Hardt & Romero (2004) similarly proposed that verb-phrase ellipsis in a clause can only be licensed if the antecedent stands in a (certain kind of) parallel relation with the target.

Attempting to precisely characterize the set of coherence relations that non-restrictive modifiers can express is an attractive topic for future investigation. The discussion in §1.6 suggests that the range must be diverse, comprising at least EXPLANATION, VIOLATED EXPECTATION, and a handful of others. At the same time, though, there may be non-trivial limitations. Consider RESULT, for example: can the content of a non-restrictive adjective express the result state of the event described by the sentence in which it occurs? I personally find it difficult to understand 3.101 without imagining my father as a quadriplegic *prior to* his injury—in spite of the radical implausibility of such a situation.²⁴

(3.101)The battle for Mars nearly killed—and permanently disabled—my quadriplegic father.

²⁴ The status of (i) is interestingly less clear.

- (i) My quadriplegic father sustained his injury in the battle for Mars.

To the extent that a RESULT interpretation is possible, this may be due to a quasi-anaphoric relation between an implied injury that caused paralysis and the noun phrase *his injury*. The RESULT reading becomes less salient if *his* is replaced by *an*, for example. Also potentially relevant is that the non-restrictive modifier in (i) is in the subject, but in 3.101 is in the object.

Such judgments must be established quantitatively before firm conclusions can be drawn, but if result interpretations are indeed impossible, this would suggest that non-restrictive adjectives cannot express just any coherence relation. The internal structure of adjectival modifiers also plausibly limits the set of relations they can express: it is difficult to imagine an attributive adjective forming part of a temporally advancing narrative with the clause in which it occurs. Thorough investigation might also reveal whether non-restrictive adjectives express a different set of relations than appositives, parentheticals, sequenced sentences, etc.

3.5.3 Operationalizing “non-restrictive” for behavioral stimulus creation

While the definition of non-restrictiveness offered in this chapter is mathematically unambiguous, it is nevertheless difficult to empirically determine whether a given modifier in a given utterance is restrictive or not. The reason is that there is no mechanical procedure to determine what a speaker believes about whether the extension of one phrase is a proper subset of the extension of another phrase. This presents a problem for behavioral and experimental methodologies, which require large, normed sets of linguistic stimuli. It is therefore useful to develop operational tests for determining if a particular modifier is non-restrictive or restrictive. The results of such tests could be used to select stimuli for studies investigating the role of restriction in language processing and communication. An ideal operational test for restrictiveness might collect truth-value or felicity judgments, and could be deployed with a online survey-based platform like Amazon Mechanical Turk. This method would provide quantitative assurance that a restrictiveness manipulation in an experimental design would actually manipulate restrictiveness.

Here is one sketch of how non-restrictiveness could be operationalized in the way described above (alternatives should certainly be sought). Suppose we want to know, for a collection of adjective-noun combinations, which of them can be naturally used for non-restrictive modification. The truth-conditional irrelevance characteristic of non-restrictive modifiers can be exploited as follows: given a sentence containing an unmodified noun phrase, adding a non-restrictive modifier to the noun should never affect the truth-value of the sentence. Put another way, the modified and unmodified versions should always have the same truth-value. Subjects could therefore be posed questions such as the following:

PosMod: Given that *John has a ball*, is it true that *John has a round ball*? YES // NO // NOT SURE

Because of the lack of specific information about John’s ball, a consistent answer of YES for this trial would indicate that *round* makes a good non-restrictive modifier for *ball*. On the other hand, an answer of NO or NOT SURE would indicate a poor combination for non-restrictive modification. For example I would

expect that the following question would be reliably answered in the negative, reflecting that *blue* is a poor non-restrictive modifier for *ball* in the general case (cf. *round* and *ball*).²⁵

PosMod: Given that *John has a ball*, is it true that *John has a blue ball*?

The polarity of a noun's environment affects the entailments it generates. For example *I saw a black crow* entails *I saw a crow* but not vice versa, while *I didn't see a crow* entails *I didn't see a black crow* but not vice versa. Therefore, reversing the polarity of the clauses while also exchanging the background-target positions (basically contraposition) should also yield affirmative answers for good non-restrictive combinations and negative or indeterminate answers for poor non-restrictive combinations. For example:

NegNoMod: Given that *John does not have a round ball*, is it true that *John does not have a ball*?

NegNoMod: Given that *John does not have a blue ball*, is it true that *John does not have a ball*?

To ensure subjects are performing the task as expected, the background and target sentences could be reversed or the polarity of the sentences could be reversed (but not both simultaneously). For these trials, affirmative responses are uniformly expected for non-restrictive and restrictive modification:

PosNoMod: Given that *John has a blue/round ball*, is it true that *John has a ball*?

NegMod: Given that *John does not have a ball*, is it true that *John does not have a round/blue ball*?

The full polarity-by-modification paradigm is illustrated represented in the following box.

PosMod: Given that *John has a ball*, is it true that *John has a round ball*?

PosNoMod: Given that *John has a blue/round ball*, is it true that *John has a ball*?

NegMod: Given that *John does not have a ball*, is it true that *John does not have a round/blue ball*?

NegNoMod: Given that *John does not have a blue ball*, is it true that *John does not have a ball*?

With this strategy in mind, one constructs a variety of modifier-noun combinations, embeds them into carrier sentences, and creates four trials for each modifier-noun pair (corresponding to the four conditions above). The results can then be interpreted as indicated in the boxes below, for any modifier-noun pair. An improved implementation of the strategy might adjust the profile of putative non-restrictive stimuli so that some negative or indeterminate responses are expected. Without such an adjustment to the design, subjects who simply answer YES uniformly would misleadingly appear to be performing flawlessly on non-restrictive trials.

²⁵ This again highlights the asymmetry between definites and non-definites: if only one ball is salient *blue* can be “non-restrictive” in the *blue ball*. But general characteristics of balls become relevant in non-definites, which is why *blue* does not easily modify *ball* non-restrictively in non-presuppositional DPs.

| |
|--|
| profile of non-restrictive mod (<i>round ball</i>) PosMod: YES PosNoMod: YES NegMod: YES NegNoMod: YES |
|--|

| |
|---|
| profile of restrictive mod (<i>blue ball</i>) PosMod: NOT SURE/NO PosNoMod: YES NegMod: YES NegNoMod: NOT SURE/NO |
|---|

A norming study of this kind—though not necessarily this specific task—would enhance the precision of any psycholinguistic study for which the restrictive/non-restrictive distinction is relevant. One important factor that the sketch above ignores is the effect of determiner type: as discussed above, non-restrictive modification in definite descriptions arguably interacts with the lexical presupposition of *the* while non-restrictive modification in non-definites requires appeal to world-knowledge (see also fn. 25).

Finally, note that the paradigm could be elaborated in various ways depending upon the relevant research questions. For example, to empirically test the hypothesis that postnominal modifiers in English lack non-restrictive interpretations, one could compare prenominally modified targets to postnominally modified targets, while holding constant the content of the modifier (e.g. *blue ball* versus *ball that's blue*).

We now proceed to Chapter 4, in which the direct/implicit relative Bolinger contrast for modal adjectives is described and analyzed in terms of the theory of DP developed in Chapter 2. Chapter 5 then considers a selection of additional modification-related puzzles, and shows how the analytical framework advanced in this dissertation can shed light on them as well.

Chapter 4

Case study: an ambiguity in modal adjectives

4.0 Introduction

This chapter investigates a Bolinger contrast that exists for a limited class of adnominal modal adjectives, most famously *possible*. This contrast has superficially the same distribution as many of the others discussed in Chapter 1: prenominal *possible* in (4.1a) admits of two distinct readings, which I will refer to as the “direct” and the “indirect” or “implicit relative” interpretations. Postnominal *possible* in (4.1b) admits of only the indirect interpretation (Larson 2000a).

- (4.1) a. Mary interviewed every possible candidate. (Larson 2000a:1)
✓*Indirect*: ‘Mary interviewed every candidate it was possible to interview.’
✓*Direct*: ‘Mary interviewed everyone who was possibly a candidate.’
- b. Mary interviewed every candidate possible.
✓*Indirect*, ✗*Direct*

Importantly, semantically similar adjectives such as *potential* do not exhibit this ambiguity: *potential* has no interpretation parallel to the indirect reading of *possible*, and is ungrammatical altogether postnominally (Larson 2000a). The fact that the pattern in 4.1 is only observed for a limited class of lexical items sets it apart from the more general restrictive/non-restrictive opposition from Chapter 4. And like i-level/s-level, intersective/adverbial, etc., ordering restrictions on iterated prenominal adjectives in English exist (i.e. *every possible possible...* is unambiguous with respect to this alternation; see §2), which I have argued is not the case for restrictive/non-restrictive.

Because the indirect interpretation of (4.1a) is accurately paraphrasable with a relative clause, this phenomenon has been used as evidence that postnominal adjectives in English are reduced (“implicit,” “under-

lying”) relative clauses.¹ The distribution of *possible* has been interpreted as strong evidence in favor of the two-domains model of nominal modification (Cinque (2010); Larson (1999,2000a); Demonte (2008))). In this chapter, though, I will argue that the distribution of *possible* and related items can be accounted for under the more conservative syntactic assumptions developed in Chapter 2.

The goals and structure of the chapter are as follows: in §1.1, I introduce the defining properties of the direct/indirect contrast, including the lexical and syntactic restrictions it is subject to. §1.2 summarizes various arguments found in the literature that properties of direct/indirect modals across languages provide evidence for a two-domains model of modification along the lines of Cinque (2010) or Larson (1999). In §1.3 I provide a summary and critical review of other existing analyses of this contrast. In §2 I show that direct/indirect readings of *possible* are not correlated with inner/outer modifier positions, contrary to the basic premise of two-domains based approaches. A single, polymorphic lexical entry for *possible* is introduced and shown to deliver correct results for direct readings while also allowing *possible* to compose with a propositional argument. §2.1 summarizes some concrete predictions of the two-domains model for *possible*. §2.2 shows that direct *possible* is not restricted to inner positions. §2.3 shows that a reduced relative clause structure for *possible* is not sufficient to derive correct truth-conditions for the implicit relative reading. Having established motivation for an alternative theory of direct versus implicit relative readings, the remainder of the chapter develops a new analysis of modal attributive adjectives, focusing on *possible* and *necessary*. §3 motivates and sketches a degree-based syntax and semantics for implicit relative adjectives, a generalization of Romero’s (2013) proposal for modal superlatives. §4.1-4.3 formalizes the new degree-based analysis and demonstrates via sample derivations that it generates correct truth-conditions for the basic cases. §4.4 revisits the distribution of attributive modal adjectives as described by Larson (2000a,2000b) and Cinque (2010), and shows point-by-point how the present theory captures this distribution accurately. §5 concludes.

The results of this chapter have implications for the theory of nominal modification. Two are that the inner/outer distinction among modifiers does not track the direct/indirect readings of modal adjectives as previous theorists have held; and that reduced relative clause structure sometimes does have strong motivation for postnominal adjectives. Therefore, both the two-domains model and the dual source assumption are shown to be involved crucially in the analysis of some Bolinger contrasts, but not others.

¹ Larson (2000a) dubbed this reading the “implicit relative reading,” a label that I will use interchangeably with “indirect” when the latter term is a potential source of confusion.

4.1 Background on *possible*: distribution and approaches

4.1.1 Restrictions on the distribution of indirect readings

In English, the availability of what I am calling the indirect or implicit relative reading of an attributive adjective depends upon something of a perfect storm of grammatical factors:² the adjective must appear either postnominally or (roughly) adjacent to D to receive an implicit relative reading (the *locality restriction*); the adjective must be a modal like *possible*, *conceivable*, or (as I will show) *necessary* (the *lexical restriction*); and finally (roughly) a universal quantifier must appear in the DP that the adjective is contained in. More precisely, only the Ds *every* and *all*, the superlative morphemes *-est/most*, and (as I will show) DP-internal *only* license the indirect reading of *possible*—this is the *determiner restriction* (I argue in §3.3 that a slightly looser restriction holds for *necessary*). In this subsection I illustrate each of these properties for English. Later sections develop explanations for why the restrictions ought to exist in the first place.

4.1.1.1 Determiner restriction

The direct/indirect opposition is rare among related contrasts in that one interpretation—the indirect or implicit relative reading—is licensed only in the presence of certain DP-internal material (Larson 2000a). Determiners that denote universal quantifiers (e.g. *every*, *all (the)*) support the implicit relative reading,³ as do definite descriptions containing a superlative (e.g. *the tallest*, *the most*), and definites with *only*. This is illustrated for universals in 4.1 above, and for superlatives and *only* in 4.2-4.3.⁴

(4.2) John climbed the tallest possible mountain.

✓*Indirect*: ‘John climbed as tall a mountain as it was possible to climb.’

✓*Direct*: ‘John climbed the tallest *x* such that *x* is possibly a mountain.’

(4.3) John climbed the only possible mountain.

✓*Indirect*: ‘John climbed the only mountain it was possible to climb.’

✓*Direct*: ‘John climbed the only *x* such that *x* is a possible mountain.’

² Larson (2000a) first identified the existence of these three restrictions, although my description of them here is distinct from his formulation, as the present characterization is informed by new observations to be introduced in this chapter.

³ The indirect reading is most clearly evident with *every*; judgments vary quite a bit with *each* and *all (the)*. Because of their unclear status with respect to indirect readings, I will set aside discussion of these latter two determiners in this dissertation.

⁴ Throughout I mainly use *possible* to illustrate various facts about implicit relative readings of adjectives. This is in part because the intuitions seem to be clearest with *possible*, partially because the previous literature is restricted almost entirely to examples with *possible*, and partially to avoid any confounds that may result from accidentally differing syntactic or semantic properties of different adjectives. The observation that *the only* supports implicit relative readings is new as far as I can tell.

Other determiners do not license indirect interpretations (4.4), even though there is nothing semantically incoherent about such hypothetical readings: the first paraphrase of 4.4 is exactly what the sentence would mean if *possible* could receive this reading in this context.

(4.4) John climbed a/three/many/few/no possible mountain(s).

✗*Indirect*: ‘John climbed a/three/many/few/no mountain(s) it was possible to climb.’

✓*Direct*: ‘John climbed a/three/many/few/n x s.t. x is a possible mountain.’

The determiners that license implicit relative readings all syntactically allow for postnominal *possible*, as shown in 4.5 (Larson (2000a)). 4.6 illustrates Larson’s (2000a) observation that postnominal *possible* is unambiguously indirect for all determiners that license it, a characteristic distributional property of Bolinger contrasts (see Chapter 2:§1).

(4.5) a. Mary sampled every/the sweetest/the only food(s) possible.

b. *Mary sampled a/three/many/few/no food(s) possible.

(4.6) a. Mary interviewed every candidate possible.

✓*Indirect*, ✗*Direct*

b. Mary interviewed the smartest candidates possible.

✓*Indirect*, ✗*Direct*

c. Mary interviewed the only candidate possible.

✓*Indirect*, ✗*Direct*

Additionally, the nominals that do support implicit relative readings of *possible*, *conceivable*, *imaginable*, etc. all appear to involve universal quantification at the level of truth-conditional meaning. To see this, consider the following simplified but more-or-less accurate lexical entries for the determiners that support modal readings of lower adjectives:

(4.7) a. $\llbracket \text{every} \rrbracket = \lambda P[\lambda Q[\boxed{\forall y} [P(y) \rightarrow Q(y)]]]$

b. $\llbracket \text{-est} \rrbracket = \lambda C_{\langle e,t \rangle}[\lambda D_{\langle d, \langle e,t \rangle \rangle}[\lambda x[\exists d[D_d(x) \wedge \boxed{\forall y} [[C(y) \wedge y \neq x] \rightarrow \neg D_d(y)]]]]]$

c. $\llbracket \text{only} \rrbracket = \lambda P[\lambda x[P(x) \wedge \boxed{\forall y} [P(y) \rightarrow y = x]]]$

In short, syntactic or semantic properties of determiners and determiner-like complexes of elements are relevant for the analysis of implicit relative *possible*. Coincidentally, properties of D also appear to be relevant for the licensing of postnominal adjectives in general. An important question, then, is why there should be any connection at all between universal quantification, implicit relative readings, and the licensing of postnominal adjectives. This question is partially addressed in §3 below.

4.1.1.2 Adjective restriction

Only a small number of adjectives can get implicit relative readings (Larson 2000a). Examples include *possible*, *imaginable*, *conceivable*, and *permissible*. This non-exhaustive list may give the impression that *-ible/-able* is crucially involved here, a generalization suggested by Larson (2000b) and argued for explicitly by Harris (2012). However, 4.8 demonstrates that the adjective *necessary* can receive an implicit relative reading as well, a fact heretofore unobserved and unaccounted for on existing theories. Therefore, morphological structure cannot be a necessary condition for the licensing of implicit relative readings.

(4.8) Mary interviewed every necessary candidate.

✓*Direct*: ‘Mary interviewed every person who was/is necessarily a candidate.’

✓*Indirect*: ‘Mary interviewed every candidate it was necessary to interview.’

Conversely, morphological structure cannot be a sufficient condition for the availability of a implicit relative reading, either: there exist many *-ible/-able* adjectives that lack modal readings altogether, e.g. *responsible* and *understandable*. Again, this is despite the fact that such readings would be perfectly coherent, as illustrated in the reasonable but unavailable paraphrases of (4.9a) and (4.9b) below.

(4.9) a. Mary interviewed every responsible/understandable candidate.

✓*Direct*: ‘every candidate who was responsible/understandable’

✗*Indirect*: ‘every candidate it was responsible/understandable for her to interview.’

b. Mary interviewed every candidate responsible/*understandable.

✓*Direct*; ✗*Indirect*

A number of adjectives semantically related to *possible* cannot receive implicit relative readings, showing that “modality” is also not sufficient to license this indirect reading (a *potential candidate* is something quite similar to a *possible candidate*, as Larson (2000a) observes). To see this, compare (4.10a) with (4.10b).

(4.10) a. Mary interviewed every likely/certain/probable/potential candidate.

✗*Indirect*, ✓*Direct*

b. Mary interviewed every possible/necessary/conceivable/imaginable candidate.

✓*Indirect*, ✓*Direct*

Larson (2000a) correlates availability of implicit relative reading with syntactic ability to take an infinitival complement, noting that predicative *possible* can take an infinitival complement whereas *potential/probable* cannot.

- (4.11) a. It is possible/conceivable (for Mary) to interview that candidate.
 b. *It is potential/probable (for Mary) to interview that candidate.

I believe that the significance of this correlation is not as great as it may appear. Turning again to *responsible/understandable*, observe that these adjectives—for which an implicit relative reading is impossible—nevertheless syntactically select infinitival complements (cf. 4.11):

- (4.12) It is responsible/understandable (for Mary) to interview that candidate. (cf. (4.9a))

That said, it does appear to be the case that all adjectives admitting of implicit relative readings are also capable of taking an infinitival complement, i.e. they are grammatical in the frame [It be $_ _$ [CP \dots [TP PRO to \dots]]]. In other words, compatibility with an infinitival complement may be a *necessary* condition for an implicit relative reading to be possible, but *not* a sufficient condition.

Exactly how to characterize the class of adjectives that can have modal readings is thus far from obvious. At least the following appear to be true: (i) being a “modal” adjective (making reference to non-actual states of affairs) is a necessary but not sufficient condition for having an implicit relative reading; (ii) presence of the suffix *-ible/-able* is neither a necessary nor a sufficient condition for having an implicit relative reading; and (iii) compatibility as a predicative adjective with an infinitival complement is a necessary but not sufficient condition for having an implicit relative reading.

4.1.1.3 Syntactic position and “locality restrictions”

Implicit relative readings of prenominal adjectives are subject to stricter “locality conditions” than are other indirect readings (Larson 2000a; Schwarz 2005). While ordering restrictions are observed in, e.g. i-level/s-level and intersective/adverbial (see Chapter 1), a stricter condition of *adjacency* is required to hold between an implicit relative adjective and the higher element that licenses it (i.e. *every*, *-est*, etc.).

- (4.13) a. I bought every possible affordable present. ✓Indirect, ✓Direct
 b. I bought every affordable possible present. ✗Indirect, ✓Direct

- (4.14) a. I bought the largest possible affordable present. ✓Indirect, ✓Direct

- b. I bought the largest affordable possible present. *✗Indirect, ✓Direct*
- (4.15) a. I bought the only possible affordable present. *✓Indirect, ✓Direct*
 b. I bought the only affordable possible present. *✗Indirect, ✓Direct*
- (4.16) a. I bought the most possible large presents. *✓Indirect, ✓Direct*
 b. I bought the most large possible presents. *✗Indirect, ✓Direct*

German *möglich* ‘possible’ exhibits a parallel direct/implicit relative ambiguity (Corver 1997; Schwarz 2005). According to Schwarz’s (2005) description of the data, implicit relative readings of *möglich* appear in a more restricted environment than for *possible* in English. Specifically, Schwarz (2005) reports that *möglich* can receive an implicit relative interpretation only in definite superlative DPs, and only if it “shares” an agreement suffix with the superlative adjective, as in example (4.17a). Agreement sharing in German pronominal adjectives, Schwarz states, is only possible when superlative *-st* ‘-est’ is adjacent to the adjective *möglich*.⁵

- (4.17) a. Ich habe das größt möglich.e Geschenk gekauft.
 I have the largest possible.INFL present bought
 ‘I bought the largest present possible.’ (*unambiguously modal*) (German)
- b. Ich habe das größt.e möglich.e Geschenk gekauft.
 I have the largest.INFL possible.INFL present bought
 ‘I bought the largest of the possible presents.’ (*unambiguously direct*)

Thus the German pattern is similar to but more restricted than the English pattern.⁶ However, in German superlative adverbial phrases, the morpheme *-st* can appear directly suffixed to the modal adjective *möglich*, as in the following example from the internet. Note that in this kind of example, the adjective whose degree is being compared (here *schnell* ‘fast’) is morphologically unmarked, as it is not an attributive form.

⁵ In German, every attributive adjective bears gender, number, and case morphology.

⁶ One contested issue in the syntax of modal adjectives is the question whether English analytic superlatives (*the most Adj*) are subject to the same kind of adjacency condition that synthetic superlatives (*the Adj-est*) are. Schwarz (2005) and Romero (2011) report that (i)ab do not have implicit relative readings. However, as a native speaker, I find this interpretation to be the most prominent, as is confirmed by other English speakers I have consulted.

- (i) a. I bought the most expensive possible present. *✓Indirect, ✓Direct*
 b. I drink the most expensive possible coffee. *✓Indirect, ✓Direct*

The status of (i) is important to determine for the analysis of *possible*. I will henceforth assume that the indirect reading in these cases is available.

- (4.18) Die Touchscreen-Anwendung sollte möglich-st schnell laufen.
 the touchscreen-application should possible-est fast run
 ‘The touchscreen application should run/function as fast as possible.’

The corresponding form **possiblest* is unattested in (Modern) English. A systematic comparison of superlative morphology in nominals and adverbials in English versus German may well shed light on this peculiar difference. Such investigation is unfortunately beyond the scope of this dissertation.

To summarize, *possible* is unambiguously indirect postnominally, and requires a certain proximity to the determiner (alternatively, perhaps distance from the noun) to be interpreted indirectly prenominally. The facts are slightly different in German, in which *möglich* apparently requires an even tighter syntactic relationship to the superlative morpheme in order to receive the indirect/implicit relative reading. These properties are not shared by other Bolinger contrasts, in English or as far as I know, in any other language. This makes the direct/indirect ambiguity in modal adjectives a particularly unique case study, and suggests that its properties may be more idiosyncratic than other Bolinger contrasts’.

4.1.2 The ambiguity as evidence for the two-domains theory

As discussed in Chapter 2, Cinque (2010) has argued that postnominal adjectives in Romance languages are often ambiguous in the same way that prenominal adjectives are in English—and the readings that are only available prenominally in English are the only readings available prenominally in Italian (see Cinque 2010:Ch2). Cinque (2010:7) reports that postnominally, Italian *possibile* ‘possible’ is ambiguous but prenominally it retains only the direct reading:

- (4.19) a. Maria ha intervistato ogni possibile candidato.
 M. has interviewed every possible candidate
 ‘Maria interviewed every potential candidate.’ (*unambiguous*)
- b. Maria ha intervistato ogni candidato possibile.
 M. has interviewed every candidate possible
 ‘Maria interviewed every possible candidate.’ (*ambiguous*) (Italian, Cinque 2010)

This is exactly the kind of mirror-image patterning expected on Cinque’s (2010) theory of modification in DP, and hence can be interpreted as evidence for his particular analysis.

Bolinger effects are also observed with multiple modifiers, such that—in terms of the two-domains model—English prenominal inner modifiers must be structurally closer to the noun than prenominal outer modifiers. Cinque’s (2010) theory predicts that the reverse relationship should hold among postnominal

adjectives in Romance languages. This theory therefore leads us to expect that implicit relative *possible* cannot occupy a position closer to the noun than can direct *possible* or any other inner adjective. This expectation appears to be borne out, a generalization observed by Larson (2000a) and illustrated in 4.20.

- (4.20) a. I bought the largest_{ind.} possible_{ind./dir.} affordable_{dir.} present. ✓*Indirect*, ✓*Direct*
 b. I bought the largest_{ind.} affordable_{dir.} possible_{*ind./dir.} present. ✗*Indirect*, ✓*Direct*

However, it is important to keep in mind the locality constraints from §1.1.3 above. The data in 4.20 could simply be a consequence of this restriction, thus rendering the inner/outer modifier distinction orthogonal. The following contrast from Larson 2000a, exactly parallel to Larson’s (1998) *nonvisible visible/*visible nonvisible* case, constitutes perhaps the most direct evidence for a two-domains based explanation of the direct/indirect ambiguity in modal adjectives (or at least an explanation that collapses the *visible* and *possible* patterns).

- (4.21) a. Mary interviewed every possible_{indirect} possible_{direct} candidate.
 b. *Mary interviewed every possible_{direct} possible_{indirect} candidate.

To conclude: the distribution of *possible* and other modal adjectives fits nicely within Larson’s (1999) and Cinque’s (2010) two-domains theories. However, postulating a split domain of modification alone says nothing about why determiner restrictions ought to exist, nor does it predict the stricter locality constraints observed with *possible* when compared to the looser inner/outer syntactic constraints observed in other Bolinger contrasts. We will revisit these points in §3, where I propose that independent properties of modal adjectives and their complements alone can explain the difference in locality constraints.

4.1.3 Summary and critical review of existing analyses

The direct/indirect ambiguity in *possible* and related items was initially discovered by Larson (2000a), who proposed that indirect *possible* is underlyingly a reduced relative clause (hence “implicit relative”). The phenomenon was then further analyzed and supplemented with cross-linguistic data by Schwarz (2005); Cinque (2010); and Romero (2011,2013). Harris (2012) recently argued that direct/indirect *possible* should receive the same semantic treatment as adjectives displaying the i-level/s-level ambiguity. This section provides a cursory overview of this body of literature, focusing on existing formal semantic analyses and their relationship to the theory of modification in the noun phrase.

Larson (2000a) argues that that the relative clause underlying *possible* has the form $[_{CP} \text{that it be } [_{AP} \text{possible } [_{CP_\blacktriangle} \dots]]]$, containing an antecedent-contained deletion (ACD) gap \blacktriangle . This gap is resolved on his theory by QRing the DP containing *possible* and reconstructing CP_\blacktriangle with an infinitive form of the matrix clause with a gap in object position.

- (4.22) a. Mary interviewed every candidate possible.
 b. $[_{CP} \text{Mary PST interview } [_{DP} \text{every candidate } [_{OP_i} \text{possible } [_{CP} \emptyset]]]]$
 \xrightarrow{QR} $[_{DP_i} \text{every candidate } [_{OP_i} \text{possible } [_{CP} \emptyset]]] [_{CP} \text{Mary PST interview } t_i]$
 $\xrightarrow{recon.}$ $[_{DP_i} \text{every candidate } [_{OP_i} \text{possible } [_{CP} \text{for Mary to interview } t_i]]]$
 $.[_{CP} \text{Mary PST interview } t_i]$

A notable property of this derivation is that after reconstruction occurs, the higher CP *for Mary to interview* t_i has an infinitive form, whereas the matrix CP *Mary PST interview* t_i is finite. In other words, this derivation requires the option of reconstructing a non-finite clause from a finite one. In Larson's analysis, an entire CP containing the matrix subject is reconstructed. This has the consequence of forcing *possible* to be relativized to an individual (here, the subject). In other words, on this analysis (4.22a) means that Mary interviewed every candidate that *Mary* was able to, not every candidate that *anyone* was able to.

On the basis of the argument-sharing data in 4.17, Schwarz (2005) proposes that *-st möglich* (and presumably *-(e)st possible* by extension) is a non-decomposable lexical item, with the semantics in 4.23 (where $[[R]] \subseteq W \times W$ an accessibility relation). Schwarz's analysis defines *-est possible* in terms of universal quantification over degrees, a technique common in the theory of superlatives (more in §4).

$$(4.23) \quad \llbracket \text{est.possible} \rrbracket^w = \lambda P_{\langle s, \langle d, t \rangle \rangle} [\forall d [\exists w' [wRw' \wedge P(w')(d) = 1] \rightarrow P(w)(d) = 1]]$$

This semantics results in the truth-conditions in (4.24c) for (4.24a), given the LF in (4.24b) (importantly, if $P(d)$ and $d' < d$, then $P(d')$ for any P and d), where A is the abstract indefinite determiner in comparative/relative superlatives that gets pronounced $[\delta\partial]$ (see Szabolcsi 1986; Heim 1999; Romero 2011; below).

- (4.24) a. John climbed the highest possible mountain.
 b. $[_{DegP} \text{est possible}] \lambda_1 [_{John} \text{climbed A } [_{AP} e_1 \text{high}] \text{mountain}]$
 c. $\forall d [\exists w' [@Rw' \wedge [\mathbf{John} \text{ climbed a } d\text{-high mountain}](w') = 1] \rightarrow P(@)(d) = 1]$
 'For any degree d , if it's possible that John climbed a d -high mountain, then John climbed a d -high mountain.'

On this semantics, (4.24a) means “John climbed as high a mountain as he possibly *climbed*,” which is problematic since (4.24a) actually means something closer to “John climbed as high a mountain as it’s possible *to climb*.” Note also that these truth-conditions too are relativized to the subject, *John*, and say something about *his* mountain-climbing abilities, not about mountain-climbing abilities in general.

In a series of papers, Romero (2011,2013) developed a degree-based analysis of implicit relative adjectives in superlative DPs (‘modal superlatives’). She adopts Schwarz’s [-*(e)st possible*] constituency, but the general strategy of Larson’s ACD-based derivation for the resolution of \blacktriangle . For her, [-est [λd_1 [*possible* \blacktriangle]]] is an LF constituent that QRs above the matrix subject, a position from which the elided material \blacktriangle can be recovered. Romero analyzes a type-shifted version of [λd_1 [*possible* \blacktriangle]] as the comparison class argument of Heim’s (1999) two-place *-est*, whose semantics is defined in (4.25a). A covert SHIFT operation which turns [λd_1 [*possible* \blacktriangle]] into a suitable argument for *-est* is defined in (4.25b).

$$(4.25) \quad \begin{array}{l} \text{a. } \llbracket \text{-est} \rrbracket = \lambda Q_{\langle\langle d,t \rangle, t \rangle} [\lambda P_{\langle d,t \rangle} [\exists d [P(d) \wedge \forall Q \in \mathcal{Q} [Q \neq P \rightarrow \neg Q(d)]]]] \\ \text{b. } \text{SHIFT}_{\langle d,t \rangle \rightarrow \langle\langle d,t \rangle, t \rangle}^\downarrow = \lambda D_{\langle d,t \rangle} [\lambda D'_{\langle d,t \rangle} [\exists d' [D(d') \wedge D' = [\lambda d'' [d'' \leq d']]]]] \end{array}$$

Romero’s LF for (4.26a) is given in (4.26b) and the truth-conditions are given in (4.26c). The surface constituent *possible* is a covert degree relative clause with semantic type $\langle\langle d,t \rangle, t \rangle$ after the application of SHIFT (a feature I will argue in §4 can be generalized to universals).

$$(4.26) \quad \begin{array}{l} \text{a. } \text{John climbed the highest possible mountain.} \\ \text{b. } \llbracket \text{DegP -est} \llbracket \text{XP } \lambda d_1 \llbracket \text{possible } \blacktriangle_{=IP^*} \rrbracket \rrbracket \\ \quad \llbracket \text{IP } \lambda d_2 \llbracket \text{IP}^* \text{ John} \llbracket \text{VP climbed} \llbracket \text{NP A } t_2\text{-high mountain} \rrbracket \rrbracket \rrbracket \\ \text{c. } \exists d [\exists x [\mathbf{mtn}(x) \wedge \mathbf{climb}(x)(\mathbf{j}) \wedge d\text{-high}(x)] \wedge \\ \quad \forall D' [\exists d' [\diamond \exists x [\mathbf{mtn}(x) \wedge \mathbf{climb}(x)(\mathbf{j}) \wedge d'\text{-high}(x)] \\ \quad \wedge D' = \lambda d'' [d'' \leq d'] \wedge D' \neq \lambda d_3 [\exists x [\mathbf{mtn}(x) \wedge \mathbf{climb}(x)(\mathbf{j}) \wedge d_3\text{-high}(x)]]]] \\ \quad \rightarrow \neg D'(d)]] \end{array}$$

‘There’s a degree (of height) *d* s.t. John climbed a *d*-high mountain and there’s no degree higher than *d* s.t. it is possible for John to climb a mountain of that height.’

While an innovative and well-motivated approach to modal superlatives, Romero’s (2013) degree-based analysis does not automatically extend to implicit relative adjectives in universals or in definites. In other words, what makes the analysis work is the degree-based semantics for \blacktriangle combined with the standard

degree-based semantics for *-est*. Plugging in Romero’s semantics for [λd_1 [*possible* ▲]] and a standard lexical entry for *every* in *every possible mountain*, for example, would yield a type-clash. This approach alone therefore does not explain the determiner restriction described above, a defining property of the implicit relative reading of modal adjectives. Postnominal occurrences of implicit relative adjectives, where *-est* and *possible* are not adjacent or even ‘local’, are also not predicted to be possible on this analysis (without further amendment), due to the assumption that *-est* and *possible* must form a syntactic unit.

4.2 The independence of direct/indirect *possible* from the two-domains model

In this section, I argue that there is no direct correlation between the inner versus outer position of a modal attributive adjective on the one hand, and whether it receives a direct or an indirect interpretation on the other. In effect, this shows that a two-domains model of nominal modification alone cannot explain the distribution of *possible* in English. The importance of this result is that it highlights the need for an alternative set of assumptions about what licenses direct and implicit relative readings in the first place.

4.2.1 Direct modal adjectives are not restricted to inner positions

The word *possible* has a syntactically diverse distribution: it can appear without a clausal complement in attributive position but generally not in predicative position 4.27,⁷ it participates in the *tough*-construction 4.28, and it can appear with a finite sentential complement as the predicate of a copular clause with an expletive subject 4.29.

(4.27) a. John is a possible winner. b. *The winner is possible.

(4.28) a. It is possible PRO/for Bill to please John. b. (?)John is possible PRO/for Bill to please.

(4.29) a. It is possible that Bill pleased John. b. *John is possible that Bill pleased.

The distributional properties illustrated by 4.27 imply *possible* must have a use on which it maps noun denotations to noun denotations. And given 4.28 and 4.29, it must also have a use on which it maps propositions to propositions. This kind of dual function is not rare among adjectives, for example *sad*, *gross*, *enticing*, and many others can compose with both common nouns in attributive position (*sad clown*), and clauses of some kind in predicative position (*It was sad for the clown to go to prison/that the clown went to prison*). However, all semantically explicit existing analyses of the direct/indirect alternation with *possible* assume

⁷ The exception: *possible* can appear in predicative position if the subject denotes a proposition, e.g. *That John will win is possible*, or *That is possible* where *that* refers to some salient utterance/proposition.

distinct lexical items for each reading. For example, Harris (2012) and Schwarz (2005) both assume that the occurrences of *possible* in (4.30a) and (4.30b) involve distinct lexical items. In both accounts, propositional *possible* in (4.30a) is assigned the type of a sentential operator, and attributive *possible* in (4.30b) is assigned the type of an operator on noun meanings.

- (4.30) a. It is possible that the senator left. (propositional *possible*)
 b. John is a possible senator. (attributive *possible*)

This ambiguity is crucially invoked in the analysis of direct versus indirect readings of *possible*. On Harris's (2012) and Schwarz's (2005) proposals, the former involves attributive *possible* and the latter involves propositional *possible*. Harris's semantics for the two items, for example, are as follows:

- (4.31) Harris's (2012) lexical entries for two *possibles*
 a. $\lambda P_{\langle e, \langle s, t \rangle \rangle} [\lambda x [\lambda s [\exists s' [s \leq_P s' \wedge P(x)(s')]]]]$ (propositional *possible*)
 b. $\lambda P [\lambda s [\text{EXH}^\diamond(P)(s)]]$ (attributive *possible*)

This is a genuine lexical ambiguity because the two entries contain distinct operators in the body of the lambda-term—(4.31a) and (4.31b) cannot be two instances of a single, generalized schematic term.

I would like to suggest that postulating this kind of lexical ambiguity is unnecessary. Although 4.30 shows that *possible* must be capable of composing with constituents of different semantic types, that alone does not imply lexical ambiguity: coordinators like *and*, for example, are polymorphic but not ambiguous. And furthermore, if 4.30 is sufficient evidence to posit two distinct *possibles*, the argument for ambiguity could be replicated for *sad*, *gross*, etc.—an ontological proliferation.

I propose the following polymorphic semantics for *possible*. Here, R is a free variable over accessibility relations whose value is either grammatically specified (e.g. by a PP adverbial like *in view of what is known/allowed/etc.*) or else provided contextually. This is motivated by the fact that *possible* is compatible with a variety of modality types, including deontic, epistemic, alethic, etc. I allow for the possibility that the string of σ 's in 4.32 can be empty, in which case *possible* will be an operation on propositions (see (4.33b)).

- (4.32) **Polymorphic semantics for *possible***

$$\llbracket \text{possible} \rrbracket = \lambda P : \langle s, \langle \sigma_0, \dots \langle \sigma_n, t \rangle \rangle \rangle [\lambda s : s [\lambda v_0 : \sigma_0 \dots [\lambda v_n : \sigma_n [\exists w [sRw \wedge P_w(v_0) \dots (v_n)]]]]]]$$

Consider the semantics that 4.32 generates for attributive and propositional uses of *possible*, shown in (4.33a) and (4.33b), respectively. These are exactly the kinds of meanings that we want under a classical possible-worlds semantics for modals (such as Kratzer’s (1981)).

- (4.33) a. $\llbracket \text{possible senator} \rrbracket = \llbracket \text{possible} \rrbracket (\llbracket \text{senator} \rrbracket)$
 $= \lambda P_{\langle s, \langle e, t \rangle \rangle} [\lambda s [\lambda x [\exists w [sRw \wedge P_w(x)]]]] (\lambda s_2 [\lambda x_2 [\text{senator}_{s_2}(x_2)]])$
 $= \lambda s [\lambda x [\exists w [sRw \wedge \text{senator}_w(x)]]]$
 \approx ‘the set of individuals that are senators in some accessible world’
- b. $\llbracket \text{possible} [\text{CP that } \varphi] \rrbracket = \llbracket \text{possible} \rrbracket (\llbracket [\text{CP that } \varphi] \rrbracket)$
 $= \lambda p_{\langle s, t \rangle} [\lambda s [\exists w [sRw \wedge p(s)]]] (\lambda s_2 [\varphi(s_2)])$
 $= \lambda s [\exists w [sRw \wedge \varphi(s)]]$
 \approx ‘the proposition that φ is true in some accessible world’

These simple computations are meant to illustrate that the syntactic variability of *possible* and related adjectives does *not* entail a corresponding semantic variability: *possible* has a schematic but invariant lexical semantics that can be instantiated by various types. Different syntactic environments will require different instantiations of the general semantics, just as different environments call for different argument types for *and* or *or*. The semantics for *the* from Chapter 2 is another example.

Having introduced an unambiguous analysis of *possible*, we can now ask whether such an analysis sheds any light on the direct/implicit relative ambiguity. Assuming the framework from Chapter 3, here is one interesting and novel prediction: there are two distinct syntactic structures for direct *possible*, namely NP-adjunction and head-adjunction. As shown in 4.34, these structures generate slightly different interpretations for *possible senator*.

(4.34) **Two possible analyses of direct modal adjectives**

- a. **NP-adjunction:** $[\text{NP} [\text{AP possible}] [\text{NP} [\varphi [\text{SG}]]] [\text{N senator}]]]$
 $\lambda s [\lambda x [\exists w [sRw \wedge \text{senator}_w(x)]]]$
‘the set of individuals that are possibly atomic senators’
- b. **Head-adjunction:** $[\text{NP} [\varphi [\text{SG}]] [\text{N} [\text{A possible}] [\text{N senator}]]]$
 $\lambda s [\lambda x [^\circ (\lambda y [\exists w [sRw \wedge \text{senator}_w(y)]])(x)]]$
‘the set of atomic individuals that are possible senators’

Nothing from Chapter 2 blocks either of these structures in 4.34. Therefore, in the absence of reason to think otherwise, I assume that both structures should be made available by the grammar. The slight meaning difference between the structures in 4.34 are, I believe, unimportant. As far as I can tell, it is innocuous to assume that if one is an atomic individual in one world, then *s/he* is in all worlds; similarly for non-atomic individuals.⁸ If correct, this assumption renders the two options in 4.34 equivalent.

The availability of both (4.34a) and (4.34b) as grammatical phrases has an important consequence for the architecture of modification: modal adjectives like *possible* can, on the present proposal, receive a direct reading in both an inner position (as in (4.34b)) and in an outer position (as in (4.34a)). The theories of Cinque (2010) and Larson (1999) predict this to be impossible, since on these approaches the syntactic position of a DP-internal adjective is what determines its interpretation. According to the principles from Chapter 2, though, inner versus outer syntactic position does *not* disambiguate the between the two readings of modal adjectives. No non-*ad hoc* principle blocks a direct reading of *possible* in an outer modifier position, and therefore the direct reading is generated in both inner and outer positions.

4.2.2 Outer position and reduced relative structure do *not* imply ‘implicit relative reading’

The syntactic position of adjectives with implicit relative readings is less straightforward to assess. On Larson’s (1999) and Cinque’s (2010) dual source model of modification, outer adjectives are the predicates of a covert, reduced relative clause structure. This assumption is taken to explain why outer adjectives differ semantically from inner adjectives: the former are predicative and intersective, while the latter are attributive and may or may not be intersective. At face value, then, this theory predicts that *every candidate possible* should be semantically equivalent to *every candidate that was possible*. Besides being at least mildly ungrammatical, though, *every candidate that was possible* just does *not* intuitively mean the same thing as *every candidate possible*. Instead, as Larson (2000a) noted, *every candidate possible* means something closer to *every candidate that it was possible to ▲*, where ▲ is an ellipsis site. This discrepancy implies that even within a two-domains model, something extra needs to be said about implicit relative *possible*—if the reduced relative structure comes from the outer position, then where does the complement of *possible* along with the ellipsis site come from? The correct conclusion, I believe, is that the licensing of implicit relative readings is dependent upon not just outer position and reduced relative structure, but also upon whatever conditions license an infinitival complement of *possible*. We now turn to this and related issues.

⁸ A conclusive argument to this effect must take issues about trans-world identity seriously. Such issues are beyond the scope of this dissertation.

4.3 The internal structure of implicit relative modal adjectives in English

In this dissertation, I have so far made no use of reduced relative structures. This is not because I believe there is clear evidence *against* them, but rather because I have attempted to show how all relevant facts can be accounted for with only the null assumption that an adjective is simply an adjective, independent of its position or function in a clause. However, Larson (2000a) showed that there are strong reasons to suppose that indirect readings of *possible* do involve a certain kind of reduced relative clause structure. Here I will provide additional arguments for a similar conclusion, but because these arguments invoke specific properties of modal adjectives, they cannot be replicated for non-modal adjectives or generalized much at all. The following claim hence pertains only to a small class of lexical items: both pre- and postnominally, indirect readings of modal adjectives arise precisely when the adjective is surrounded by a covert, reduced relative clause in which the modal is an AP predicate that has an silent infinitive clause complement with a degree-based semantics (the silent complement is called ▲).

4.3.1 Modality type restrictions as evidence for covert structure

The first argument for a covert relative clause in implicit relative readings is based on a difference in the set of modality types or flavors compatible with direct versus indirect modal adjectives. *Possible* is compatible with a number of different modality types, and syntax at least partially restricts the kinds of modality types *possible* can express. In (4.35a), *possible* is compatible with epistemic (knowledge) only; in (4.35b), deontic (permission) and circumstantial (ability) only; in (4.35c), all of epistemic, deontic, and circumstantial.

- (4.35) a. It is possible that Mary left.
b. It is possible for Mary to leave.
c. He's a possible candidate.

The indirect reading of 4.36 supports only deontic and circumstantial modalities, which is exactly the set of interpretations available when predicative *possible* takes an infinitival complement (cf. (4.35b)). Crucially, this is not the same set of available modality types when *possible* has a direct reading, as in *John is a possible murderer*; in that case, an epistemic interpretation is also available.

- (4.36) Mary interviewed the tallest candidate possible.
Circumstantial: 'the tallest candidate she was able to interview'

Deontic: ‘the tallest candidate she was allowed to interview’

Epistemic: #‘the tallest candidate that she might have interviewed (for all we know)’

The same generalization can be made about *necessary*, suggesting that implicit relative *possible* and *necessary* have the same semantic role that they do with an overt infinitive complement. To see this, note first that examples parallel to 4.35 can be constructed for *necessary* and we get the same results: (4.37a) is epistemic only (to the extent that the sentence is felicitous); (4.37b) is deontic (“she’s not allowed here”) or circumstantial (“she’ll be late if she doesn’t”); and (4.37c) appears to be compatible with epistemic, deontic, and circumstantial modal flavors.

- (4.37) a. It is necessary that Mary left (given what we know about the world).
b. It is necessary for Mary to leave.
c. She’s a necessary candidate.

Again, *necessary* can only be interpreted deontically or circumstantially, coinciding exactly with the set of modality types supported by predicative *necessary* when it occurs with an infinitive clause as complement.

- (4.38) Mary interviewed the fewest candidates necessary.⁹

Circumstantial: “the smallest *n* such that interviewing *n* candidates achieved some goal”

Deontic: “the smallest *n* such that interviewing *n* candidates is permissible”

Epistemic: #“the smallest *n* such that for all we know, Mary interviewed *n* candidates”

4.35-4.38 constitute strong evidence that on the implicit relative reading, *possible* and *necessary* have the syntactic role of a predicative adjective with an infinitive (not finite) complement clause, and *not* the syntax of a bare attributive adjective. Therefore, we have strong motivation for postulating a silent relative clause structure ▲ surrounding indirect *possible/necessary/etc.*¹⁰

4.3.2 Arguments that ▲ is always a degree relative

Having established that indirect modal adjectives involve some kind of covert relative clause structure, we are now in a position to ask precisely what kind of structure that might be, and precisely what kind of meaning it might have. It will turn out that, in answering these questions, we will be led to an important

⁹ A quantity superlative is used here to increase pragmatic plausibility; I do not believe this is crucial.

¹⁰ I will not investigate in detail the mechanism by which the empty element receives its meaning; a number of possibilities exist; see Larson 2000a and Romero 2013 for a move-and-reconstruct analysis.

revision of Larson’s determiner restriction (see §3.3) on implicit relative readings. This in turn will provide strong evidence that the complement of *possible* is always a degree-denoting relative clause, an idea that generalizes Romero’s (2011,2013) analysis of modal superlatives. In virtue of this, we also generalize her explanation of the locality constraints from §1.1.3 to universal and plain definite nominals. We now return to the determiner restriction with the goal of gaining indirect insight into the semantics of ▲.

4.3.2.1 Revising the determiner restriction

Larson (2000a); Schwarz (2005); and Romero (2011,2013) have all noted that implicit relative readings of *possible* are available only in the presence of a superlative or a universal. In these contributions *possible* was taken to be a representative example of implicit relative adjectives in general, so research on this topic has been guided by the notion that this determiner restriction is completely general for adnominal modal adjectives. Here I propose that the determiner restriction is actually less stringent than it appears, but that the lexical meaning of *possible* is incompatible with an indirect reading in *the*-headed definite descriptions. More specifically, I will argue that combined with only the definite article and a head noun, implicit relative *possible* introduces a presupposition failure not introduced by other modal adjectives like *necessary*. Therefore, this approach guarantees that *the possible NP* cannot in general have an implicit relative reading, as desired. Since there is no such clash between *possible* and *every* or *-est*, though, implicit relative readings are predicted to be available in the presence of these elements. On the other hand, no clash exists between *necessary* and the definite article, so *the necessary NP* can admit of an implicit relative reading. The degree-based semantics I develop for ▲ constitute the first analysis that directly predicts this novel observation. It also explains why indirect *necessary* has a wider distribution than does indirect *possible*. The treatment of all implicit relative modals as containing a covert degree relative clause is a non-obvious but natural extension of Romero’s (2011,2013) proposal for modal superlatives.

The following collection of sentences shows that *necessary* can receive an implicit relative reading with a slightly wider range of determiners than can *possible*: in addition to universals and superlatives, the definite article appears to license an implicit relative interpretation of *necessary*, as long as the head noun is plural or mass. This is illustrated in 4.39. Note also that *the only* and *only the* with a plural noun appear to also license the relevant interpretation of the adjective.

- (4.39) *Context*: I think we can survive in the wilderness, because we have what we need...
- a. I brought every tool necessary.

- b. I brought the fewest tools necessary (so we can travel light).
- c. I brought all the tools necessary.
- d. I brought the tools necessary.
- e. I brought the three tools necessary.
- f. I brought [the only/only the] tools necessary.

Importantly, the judgments remain more-or-less unchanged if *necessary* appears in a prenominal position. This fact reinforces the intuition that the sentences in 4.39 do indeed illustrate the relevant interpretation. As expected, the direct reading also becomes available in 4.39' due to the prenominal position. Although deviant for irrelevant reasons (i.e. world-knowledge violation), the sentences in 4.39' have a reading on which *necessary tools* refers to the set of objects that are necessarily tools.¹¹

- 4.39' a.' I brought every necessary tool.
 b.' I brought the fewest necessary tools (so we can travel light).
 c.' I brought all the necessary tools.
 d.' I brought the necessary tools.
 e.' I brought the three necessary tools.
 f.' I brought [the only/only the] necessary tools.

As with *possible*, indefinite, proportional and numerical quantifiers all fail to license the implicit relative reading of *necessary*. This is shown for the postnominal position in 4.40 and for the prenominal position in 4.41 (where *I brought...* should be prefixed to each example).

- | | | | |
|--------|--|--------|---|
| (4.40) | a. *[some/several/0] tools necessary. | (4.41) | a. ??[some/several/0] necessary tools. |
| | b. *[few(er than three)] tools nec. | | b. *[few(er) than three] nec. tools. |
| | c. *[each/no] tool(s) necessary. | | c. *[each/no] necessary tool(s). |
| | d. *[the tool _{SG}] necessary. | | d. ??[the] necessary [tool _{SG}]. |

So what is to be made of this distributional difference between *possible* and *necessary*? It may appear at first to be a minor idiosyncrasy since Larson's (2000a) determiner restriction still generally holds—it is just

¹¹ The examples in 4.39'-4.41 may be complicated by an additional reading, paraphraseable as "I brought every tool that was necessary," in which *necessary* takes no complement. Such a parallel reading is not possible to test with *possible*, since the corresponding paraphrase is ungrammatical. I leave this issue open for future research.

that for certain lexical items the definite determiner in addition to universals and superlatives can support implicit relative readings.

4.3.2.2 Explaining the determiner restriction: A parallel with amount relatives

4.39-4.41 reveal a surprising empirical parallel that has consequences for the present analytical domain. The parallel is with another determiner restriction that exists for amount relatives and/or relatives out of existentials (depending on whose analysis; see Carlson 1977a for foundational discussion, and for recent theoretical analyses Grosu & Landman 1998; Grosu & Krifka 2007; McNally 2008; Herdan 2008; Bylinina 2013). Amount readings of relative clauses are semantically distinct from ordinary restrictive readings in that they intuitively denote *quantities* of stuff instead of properties of individuals. The amount reading can be illustrated by the following ambiguous sentence, due to Heim (1987).

- (4.42) It would take weeks to drink the champagne they spilled that evening.
- a. *Restrictive*: ...to drink said champagne off the ground
 - b. *Amount*: ...to drink an *amount* of champagne equal to the amount they spilled

Sometimes alternatively called “degree relatives,” amount relatives are often discussed alongside two other sentence-types: relatives out of existentials (e.g. 4.43), and antecedent-contained deletion (ACD) relatives (e.g. 4.44).¹²

- (4.43) a. Bill took with him the three books [RC there were __ on the table].
b. You’ve eaten every cookie [RC there was __ in the house].
c. Every man [RC there was __ on the life raft] died.

- (4.44) Marv put everything he could in his pocket.
- a. *Restrictive*: ‘every object he was able to put in’
 - b. *Amount*: ‘as much stuff as he was able to put in’

All three constructions—amount, existential, and ACD relatives—have been argued to semantically involve quantification/abstraction over degrees (Carlson 1977a; Heim 1987; Grosu & Landman 1998; a.o.). On this view, this type of relative clause contains in its gap position a silent degree variable or else a covert

¹² McNally (2008) has shown that there are important differences between amount, existential, and ACD relatives. But we will not go into detail, as the goal of this discussion is only to illustrate that a parallel determiner restriction exists for amount relatives and related constructions.

expression *d-many NP(s)*. For example an abbreviated LF representation for (4.43a) would be roughly as follows (simplified from Grosu & Landman 1998).

(4.45) [DP the [NP books_i λd[[CP that [IP there be *d-many* books_i [PP on the table]]]]]]

The parallel that I suggest between these constructions and implicit relative *necessary* is quite direct: only the definite article and universals like *every* (and *any*) are compatible with amount relatives, broadly construed (Carlson 1977a). This pattern is nearly identical to the one observed with *necessary* above, strongly suggesting that the two construction types exemplify a common phenomenon:

- (4.46) a. Every hour this movie lasts beyond my bedtime means more aggravation for me.
 b. [The, Those, Any] hours the movie lasts beyond my bedtime make little difference.
 c. *Several hours that the movie lasted past my bedtime passed quickly.
 d. *We whiled away some hours which the movie lasted past my bedtime.

(Carlson 1977a:530)

The same determiner restriction is observed in relatives-out-of-existentials, and the judgments turn out to feel somewhat sharper. The following contrast additionally shows that attributive superlatives are possible in this frame (see (4.47d)), thus strengthening the parallel between implicit relative modal adjectives on the one hand, and degree-denoting relative clauses on the other.

- (4.47) a. every book there was on the table
 b. the (three) books there were on the table
 c. the [(one)/(only)] book there was on the table
 d. the longest books there were on the table

- (4.48) a. *several/some/∅ books there were on the table
 b. *few(er than three)/many/most books there were on the table

And finally, recall from (4.40d) above that *necessary* cannot have an indirect reading with plain *the* if the head noun is count and singular (e.g. **the tool necessary*). Such a restriction is not observed in the presence of *only*, however (cf. *the only tool necessary*). This pattern is mirrored by the relative-out-of-existential construction, emphasizing the depth of the parallel:

- (4.49) a. The [one/only] man there was on the life raft died.
 b. *The man there was on the life raft died.

There is much more to say about the distribution of amount relatives and related constructions (see Carlson 1977a; Grosu & Landman 1998; McNally 2008), but of most relevance for us here is the following generalization suggested by the last handful of examples introduced above:

(4.50) **Determiner restriction for amount relatives**

Relative clauses with a degree-based semantics are only licensed by *the*, (as long as the noun is not count and singular), by *the only*, by *-est*, and by *every*.

Thus, if the implicit relative adjectives *possible* and *necessary* always contain a covert degree-referring relative clause structure, then the determiner restriction introduced in §3 above actually follows from 4.50 (with one exception, to be discussed shortly). Notice that this conclusion is vaguely reminiscent of Larson's (2000a) proposal that implicit relative *possible* is an ACD construction. In fact, if ACD relatives in general also have a degree-based semantics, that would explain why they too share the determiner restriction.

So why is it then that *necessary* but not *possible* can have an implicit relative reading in simple definite descriptions like the following?

- (4.51) a. Mary interviewed the candidates necessary.
 b. *Mary interviewed the candidates possible.

I will argue in §4 that the licensing of an implicit relative reading depends upon the presence of a universal quantifier somewhere in DP, and that this fact explains the asymmetry: *possible* is an existential quantifier over worlds/situations, while *necessary* is a universal. I will spell out this account in full detail in §4.4s, after the general analysis of implicit relative readings is developed. But we first motivate the generalized degree-based analysis via some background on the syntax of attributive adjective-infinitive structures.

4.3.3 The syntax of DP-internal infinitives, and explaining the locality restriction

Constructions of the form [DP D [NP N [AP A [CP ...]]]] are plentiful in English. For example, the following noun phrases are similar to *every candidate possible* ▲ in containing a determiner, a noun, an infinitive clause, and an adjective. Furthermore, *for DP* can be inserted between the noun and the infinitive in each case, e.g. *a hard nut for me to crack*, just like *...possible for Mary to interview*.

(4.52) **Clausal AIC**

- a. a hard nut to crack
- b. an easy book to read
- c. a good person to talk to

(4.53) **Nominal AIC**

- a. a short guy to draft
- b. a long book to assign
- c. a pricey car to drive like that

In a series of papers Fleisher (2008,2011) has shown that the nominals in 4.52 and in 4.53, despite their superficial similarity, fall into two subclasses with distinct syntactic and semantic properties. These classes are called the clausal and nominal attributive-with-infinitive constructions (AICs), respectively. One difference between the two classes can be seen in the following examples: only clausal AICs can be paraphrased with a sentence containing an expletive subject and a predicative AP containing the adjective and infinitive (here ‘#’ is # on intended reading).

- (4.54)
- a. It is hard to crack this nut.
 - b. It is easy to read this book.
 - c. It is good to talk to this person.

- (4.55)
- a. #It is short to draft this guy.
 - b. #It is long to assign this book.
 - c. #It is pricey to drive this car like that.

The nominal attributive-with-infinitive construction differs from the *tough*-construction in involving attributive adjectives instead of predicative adjectives (Fleisher 2011). The clausal AIC, however, appears to be possible for all *tough*-type adjectives: the examples in 4.52 are grammatical in the *tough*-frame in 4.56, but the nominal AIC examples in 4.53 are marginal at best in this environment, as shown in 4.57.

- (4.56)
- a. This nut is hard to crack.
 - b. This book is easy to read.
 - c. This person is good to talk to.

- (4.57)
- a. ??This guy is short to draft.
 - b. ??This book is long to assign.
 - c. ??This car is pricey to drive like that.

In nominal AICs, the adjective semantically applies to the noun in a way that the adjective of a clausal AIC does not: a short guy to draft is a guy with some degree of shortness, but a car that is an easy car to drive cannot be considered “easy” independently of some capacity or activity such as driving. Fleisher (2011:345-346) states that “[N]ominal-AIC DPs behave like ordinary predicative DPs in which an attributive adjective modifies the following noun, while clausal AICs behave as if the attributive adjective does not modify the noun at all.” This consideration and a number of others led Fleisher (2008) to analyze the adjective and the infinitive of a clausal AIC as a syntactic constituent to the exclusion of the noun. The underlying structure

Fleisher posits for *a good person to talk to*, for example, is abbreviated in 4.58. On his analysis, the CP is extraposed at surface structure, so that the linear order of constituents is D-A-N-CP (*a good person to talk to*), and not D-A-CP-N (**a good to talk to person*).

(4.58) [DP [D a] [NP [AP [A good] [CP to talk to]] [NP [N person]]]]

Such a constituency reflects the intuition that in clausal AICs, the adjective does not modify the noun directly, but rather acts something like the predicate of a reduced relative clause (cf. *a person who is good to talk to*).

We can use Fleisher's diagnostics to shed light on the structure of DPs containing implicit relative adjectives: for a noun phrase like *every possible candidate*, we can ask (i) does *possible* apply semantically to *candidate* or not? and also (ii) does *every possible candidate* pattern like a clausal or nominal AIC with respect to the frames in 4.54-4.57? The answer to (i) appears to be that *possible* does *not* apply directly to *candidate*. If it did, then the following syllogism would be valid, since Bill's being a candidate makes him a possible candidate. This contradicts the intuition that 4.59 is clearly invalid on the relevant interpretation of P₁.

(4.59) P₁ Mary interviewed every possible candidate. P₂ Bill is a candidate.
 ∴ Mary interviewed Bill.

The invalidity of 4.59 suggests that implicit relative *possible* cannot have the syntax of a nominal AIC adjective like *short* in 4.53.

Concerning question (ii), the judgments in 4.60 further suggest that implicit relative *possible* patterns more like a clausal AIC adjective than it does a nominal AIC adjective (cf. 4.56 and 4.57):

(4.60) a. It is possible (for Mary) to interview this candidate.
 b. ?This candidate is possible (for Mary) to interview.

If we are correct in considering *possible* to be a clausal AIC adjective, then something like Fleisher's structure in 4.58 could be extended to implicit relative *possible* (consistent with Larson's (2000a) and to some extent Romero's (2013) syntax). That said, there are other ways to implement the constituency and each will predict different truth-conditions depending upon the nature of *possible*'s covert argument, whether the head noun syntactically originates within said complement or is purely external to it, etc.

I propose that the syntax of DPs containing implicit relative *possible* do have roughly the constituency

in 4.58. But the considerations I will now introduce demand that the structure be slightly altered and supplemented with an additional structural layer: the degree phrase (DegP). Some items commonly assumed to occupy Deg head positions are *too*, *so*, comparative *-er*, and superlative *-est* (see Abney 1987 for early discussion). A DegP complement of D has been argued to exist for postnominal infinitive structures such as (4.61a) and (4.62a) by Dubinsky (1998). The structures that Dubinsky posits for (4.61a) and (4.62a)—(4.61b) and (4.62b), respectively—contain the silent degree operators MAX (for “maximal”) and SUFF (for “sufficient”), respectively, whose semantics are discussed in §3 below (none were provided in Dubinsky 1998).

- (4.61) a. THE ([ði]) car to drive is a Porsche.
 b. [DP [D the] [DegP [Deg MAX] [NP [NP car] [CP to drive]]]]
- (4.62) a. A person to pilot the ship is what Jason was looking for.
 b. [DP [D a] [DegP [Deg SUFF] [NP [NP person] [CP to pilot the ship]]]]

This analysis of postnominal infinitives can be extended to cases in which superlative morphology is present, as in (4.63a). There are independent reasons for analyzing superlative *-est* as a Deg head (see Heim 1999; Kennedy 1999 and references therein), so the structures in 4.61 and 4.62 model sentence (4.63a) in a natural way. Unlike the adjectiveless cases, though, the analysis of (4.63a) involves an overt degree head (*-est* instead of SUFF/MAX), and also overt movement of the adjective to the Deg position, in which a complex head *easy-est* is formed.

- (4.63) a. the easiest car to drive
 b. [DP [D the] [DegP [Deg -est] [NP [NP [N car]] [AP [A easy] [CP to drive]]]]]
 \Rightarrow [DP [D the] [DegP [Deg [A_i easy] [Deg -est]] [NP [NP [N car]] [AP t_i [CP to drive]]]]]

This syntax predicts a very high default position for adjectives that are associated with DP-internal infinitival clauses, at least when independent grounds exist for positing a DegP layer inside DP. The contrasts in 4.64 are expected on this analysis, since superlatives and *tough*-adjectives with infinitives are required to project a DegP layer above NP and since NP is the structural layer in which attributive adjectives are adjoined.

- (4.64) a. the easiest imported car to drive // *the imported easiest car to drive
 b. an easy trashy novel to read // *a trashy easy novel to read

If structures of the kind in 4.64 are syntactically parallel to implicit relative readings of modals in having the general shape [_{DP} D [_{DegP} ...]], then we also have an explanation for the locality restriction introduced in §1.1.3. Specifically, if *possible* in *every possible candidate* occupies the same position as *easy* in (4.63a), then it follows that no adjective can intervene between Deg and *possible* (provided we assume that adjectives can only adjoin to NP or N, as proposed in Chapter 3:§4). This correctly predicts the distribution of implicit relative readings prenominal in English (see §1.2 for details):¹³

- (4.65) a. Mary interviewed every possible smart candidate. ✓*implicit relative*
 b. Mary interviewed every smart possible candidate. ✗*implicit relative*

To summarize so far: I have provided theoretical and empirical motivation that implicit relative readings of modal adjectives in English contain a reduced relative clause structure with the semantics of a degree relative; that when an implicit relative adjectives is present in DP, D selects for DegP instead of NP; and that Deg can be filled by superlative *-est*, or else a silent operator like MAX or SUFF. These assumptions account for two defining syntactic properties of implicit relative readings: the determiner restriction—that only universals license this reading of modal adjectives—and the locality restriction—that an adjective must be adjacent or nearly adjacent to D in order to have this reading.

4.4 Formalizing the new analysis of implicit relative modals in English

In this section I formalize the amount relative analysis of indirect modal adjectives sketched in the previous section. I focus on implicit relative readings of the lexical items *necessary* and *possible*, the latter of which we've seen has a slightly more restricted distribution than the former (§3.3). I believe that the analysis presented can be extended to *imaginable*, *conceivable*, and any other adjective that can have the relevant kind of reading.

The presentation proceeds as follows: I first present derivations for the syntax and semantics of implicit relative readings in *every*-universals and in superlative DPs. When relevant, points of divergence between my analyses and those of Romero (2013), Larson (2000a), and Schwarz (2005) are identified. Following this, we turn to the topic of implicit relative adjectives in plain definite descriptions. Building on the discussion in §3.3, it is proposed that implicit relative readings are impossible with *possible* because no element in the DP *the NPs possible* contains a universal quantifier in its lexical semantics, and this absence leads to

¹³ In postnominal position, recall that implicit relative adjectives always scope over prenominal adjectives. This fact can also be squeezed out of the DegP analysis, a topic we address fully in §4.4.


an unsatisfiable presupposition (unsatisfiable under a commonsense understanding of what it means for something to be “possible”). Specifically, the constituent $[_{DP} \text{ the } [_{AP} \text{ possible } [_{CP_{\blacktriangle}} \dots \text{TO VERB} \dots]]]$ $NP_{[PL]}$ presupposes that there is a unique element of $[[*NP]]^{\textcircled{a}}$ that is verbed in some accessible world. The corresponding structure with *necessary*, however, presupposes the existence of a unique element of $[[*NP]]^{\textcircled{a}}$ that is verbed in *every* accessible world. This aligns with intuitions about what, e.g. *I brought the tools necessary* means. The analysis of implicit relative readings in definite descriptions is then applied to definites containing *only* (e.g. *the only tools possible*), a construction that has not been identified or analyzed in the literature to date. Finally, in §4.4, I argue that the distribution of direct versus indirect readings of modal adjectives in English follows without further stipulation from the hypotheses advanced in this chapter and in Chapter 2. In other words, the existence or non-existence of two distinct syntactic domains for attributive modifiers is logically independent from the question of why direct and indirect readings of modal adjectives distribute the way that they do in English.

4.4.1 Deriving implicit relative readings in universals

While the emphasis of this analysis is mostly on the syntactic and semantic relationship between D, Deg, AP, and \blacktriangle , and less so with DP-external processes, for concreteness I assume that the SILENT verb below CP_{amt} is licensed under identity with the matrix verb.¹⁴ I do not assume raising of the object to be necessary for the resolution of *possible*'s complement, though quantifier raising (QR), type-raising or something functionally equivalent will inevitably be required to resolve the type mismatch between the $\langle e, \langle e, t \rangle \rangle$ -type verb and the $\langle \langle e, t \rangle, t \rangle$ -type direct object.

I propose the following analysis of indirect (= implicit relative) readings of modal attributive adjectives in universal nominals, using the item *possible* in the noun phrase *I brought every tool possible* as an example. The structure of the target noun phrase in 4.66 is broken into two parts: the infinitive amount/degree CP (in b.), and the DP structure above that CP (in a.). The complement of *possible* is $\langle e, t \rangle$ -type but contains a free degree variable *d*. This variable is bound via abstraction *above* the modal adjective, yielding $\langle d, \langle e, t \rangle \rangle$ as the type for the constituent $[_{AP} \text{ possible } \blacktriangle]$.

¹⁴ This may run into problems in ditransitive sentences, where what is elided is not even a constituent, much less a sole verb.



The interpretation of CP_{amt} is the following $\langle e, t \rangle$ -type expression.

$$(4.67) \quad \llbracket \text{CP}_{\text{amt}} \rrbracket = \lambda x [d\text{-many}(x) \wedge \mathbf{bring}(x)(\text{PRO})]$$

Possible then composes with CP_{amt}, after which the free degree variable is abstracted over. This yields a $\langle d, \langle e, t \rangle \rangle$ -type expression as the semantics for AP₁, precisely the input type for MAX (recall that the presence of MAX in amount relative clauses is motivated by the arguments of Grosu & Landman (1998)):

$$(4.68) \quad \llbracket \text{AP}_1 \rrbracket = \lambda d [\lambda x [\diamond [\mathbf{bring}(x)(\text{PRO}) \wedge d\text{-many}(x)]]]$$

The Deg operator MAX then applies to this higher AP. As can be seen from the lexical entry for MAX in 4.69, adapted and simplified from Grosu & Landman's (1998), its semantics are quite similar to those of standardly assumed meanings for superlative *-est*.

$$(4.69) \quad \llbracket \text{MAX} \rrbracket = \lambda P_{\langle d, \langle e, t \rangle \rangle} [\lambda x [\exists d [P_d(x) \wedge \forall y [P_d(y) \rightarrow y = x]]]]$$

Application of the maximalization operation to $\llbracket \text{AP}_1 \rrbracket$ yields

$$(4.70) \quad \llbracket \text{Deg}' \rrbracket = \llbracket \text{MAX} \rrbracket (\llbracket \text{AP}_1 \rrbracket) = \\ \lambda x [\exists d [(\diamond [\mathbf{bring}(x)(\text{PRO}) \wedge d\text{-many}(x)]) \wedge \forall y \neq x [\neg (\diamond [\mathbf{bring}(y)(\text{PRO}) \wedge d\text{-many}(y)])]]]]$$

Since the meaning of Deg' has type $\langle e, t \rangle$ —the same as a noun modulo intensionality—it can compose with NP via intersective modification (predicate modification), yielding the meaning for DegP shown in 4.71:

$$(4.71) \quad \llbracket \text{DegP} \rrbracket = \mathbf{PM}(\llbracket \text{Deg}' \rrbracket)(\lambda x_1 [\mathbf{tool}(x)]) = \\ \lambda x [\mathbf{tool}(x) \wedge \exists d [(\diamond [\mathbf{bring}(x)(\text{PRO}) \wedge d\text{-many}(x)]) \wedge \\ \forall y \neq x [\neg (\diamond [\mathbf{bring}(y)(\text{PRO}) \wedge d\text{-many}(y)])]]]]$$

To finish the DP structure, we integrate *every*'s standard meaning:

$$(4.72) \quad \llbracket \text{DP} \rrbracket = \llbracket \text{every} \rrbracket (\llbracket \text{DegP} \rrbracket) = \\ \lambda Q [\forall x [\mathbf{tool}(x) \wedge \exists d [(\diamond [\mathbf{bring}(x)(\text{PRO}) \wedge d\text{-many}(x)]) \wedge \\ \forall y \neq x [\neg (\diamond [\mathbf{bring}(y)(\text{PRO}) \wedge d\text{-many}(y)])]]] \rightarrow [Q(x)]]]]$$

Zooming out, the matrix subject and verb are incorporated into the structure via something like QR.

$$(4.73) \quad [_{\text{TP}} [_{\text{DP}} \text{I}] [_{\text{T}'} \text{PST}] [_{\text{VP}} \text{bring} [_{\text{DP}} \text{every tool MAX} \\ \lambda d \text{ possible } \lambda x [\text{FOR PRO TO BRING } d\text{-MANY } x]]]]]$$

The truth-conditions for *I brought every tool possible*, then, are the following:

$$(4.74) \quad \forall x (\mathbf{tool}(x) \wedge \exists d [(\diamond [\mathbf{bring}(x)(\text{PRO}) \wedge d\text{-many}(x)]) \wedge \\ \forall y \neq x [\neg (\diamond [\mathbf{bring}(y)(\text{PRO}) \wedge d\text{-many}(y)])]]] \rightarrow \mathbf{brought}(x)(\mathbf{me})) \\ \text{'For any } x, \text{ if } x \text{ is the maximal tool-sum it was possible for me to bring, then I brought } x\text{'}$$

This analysis extends to postnominal implicit relative *necessary*, as the reader may verify.

A nice consequence of this analysis is that prenominal implicit relative adjectives can be assigned the same structure—the only difference is that in this case the NP specifier of Deg will appear to the left instead of to the right of Deg. Because the hierarchical structure is the same, though, so will be the truth-conditions.

$$(4.75) \quad [_{\text{DP}} [_{\text{D}} \text{every}] [_{\text{DegP}} [_{\text{Deg}'} [_{\text{Deg}} \text{MAX}] [_{\text{AP}_1} \lambda d [_{\text{AP}_0} [_{\text{A}} \text{possible}] \text{CP}_{\text{amt}}]]]] [_{\text{NP}} \text{tool}]]]$$

This analysis derives the locality restriction: assuming the conclusion of Chapter 3:§3 that attributive adjectives can only adjoin to NP or N, any additional adjectives inserted into 4.75 will linearly follow *possible*. And there is no appropriate adjunction site for additional adjectives within the Deg' constituent. In comparison to Romero's (2013) and Schwarz's (2005) analyses, mine is unique in deriving both the locality restriction and the fact that postnominal adjectives can receive implicit relative readings in the first place.

4.4.2 Implicit relative readings with superlatives ('modal superlatives')

A brief refresher on degree-based semantics: equative and comparative sentences like 4.76 are paired with structures as below, where the morphemes *as* and *-er* are analyzed as heads of the DegP projection.¹⁵

¹⁵ *Mary* is labeled as XP here to remain neutral as to whether there is elided material surrounding this DP.

- (4.76) a. John is as tall as Mary.
 [TP John [VP is [DegP [Deg' [Deg as] [AP tall]] [PP [P as] [XP Mary]]]]]
- b. John is taller than Mary.
 [TP John [VP is [DegP [Deg' [Deg -er] [AP tall]] [PP [P than] [XP Mary]]]]]

Such structures are interpreted as statements about relations between degrees; here the relevant degrees would be the respective heights of John and Mary. The following semantics can then be assigned to *-er* and *as* (adapted and simplified from Seuren (1973); Bierwisch (1989); and Heim (2000)).

- (4.77) a. $\llbracket -er \rrbracket = \lambda D_{\langle d, \langle e, t \rangle \rangle} [\lambda x [\lambda y [\exists d [D_d(y) \wedge \neg D_d(x)]]]]$
 b. $\llbracket as \rrbracket = \lambda D_{\langle d, \langle e, t \rangle \rangle} [\lambda x [\lambda y [\forall d [D_d(x) \rightarrow D_d(y)]]]]$

Combination of comparative or equative morphology with a positive-form $\langle d, \langle e, t \rangle \rangle$ -type adjective yields a transitive relation on individuals.¹⁶

Superlative *-est* can be analyzed in a similar fashion, but with two crucial differences: first, superlative adjectives, whether analytic (*most expensive*) or synthetic (*cheapest*), should denote *intransitive* properties of individuals; and second, the interpretation of a superlative is always relative to some comparison class of potentially unbounded size (in contrast to comparatives). Because of comparison class relativity, we introduce an additional $\langle e, t \rangle$ -type argument place into *-est*'s semantics, which represents the comparison class (hence the variable name 'C'). The following three-place semantics *-est* is adapted from Heim 1999:2.

- (4.78) $\llbracket -est \rrbracket = \lambda C_{\langle e, t \rangle} [\lambda D_{\langle d, \langle e, t \rangle \rangle} [\lambda x [\exists d [D_d(x) \wedge \forall y [(C(y) \wedge y \neq x) \rightarrow \neg D_d(y)]]]]]$

As per the discussion above, I assume that in superlative DPs, D selects for a DegP complement, headed by *-est* (or *most* in the case of analytic superlatives). This means that the direct object in *John climbed the highest possible mountain*, for example, will have the following structure:

- (4.79) [DP [D the] [DegP [Deg' [Deg -est] [AP [A possible] λx [CP \blacktriangle FOR PRO TO CLIMB x]]]]
 λd [NP [AP [A high] d] [NP [N mountain]]]]]

Because the syntax contains the overt degree head *-est*, no covert MAX operator enters into the derivation as it does in universals. The compositional semantics for the structure in 4.79 works out to the *e*-type term

¹⁶ The lexical semantics of equative *as* intuitively differs from that of *-er* only in requiring that every degree holding of one individual argument also holds of the other.

in 4.80. Integrating the rest of the matrix clause (*John climbed ...*), we then get the formula 4.81 (NB: the notations P_d and $d-P$ are equivalent to $P(d)$).

(4.80) $\llbracket \text{DP} \rrbracket = \iota x [\exists d [(d\text{-high}(x) \wedge \text{mtn}(x)) \wedge$

$\forall y [(\diamond \text{climb}(y)(\text{PRO}) \wedge y \neq x) \rightarrow \neg(d\text{-high}(y) \wedge \text{mtn}(y))]]]$

‘the unique d -high mountain x such that anything other than x that is possible to climb is either not a mountain or else is less than d -high’

(4.81) John climbed the highest mountain possible. \implies **climbed**_{4.80(john)}

Adopting the kind of constituency I have above creates a difficulty in the derivation of postnominal modal superlatives. Roughly, if we assume the following syntax, then we are required to assume that *-est* is interpreted in the Deg position but is pronounced along with *high* in its surface position. If linear order is read left-to-right from the tree, we would predict *the high mountain -est possible* to be the pronunciation of the target noun phrase.

(4.82) a. John climbed the highest mountain possible.

b. $[\text{DP the } [\text{DegP } \lambda d [\text{NP } [\text{AP } [\text{A high}]d] [\text{NP } [\text{N mountain}]]] [\text{Deg}' [\text{Deg } \text{-est}]$

$[\text{AP } [\text{A possible}] \lambda x [\text{CP}_\blacktriangle \text{FOR PRO TO CLIMB } x]]]]$

I will simply assume that some process applies to (4.82b) and yields the observed linear order of morphemes, setting aside for future investigation any potential side-effects of this decision. I am not sure if this kind of assumption has undesirable consequences other than that it is slightly inelegant. Such an ad hoc appeal, I would say/hope, is on a par with stipulations made in competing theories. In the present analysis, though, this problem arises only for *-est*; the analyses of universals and definites are unaffected.

Finally, some remarks are in order regarding the relationship of the above derivations to Romero’s (2013) quite similar one for the same construction. On this proposal, three-place *-est* is crucially used, whereas Romero argues that two-place *-est* is superior. However, her arguments to this effect are based upon the assumption that (comparative) superlatives involve a non-standard use of the definite article in which *the* is interpreted as an indefinite determiner. On my account, by contrast, I maintain a uniform semantics for the definite article. As far as I can tell, this patches the problem that would exist for a three-place *-est* analysis. My account also interprets *-est* in (roughly) its base position, requiring no movement-based scoping mechanism. The comparison class of *-est* is instead fixed precisely when \blacktriangle is resolved. At the same

time, I do follow Romero in analyzing [_{AP} possible ▲] as the overt expression of *-est*'s comparison class. The extensions of the degree-based semantics to universals above and to definites below are innovations of the present work.

4.4.3 Definites: why are *possible* and *necessary* different?

Recall the observation from §3 that *necessary* but not *possible* can receive an implicit relative reading in ordinary definite descriptions:

- (4.83) a. I brought with me the necessary tools. ✓*direct*, ✓*implicit relative*
 b. I brought with me the tools necessary. ✗*direct*, ✓*implicit relative*
- (4.84) a. I brought with me the possible tools. ✓*direct*, ✗*implicit relative*
 b. *I brought with me the tools possible. (N/A)

Focusing on the unambiguous postnominal construction, I propose that the direct object *the tools necessary* of (4.83b) has the syntactic structure in 4.85. The most salient difference between the structure in 4.85 and the analyses of universals and superlatives is the lack of a DegP layer between D and NP. The motivation for this omission is semantic in nature, as I will explain at the conclusion of this (sub)section.

- (4.85) [_{DP} [_D the] [_{NP₁} [_{NP₀} PL [_N tools]] [_{AP} [_A necessary] [_{CP_{amt}} λx [[_C FOR] [_{TP} [_{DP} PRO] [[_T TO] [_{VP} BRING *d*-MANY *x*]]]]]]]]]

The derivation of the semantics of (4.83b), based upon the structure in 4.85, runs in the following way (for readability we integrate modal reference via \Box/\Diamond instead of $\forall w/\exists w$): the covert CP complement ▲ (=CP_{amt}) of *necessary* is an amount relative clause, but because of the lack of a DegP layer in plain definites, no operator binds the degree variable inside ▲ and so it remains free. ▲ is therefore of type $\langle e, t \rangle$ and can serve as the argument of *necessary* (assuming a polymorphic semantics parallel to *possible*'s in 4.32):

- (4.86) $\llbracket \text{AP} \rrbracket = \lambda x [\Box [d\text{-many}(x) \wedge \text{bring}(x)(\text{PRO})]]$

The meaning in 4.86 for the constituent [_{AP} necessary ▲] is thus also of type $\langle e, t \rangle$ and so can compose via intersection with the NP [_{NP₀} PL [_N tools]], yielding yet another property-denoting expression as the meaning of NP₁:

- (4.87) $\llbracket \text{NP}_1 \rrbracket = \lambda x [*tools(x) \wedge \Box [d\text{-many}(x) \wedge \text{bring}(x)(\text{PRO})]]$

Finally, 4.87 serves as the argument to *the*, resulting in the *e*-type expression in (4.88a) for the meaning of *the tools necessary*. We spell out the corresponding expression with world variables in (4.88b), which is additionally reduced from type $\langle s, e \rangle$ to *e* via evaluating it at the actual world @.

- (4.88) a. $[[\text{DP}]] = \iota x[*\text{tools}(x) \wedge \Box[d\text{-many}(x) \wedge \text{bring}(x)(\text{PRO})]]$
 b. $[[\text{DP}]] = \iota x[*\text{tools}_@ (x) \wedge \forall w[@Rw \rightarrow d\text{-many}_w(x) \wedge \text{bring}_w(x)(\text{PRO})]]$

It is perhaps easiest to grasp the predicted meaning of *the tools necessary* when reference to worlds is encoded quantificationally, as in (4.88b). Suppose for concreteness that the modal flavor of *necessary* is teleological, so that *I brought the tools necessary* means something like “In order to accomplish what I need to accomplish (say, with tools), there is a set of tools *x* that I need to bring, and I brought *x*.” The expression in (4.88b) can then be paraphrased as “the unique set of tools *x* such that in every situation *w* in which I accomplish what I need to accomplish, *x* has *d*-many parts and I brought *x* in *w*.”¹⁷ This appears to be just what the target DP intuitively means: if it is necessary to bring a screwdriver and tinfoil, and if I brought both items, then regardless of what other tools I may or may not have brought, it is true that I brought the tools necessary. In this case, the sum of the screwdriver and the tinfoil would correspond to *x* in the previous paraphrase, i.e. in any alternative situation in which I bring what I need to, I bring at least these two items. Plugging in the matrix subject and verb, then, the truth conditions of (4.83b) on the teleological reading are predicted to be those specified in 4.89.

- (4.89) $[[\text{(4.83b)}]] = \text{brought}_@ (\iota x[*\text{tools}_@ (x) \wedge \forall w[@Rw \rightarrow d\text{-many}_w(x) \wedge \text{bring}_w(x)(\text{PRO})]]) (\text{me})$
 ‘I brought with me the unique group of tools *x* such that it is necessary to bring (at least) *x* to achieve what I need to achieve.’

These truth-conditions assert that I brought the set of tools *x* that are required for achieving what I need to.

I will now show that applying the above analytical strategy to implicit relative *possible* results in a meaning that has an unsatisfiable (in reasonable circumstances) presupposition. I propose that it is exactly this semantic deviance that accounts for the judgment that *possible* does not admit of an implicit relative reading in definite descriptions (cf. (4.51a), (4.51b)). Assuming the phrase structure in 4.85 applies to *possible* and other modal adjectives as well, we generate the structure and truth-conditions for 4.90 in (4.90a)

¹⁷ Note that the $d\text{-many}_w(x)$ conjunct is just an artifact of the degree-based analysis of \blacktriangle ; as far as I can tell, this conjunct is harmless so long as we assume that if *x* has *n* parts in @, then *x* has *n* parts in all alternative worlds/situations in which it exists. Again, issues of trans-world identity reach beyond the goals of this dissertation.

and (4.90b), respectively. Again, we use quantification over worlds in the object language to represent modal reference; here *possible* introduces an existential world-quantifier.

(4.90) */#I brought with me the tools possible.

- a. [DP [D the] [NP₁ [NP₀ PL [N tools]] [AP [A possible]
 $[\text{CP}_{\text{amt}} \lambda x [[C \text{ FOR}] [\text{TP} [\text{DP} \text{ PRO}] [[T \text{ TO}] [\text{VP} \text{ BRING } d\text{-MANY } x]]]]]]]$
- b. **brought**_@($\iota x[*\text{tools}(x) \wedge \exists w[@Rw \wedge d\text{-many}(x) \wedge \text{bring}_w(x)(\text{PRO})]]$)(**me**)

Consider now what statement (4.90b) says, setting aside x 's numerosity: "I brought the unique set of tools x such that for some situation in which I accomplish what I need to, I bring x ." The presupposition of (4.90b) is therefore that there is only a single set of tools that is possible for me to bring. This contradicts the strong intuition that 4.90 carries with it no such implication. Instead, 4.90 means something like "I brought with me the unique set of tools x such that for every atomic part x_n of x , there is some situation in which I accomplish what I need to and in which I bring x_n ." In other words, what is missing here is a universal quantifier over subparts of x . The following formula, for example, accurately captures what 4.90 might mean if it were a grammatical sentence:

(4.91) **brought**_@($\iota x[*\text{tools}(x) \wedge \forall y < x[\exists w[@Rw \wedge \text{bring}_w(y)(\text{PRO})]]]$)(**me**)

But because no quantifier akin to $\forall y < x$ is present in any of the lexical items in **the tools possible*, the stronger, implausible presupposition is generated for this noun phrase. A quantificational device over subparts is not required to generate a coherent meaning for *I brought the tools necessary*, since the truth of this statement requires that there is a specific set of tools that I bring in *every* situation of the relevant kind.

Further support for the proposed analysis of *possible* versus *necessary* comes from DP-internal *only*. It was observed above that *possible* can receive an implicit relative interpretation in the frames [DP the only [NP N --]] and [DP the only [NP -- N]]. For example, there is an implicit relative interpretation of *possible* in *I brought the only tools possible*. Provided a standard, scalar semantics for *only*, the meaning we derive for this sentence has no unsatisfiable presuppositions, in contrast to the same sentence with *only* omitted. Comparing again (4.90b) with 4.91, it is evident that the presence of a universal quantifier is required to capture the meaning of the relevant sentence. *Only*, I suggest, provides a suitable quantifier. The first step towards showing this is to define a lexical entry for *only*; the following is a presuppositional version of Rooth's (1985) denotation:

$$(4.92) \quad \llbracket \text{only} \rrbracket = \lambda P_{\langle e,t \rangle} [\lambda x : [P(x)]. [\forall y \neq x [\neg P(y)]]]$$

Assuming the bracketing of *the only tools possible* is strictly right-branching, then this phrase's derivation is parallel to that of *the tools possible*, differing only in that before NP composes with D, *only* operates on NP. This yields the following term:

$$(4.93) \quad \llbracket \text{the only tools possible} \rrbracket \\ = \iota x : [*tools(x) \wedge \exists w [@Rw \wedge d\text{-many}(x) \wedge \mathbf{bring}_w(x)(\text{PRO})]] \\ \cdot [\forall y \neq x [\neg [*tools(y) \wedge \exists w [@Rw \wedge d\text{-many}(y) \wedge \mathbf{bring}_w(y)(\text{PRO})]]]]$$

The DP in 4.93 refers to the unique x such that for any y other than x , if y is a set of tools, then in all situations w in which I achieve what I need to, either y is smaller than x , or else I do not bring y in w . Furthermore, this DP presupposes that, i.e. is defined iff, the following condition holds: x is a set of tools and it is possible to bring x ; or equivalently iff x is a set of tools that is brought in some situation in which I achieve what I need to.

The difference between the felicitous 4.93 and the deviant (4.90b) is now evident: the latter refers to the unique set x of tools that it is possible to bring, while the former presupposes that x is a set of tools possible to bring, and refers to the unique such x with the property that no other set of tools as large as x is possible to bring. The introduction of *only* therefore shifts what is asserted to be unique: without *only*, the presupposition is that there is a unique set of tools possible to bring (an implausible situation). But with *only*, the uniqueness condition applies to the asserted component, and the presupposition is weaker, namely that x is a set of tools possible to bring (and not *the* set of tools). Once again, we have seen that the presence of a genuine universal quantifier—in this case provided by *only*—is required for the licensing of implicit relative readings of modal adjectives.

To conclude this section, we return to the hypothesized syntactic difference between definite DPs with implicit relative adjectives on the one hand, and superlative and universal DPs on the other (we leave unaddressed the syntax of definites containing *only*). The syntax of plain definites differs from the syntax of modal superlatives and universals in lacking a DegP layer between D and NP. The semantic motivation for this hypothesized difference is as follows: if a DegP layer were present in the syntax of plain definites, then some kind of covert degree operator would have to be present in the Deg position. The possibilities might include the maximality operator MAX, a minimality operator SUFF, and probably a host of other

operators proposed in the literature. Introducing such operators into the syntax, however, complicates truth-conditional representations. Consider the meaning we would derive for *the tools MAX necessary* if we assumed that DegP is projected by the MAX operator in simple definites.

$$(4.94) \quad \iota x[*\mathbf{tool}(x) \wedge \exists d[(\Box(d\text{-many}(x) \wedge \mathbf{bring}(x)(\text{PRO})))$$

$$\wedge \forall y < x[\neg(\Box(d\text{-many}(y) \wedge \mathbf{bring}(y)(\text{PRO})))]]]$$

‘the unique set of *d*-many tools *x* such that *x* is necessary to bring, and such that for any subpart *y* of *x*, there is a situation in which I achieve what I need to and in which either *y* is less is than *d*-large, or else *y* is not brought’

While this denotation may not lead to blatant contradictions, it includes conditions that are either semantically vacuous (assuming the size of a set is invariant across worlds) or else involve cross-world binding configurations (e.g. consider the scopes of $\exists d$ and $\forall w$). The simpler truth-conditions delivered by the syntax in 4.85, on the other hand, correspond in a straightforward way to intuitions about what it means for *I brought the tools necessary* to be true.

Perhaps the most interesting generalization to be made from this section is that implicit relative readings of modal adjectives always require the presence of a universal quantifier of some kind. The reasons for this requirement appear to vary across constructions, but the existence of the requirement seems general. This perspective allows us to explain the novel observation that *necessary* but not *possible* can have an implicit relative reading in plain definite descriptions. Another important conclusion we have been led to is that definites can license implicit relative interpretations at all—this fact reveals a deep connection between the semantics of the implicit relative construction and the semantics of amount relatives; more specifically the previously mysterious determiner restriction can now be understood as just one instantiation of a more general determiner restriction that exists for all amount/degree relative clauses.

4.4.4 Deriving the Bolinger distribution without a two-tiered DP

To recap, the direct versus indirect/implicit relative alternation in modal attributive adjectives has syntactic properties characteristic of other Bolinger contrasts. The basic distribution of readings of *possible* as an attributive modifier are summarized in (i)-(ii). Points (iii)-(iv) indicate which interpretations are possible when *possible* appears twice in a single DP, the environment in which the syntax-sensitivity of this ambiguity is most evident.

(4.95) *Prenominally, ambiguous*

(i) [DP D *possible*_{IR/*direct} N]

(4.96) *Postnominally, unambiguously implicit relative*

(ii) [DP D N *possible*_{IR/*direct}]

(4.97) *When the two readings co-occur, modal always “outscopes” direct*

(iii) [DP D *possible*_{direct} N *possible*_{IR}] (IR > direct, *direct > IR)

(iv) [DP D *possible*_{IR} *possible*_{direct} N], but *[DP D *possible*_{direct} *possible*_{IR} N]

In this chapter I have argued that these generalizations can be derived from three basic ingredients: a well-motivated but simple theory of attributive modification and DP structure (Chapter 2); a concrete understanding of the selectional and semantic properties of *possible* (§3); and finally, a small handful of independent facts about the internal syntax and semantics of degree/amount relative clauses. Therefore, the distribution of direct and indirect readings as schematized in 4.95-4.97 does not require, or even strongly suggest, that two distinct domains of modification in DP matter here (though this kind of split-DP hypothesis is of course *consistent* with the analyses I have developed). Let us now see why this is in detail, point by point.

Point (i) states that in prenominal position, *possible* can receive both the direct and the indirect reading. The former is achieved by adjoining [_{AP} [_A *possible*]] to the head noun, which can be done at the N-level or at the NP-level. The derivation of the latter reading is given in the preceding subsections (§4.1-4.3).

Point (ii) states that in postnominal position, *possible* has only the implicit relative reading. Presumably the absence of a direct reading is just a special case of English’s general (albeit violable) tendency to disallow bare adjectival modifiers in post-noun position. While I know of no systematic theory of which adjectives can appear postnominally, the generalization appears to be roughly that the more syntactic structure a modifier has, the more likely it is to be permissible in post-noun position. For example, it is well known that whether an adjective takes a complement can affect its grammaticality as a postnominal modifier (see e.g. Abney 1987):

(4.98) a. the proud father
b. *the father proud

(4.99) a. the father [_{AP} proud of his son]
b. *the [_{AP} proud of his son] father

When bare adjectives are grammatical postnominally, as in *the stars visible*, *every man blessed*, or *the issues identified*, they tend to have a certain degree of morphological complexity, or else are participial forms. Whatever the reason for these tendencies, it appears that bare *possible* does not make the cut.

On the other hand, I hypothesize that implicit relative *possible* can appear in the postnominal position precisely because it is underlyingly the only pronounced element in a genuine (albeit covert) relative clause. And since relative clauses are uniformly postnominal in English, it follows that implicit relative *possible* can appear in this position.

Point (iii): When *possible* co-occurs pre- and postnominally, the postnominal instance (which is always indirect given point (ii)) necessarily has “scope” above the prenominal, direct instance in the sense that the postnominal instance is the surface constituent that takes the noun modified by the prenominal instance as its semantic argument.¹⁸ This generalization can be seen in 4.100, which has only the a. reading.

(4.100) Mary interviewed every possible_{direct} candidate possible_{indirect}.

✓ *indirect* > *direct*: ‘every [potential candidate] possible to interview’

✗ *direct* > *indirect*: ‘everyone who might be this: a candidate possible to interview’

The DegP analysis of implicit relative readings accounts for this fact via the assumption that D selects for DegP in this context, and not for NP. In the structure 4.101 for 4.100, the unit consisting of postnominal *possible* and MAX takes the entire NP *possible candidate* as its semantic argument. This is necessarily the case if prenominal *possible* can only adjoin to N or NP.

(4.101) Mary interviewed [DP [D every] [DegP [NP possible candidate]

[Deg' [Deg MAX] [AP [A possible] [CP ...]]]]]

A structure like 4.102 would be required to generate the opposite scoping. But again, 4.102 adjoins *possible* to DegP, which is not permissible assuming the claims of Chapter 2.

(4.102)*Mary interviewed [DP [D every] [DegP [AP possible] [DegP [NP candidate]

[Deg' [Deg MAX] [AP [A possible] [CP ...]]]]]]]

A final brief observation about point (iii): something even more general can be stated about the “scope” of implicit relative adjectives. They appear to scope over *all* other attributive modifiers. Because implicit

¹⁸ I do not consider the possibility of both instances having the implicit relative reading. It is unclear whether this interpretation of e.g. *every possible candidate possible* is available, probably due to its difficulty to parse.

relative adjectives are semantically intersective, this is most clearly shown with non-intersective adjectives, which generate distinct truth-conditions depending upon their hierarchical relationship to other modifiers. So consider the adjectives *former* and *alleged*, which may apparently scope above or below postnominal modifiers. For example, (4.103a) may refer to a former student who is located at Ohio State or an alum of Ohio State. Similarly one may or may not presuppose that the referent of *the alleged criminal who murdered my father* murdered my father, depending upon whether *alleged* adjoins directly to *criminal* to the exclusion of the relative clause, or whether it adjoins to the unit consisting of the noun and the relative clause.

- (4.103)a. a former student at Ohio State
 b. the alleged criminal who murdered my family

With the flexibility of these items' adjunction sites in mind, consider now how *former* and *alleged* interact scopally with implicit relative *possible*. The examples in 4.104 are unambiguous. In all cases, implicit relative *possible* necessarily scopes above the non-subjective adjective—including prenominal indirect *possible* ((4.104c)). This is expected on the DegP analysis.

- (4.104)a. Mary interviewed every former candidate possible.
 “every ex-candidate it was possible to interview”
 #“every *x* such that *x* used to be a candidate it was possible to interview”
- b. Mary interviewed every alleged candidate possible.
 “every *x* such that *x* was alleged to be a candidate it was possible to interview *x*”
 #“every *x* such that it is alleged that: *x* is a candidate it was possible to interview *x*”
- c. Mary interviewed every possible candidate possible.
 “every potential candidate it was possible to interview”
 #“every *x* such that *x* is possibly the following: a cand. it was possible to interview”

Thus, Larson's (2000a) observation that postnominal *possible* must scope above prenominal *possible* actually has little to do with the lexical item *possible*, and more to do with the syntactic structure of DPs containing implicit relative adjectives.

Point (iv) states that if *possible* is iterated prenominally, direct *possible* cannot precede implicit relative *possible*. This generalization follows from the discussion of point (iii) above. In particular, interpreting the first *possible* as direct in the string *every possible possible candidate*, we would need to adjoin it above

DegP, which again is not permitted by the set of assumptions from Chapter 2.

To conclude the point-by-point summary of the analysis: implicit relative *possible* takes very wide scope inside DP—so wide that if it is prenominal it must (almost) always appear adjacent to the determiner.¹⁹ This means that no modifiers can intervene between the determiner and prenominal implicit relative *possible*, including—as a special case—direct *possible*. And since direct *possible* is sometimes outscoped by postnominal modifiers, it is not surprising that postnominal indirect *possible* outscopes it, too.

4.5 Discussion and conclusion

In this chapter I have argued that the basic distributional properties of the direct/implicit relative ambiguity in modal adjectives can be captured by (i) defining a single polymorphic semantics for modal adjectives, which allows them to compose with constituents of various logical types; and (ii) postulating a silent amount/degree relative clause complement \blacktriangle for modal adjectives with indirect (implicit relative) interpretations. Existing approaches to this ambiguity have assumed more exotic formalisms in order to derive the distribution of readings, such as positing a lexical ambiguity in modal adjectives or placing the burden of explanation on the existence of two distinct structural tiers for DP-internal modifiers. A problem with the latter approach is that it stipulates where and when direct and indirect readings will arise, on a language-to-language basis without appeal to known interpretive properties of noun phrases. In my approach, the lexical semantics and syntactic subcategorization properties of specific items instead determine when and where such items will receive which kind of interpretation. I have also introduced some novel facts about implicit relative readings in English, and shown that the analysis I propose can account for them in a way that previous theories cannot. This approach, in particular component (ii), derives a puzzling determiner restriction on implicit relative readings in a motivated and empirically accurate fashion.

¹⁹ I say ‘almost’ because the following appears to admit of an implicit relative interpretation: *Mary interviewed every single possible candidate*. Explaining why this might be the case is an area of current investigation.

Chapter 5

Related issues and applications

5.0 Introduction

This chapter discusses a variety of issues peripherally related to the primary subject matter of the dissertation. Each section is more-or-less self-contained. The first section revisits the status of the type-mismatch repair mechanism introduced in Chapter 4. I show that while somewhat stipulative, the usefulness of **TMAP** extends beyond its motivating data (non-restrictive adjectives) in providing a natural analysis of certain kinds of modified proper names. §1 also provides a more general treatment of modified names, according to which names are systematically polysemous between referential and what I call “stage-predicate” interpretations. I tentatively suggest some further applications of this approach, e.g. to modifying numerals (e.g. *John’s three brothers*).

The second section is about subtriggering, or licensing by modification. I sketch a discourse-based analysis of Dayal’s (1995) essential/accidental ambiguity in relative clauses, which is shown to explain a set of facts about *any* in positive contexts that Dayal’s semantically-based analysis cannot. Without reaching firm analytical conclusions, I also introduce some new data about subtriggering with modal adjectives that cannot be explained by current theories.

Finally, §3 examines scalar-like manner implicatures that are systematically generated by restrictive but not non-restrictive modifiers in certain contexts. I sketch three ways that this asymmetry can be accounted for. The first two rely on two basic principles: first, a non-restrictive adjective in a non-definite nominal generates a universal/generic presupposition (Chapter 3:§3); and second, that the strength of (structural) alternatives must be computed blindly in the sense of Magri (2009). The final analysis sketch attributes the difference to an interaction between the semantics of focus and default syntax-prosody correspondences in

English. The discussion thus has relevance for the theory of modification, for the theory of implicature, and in the case of the last analysis sketch, for the syntax-phonology interface.

While in none of the discussions here are firm conclusions reached, each makes a specific theoretical point. §1's is that there is a diverse family of noun-phrase internal constructions whose composition looks a lot like non-restrictive modification; §2's is that discourse coherence relations can replace some of the work done by purely semantic mechanisms in (at least one) semantics-pragmatics interface puzzle; and finally, §3's is that properties of non-restrictive modification can give insight into the theory of implicature and potentially also into the semantics of syntax-phonology interactions.

The conclusion of this chapter is simple and so will not be granted its own section. In a sentence, I hope to establish the following: the set of assumptions and principles developed in this dissertation are plausible from a theoretical perspective, and have a variety of fruitful empirical applications elsewhere in the syntax, semantics, and pragmatics of English and other languages.

5.1 Composition problems and modified names (and beyond)

Recall from Chapter 4 the close semantic similarity between non-restrictive adjectives and appositive relative clauses: in many cases, these modifier types convey roughly the same information, and serve as suitable paraphrases of one another.

- (5.1) a. Cigarettes contain harmful carcinogens.
b. Cigarettes contain carcinogens, which are harmful.
- (5.2) a. I visit my sick mother.
b. I visit my mother, who is sick.

Non-restrictive adjectival modification and appositive modification are similar not only in the implications/inferences they generate, but also in a particular combinatory puzzle they pose. This puzzle, which I will dub *the composition problem for non-restrictive modification*, is most easily illustrated with appositives.

The composition problem for appositives has two components: under the assumption that an appositive modifier composed with a host DP forms a syntactic constituent, (i) that constituent has the same categorial and referential status as the unmodified DP; and (ii) there is a non-cancelable, side-issue entailment that the property denoted by the modifier holds of the individual denoted by the host DP. The situation is illustrated in the following example: where underlined content is understood as side-issue, how is it that the meanings

in (5.3a) and (5.3b) can compose to result in something like (5.3c)?¹ In other words, how is it that appositive content is computed but remains “invisible” to truth-conditional meaning?

- (5.3) John, who’s a nice guy, fixed my car for free.
- a. $\llbracket \text{John} \rrbracket = \mathbf{john} : e$
 - b. $\llbracket \text{who’s a nice guy} \rrbracket = \lambda x[\mathbf{nice-guy}_@ (x)] : \langle e, t \rangle$
 - c. $\llbracket \text{John, who’s a nice guy,} \rrbracket = \underline{\llbracket \mathbf{nice-guy}_@ (\mathbf{john}) \rrbracket} : t \llbracket \mathbf{john} \rrbracket : e$

The types of the modifier *who’s a nice guy* ($\langle e, t \rangle$) and of the host DP *John* (e) do not themselves clash, but composing them via function-argument application results in a t -type term for the meaning of the subject DP, namely (5.4a). This term subsequently introduces a type-clash at the point when the subject and the VP/T’ are to be composed, as illustrated in (5.4c).

- (5.4) a. $\mathbf{FA}(5.3a)(5.3b) = \mathbf{nice-guy}_@ (\mathbf{john}) : t$
- b. $\llbracket \llbracket \text{VP fixed my car} \rrbracket \rrbracket = \lambda x[\text{PAST}[\mathbf{fix}_@ (\mathbf{my-car})(x)]] : \langle e, t \rangle$
 - c. $\mathbf{FA}((5.4a) : t)((5.4b) : \langle e, t \rangle) = \#$ (value undefined)

5.1.1 Approaches to the composition problem

The solution for the composition problem that I proposed in Chapter 4 holds that appositives and non-restrictive adjectives require a common type-mismatch repair mechanism: **TMAP**, repeated here in 5.5. The idea is simply that the rough equivalence between pairs like (5.1a) and (5.2a) should be attributed to a common, underlying interpretive process. One way of formalizing such a black-box process is **TMAP**:

(5.5) **Type mismatch adjustment principle (TMAP):**

Where XP is not a clausal syntactic category and $\alpha : \langle \sigma, t \rangle$, $\beta : \sigma$,

$$\llbracket \llbracket \text{XP} \alpha \beta \rrbracket \rrbracket = \lambda x_\sigma [: \alpha(\beta)] . \llbracket \beta \rrbracket$$

¹ The partial-function treatment of the appositive here is purely for the sake of simplicity—I do not mean to claim that appositive content is presupposed in the technical sense (this expectation may arise from Heim & Kratzer’s (1998) influential partial-function analysis of presupposition). All that is relevant is that the referential meaning of the subject have type e and that it introduce a t -type term in addition somehow. Any kind of two-dimensional bifurcation in the semantics can achieve this, and assertion/presupposition is the most uncontroversial bifurcation of which I am aware. See Potts 2005 and Morzycki 2008 for two-dimensional, “conventional implicature”-based analyses to the composition problem for appositives and non-restrictive adjectives, respectively. The non-compositional nature of Potts’s (2005) system is identified in Kubota & Uegaki 2009, where a continuations-based treatment of side-issue content is defined. Barker et al. (2011) advance a more tightly constrained but also compositional theory of interaction between at-issue and side-issue content (also based on continuations).

Certain details of 5.5, such as the presupposition-assertion division, are negotiable and not of central importance. What defines the proposal is the idea that appositives and non-restrictive adjectives combine semantically with noun (phrases) via a single mechanism. In other words, whatever it is that solves the composition problem for appositives, should also solve the composition problem for non-restrictive adjectives.

Within the parameters of contemporary model-theoretic semantics (see Chapter 2), there are few analytical strategies available for the composition problem (provided that movements and silent syntactic operations are heavily constrained or eliminated altogether). Potential analyses can be grouped into three classes: (i) those that postulate a novel composition mode with the same status as predicate modification or function-argument application; (ii) those that hard-code a solution into the lexical semantics of some particular word/morpheme that appears in the relevant construction; and finally, (iii) those that posit a phonetically empty syntactic operator whose meaning guarantees the appropriate separation between asserted and non-asserted implications. The proposal of Morzycki (2008), as well as my introduction of **TMAP**, can be categorized as belonging to class (i). Potts's (2002) analysis of clausal appositives conforms to (ii). Something like option (iii) is a natural semantics if one adopts Cinque's (2010) two-domains theory of restrictive/non-restrictive adjectives (see Leffel 2011 for a proposal along these lines). Potts's (2005) comma-operator analysis of nominal appositives is also an instance of strategy (iii).

Arguments for or against any of these three strategies will most likely take the form of conceptual or theory-internal considerations, with elegance and generality as relevant factors.² For example, if one has a predilection against silent syntax, then option (iii) is untenable. In type-logical frameworks, where composition rules are heavily constrained, option (i) might seem unattractive. And finally, if one wishes to avoid positing lexical ambiguities, strategy (ii) would be dispreferred.

What the three strategies have most clearly in common is that each is tailor-made for a specific empirical problem. Thus in some sense, each can be seen as a "stipulation" without independent evidence. Ultimately, therefore, the preferred solution should be determined by (among other considerations) which one has the widest range of applications to phenomena found in natural language. The dialectic here is basically: if we are to add some stipulative mechanism to the grammar, it should be one that has useful applications in a maximal number of distinct domains. The goal of this section is to suggest a range of such applications, starting imminently with modified names and subsequently a handful of other constructions.

² Note that the pre- versus postnominal facts discussed in Chapter 4 can be seen as evidence for strategy (ii) or (iii) since the syntactic position of a covert or overt element can be exploited as an explanation for why non-restrictive interpretations are absent in some configurations. §3 revisits this issue.

5.1.2 Complex names

Modified names can be thought of as falling into two classes: those in which the modifier is appositive-like (*poor John*), and those in which the modifier is (arguably) restrictive on a set of entity-parts that is supplied by the name (*an exhausted John*). **TMAP** affords a very simple analysis of the first variety. The second variety, I will argue, is analyzed best in terms of an ontologically enriched semantics according to which names can denote predicates over spatio-temporally anchored individual-stages.

A basic puzzle about modified names is figuring out how it is that a property-denoting modifier can combine with an individual-denoting name to form either a suitable argument for a determiner (as in (5.6a)-(5.6c)), or a DP that can serve as an argument to a verb (as in the determinerless examples in (5.6d)).³

- (5.6) a. an exhausted John / the exhausted John / exhausted John /
b. The undefeated Chicago Bulls / an undefeated Chicago Bulls squad
c. the America of my dreams / the America that I love
d. seven-foot-four Shaquille O’Neal / ancient Rome

Notice that this puzzle has exactly the same shape as the composition problems discussed above for appositives and for non-restrictive adjectives. Composing the meanings of *John* and *poor* with function-argument application, for example, yields the incorrect propositional semantics in (5.7a) for *poor John*. What we intuitively want is better represented by the expression in (5.7b): *poor John* refers to John, and introduces the implication that John is poor (in the sense of pitiable).

- (5.7) Poor John has bronchitis again.

³ The remarks made in this section are not meant to apply to cases in which a proper name is literally used as a predicate. What I have in mind are examples like these:

- (i) a. There are three Johns in this class.
b. I’ve never met a John.
c. Every John I’ve met is a real jerk.

There are two main properties that distinguish the uses of names in (i) from the uses in 5.6. First, the names in 5.6 refer to specific individuals (or spatio-temporal parts thereof), while in (i) *John* is used as a predicate meaning something like “individual referred to by the name *John*.” To be “a John” is to be a person named *John*, while to be “John” is to be identical to a specific salient person (whose name is John). Second, the occurrences of *John* in (i) seem to be genuine common nouns, as is most plainly seen by the overt nominal morphology. Such morphology is absent from all names in 5.6, indicating that those occurrences directly refer, much like e.g. pronouns.

In this discussion, I will be exclusively concerned with non-predicational uses of names. In other words, we focus on examples like 5.6, assuming that the analysis of (i)-type cases is fairly simple: names are systematically polysemous between a referential usage and a predicational usage equivalent to “person/thing referred to as *X*.”

- a. $\mathcal{X}[\llbracket \text{poor John} \rrbracket] = \mathbf{FA}(\lambda x[\mathbf{poor}(x)] : \langle e, t \rangle)(\mathbf{john} : e) = \mathbf{poor}(\mathbf{john}) : t$
 b. $\checkmark[\llbracket \text{poor John} \rrbracket] = \blacksquare(\lambda x[\mathbf{poor}(x)] : \langle e, t \rangle)(\mathbf{john} : e) = \underline{\mathbf{poor}(\mathbf{john})} : t. \mathbf{john} : e$

For the construction exemplified in 5.7, then, the analytical challenge reduces to simply determining what composition mode should be filled in for the black box operation \blacksquare in (5.7b). The adjustment principle **TMAP** does exactly the work required. To see this, we simply plug in the translations of *poor* and *John* into schema 5.5 above.

$$(5.8) \quad \mathbf{TMAP}(\lambda x[\mathbf{poor}(x)] : \langle e, t \rangle)(\mathbf{john} : e) \\ = \underline{\mathbf{poor}(\mathbf{john})} : t. \mathbf{john} : e$$

TMAP delivers just the desired result: *poor John* is defined iff *John* is pitiable, and if defined has the same referential meaning as does the unmodified name *John*. This straightforward analysis in terms of **TMAP** illustrates another application for this novel composition mode. Viewed another way, it illustrates another *construction* that has something inherently in common with non-restrictive modification.

The puzzle of complex names does not end with the case in 5.7, unfortunately. In particular, the analysis in 5.8 does not extend to names with articles: the type of *poor John* in 5.8 is e , and determiners take $\langle e, t \rangle$ -type terms as their first argument. Furthermore, if the translation of a name, say *Hans*, is a term of type e , then what kind of semantics should the German noun phrase *der Hans* ‘the.MASC Hans’ receive? Assuming the definite article has type $\langle \langle e, t \rangle, e \rangle$, by what means can these two phrases compose?⁴ And it is not only the definite article that can appear with modified names: we also have phrases like *an exhausted John Smith*, which appear with *indefinite* articles. Finally, whether a modified name occurs with an *indefinite* or a *definite* article appears to have some kind of non-trivial effect on the ultimate interpretation of a DP, which can be illustrated by the minimal pair in 5.9. This contrast is somewhat paradoxical given that for any complex name $[\text{Mod Name}]$, $\llbracket [\text{Mod Name}] \rrbracket$ has the same referent as does the unmodified form $\llbracket [\text{Name}] \rrbracket$.

- (5.9) a. The incredibly tall John Smith will be a basketball star someday.
 b. #An incredibly tall John Smith will be a basketball star someday.

The contrast between (5.9a) and (5.9b) suggests the following analytical intuition: (i) the modifier $[\text{AP } \textit{incredibly tall}]$ denotes an individual-level property; (ii) the VP $[\text{VP } \textit{will be...}]$ denotes an individual-level

⁴ As far as I am aware, there are no languages that optionally allow unmodified (referential) names to appear with an *indefinite* article instead of a *definite* article. A theory of names should account for this fact.

predicate; (iii) *a* and *the* have their ordinary meanings; (iv) *John Smith* denotes a predicate over spatio-temporal stages/slices of the individual John Smith; and (v) sentence (5.9b) is odd because it predicates an individual-level property (future stardom) of a particular John-stage (in which he's tall).⁵ In (5.9a), by contrast, *the* picks out the maximal sum of John-stages in which John is tall; and since height does not change over time (idealizing a bit), *the incredibly tall John Smith* denotes the sum of *all* John-stages (which our theory should somehow guarantee is equivalent to the individual John Smith). Thus, the individual-level VP-property is predicated of the individual John Smith, and we have no individual-level/stage-level clash.

While these suggestions are somewhat unorthodox, similar explanations are seen elsewhere in the grammar. For example Dowty (1979) attributed the contrast between *John is tall/John knows French* and *#John is being tall/#John is knowing French* to the hypothesis that the progressive aspect is incompatible with individual-level predicates like *be tall*. Even more importantly, the intuition can be concretely tested by changing the matrix predicate of 5.9 to a stage-level VP. Since a stage-level property can hold of an individual (e.g. a person can be tired) as well as of an individual-stage (e.g. the me-slice right now is tired), the contrast observed in 5.9 should disappear in the context of a stage-level VP. (5.10a) and (5.10b) appear to verify this prediction, as both are reasonable and felicitous utterances and *blocked the shot* is not an exclusively stage-level property.

- (5.10) a. The incredibly tall John Smith blocked the game-winning shot.
 b. An incredibly tall John Smith blocked the game-winning shot.

The pattern appears to remain the same even when a strongly stage-level modifier like *incredibly exhausted* is used in place of *tall* (as in 5.11 versus 5.12). This indicates that it must be the distinction between *a* and *the*—and not any property of the modifier—that is responsible for the contrast in examples like 5.9 and 5.12.

- (5.11) a. The incredibly exhausted John Smith collapsed on the couch after work.
 b. An incredibly exhausted John Smith collapsed on the couch after work.

- (5.12) a. #An incredibly exhausted John Smith works too many hours.
 b. The incredibly exhausted John Smith works too many hours.

⁵ Another way to think about the data—one that attributes the deviance of (5.9b) to the content of the modifier—is in terms of Magri's (2009) theory of "blind" implicatures: because (5.9b) asserts the existence of *some* John-stage in which he's tall, there is an implicature generated to the effect that *not all* John-stages are ones in which he's tall. The deviance would therefore arise via a conflict between this implicature and the lexical individual-level status of *tall*. This idea will not be pursued in detail here.

Having established that the above intuition explains some tricky facts, we now sketch a formalized version of the statements (i)-(v) and apply them to the above data.

First, it is necessary to assume a certain kind of polysemy in proper names. In addition to the ordinary e -type referential semantics, I will assume that a proper name can denote a predicate over spatio-temporal individual stages/slices (“i-stages”). Ontologically, i-stages need be nothing more than ordinary objects in D_e ; any Ty_2 variable or constant of type e can have an i-stage as its value (for reasons that will become clear). Thus, for example, a name like *John* can have either of the following interpretations, where $R_{(t,l)}(x)(y)$ is a parameterized realization relation that holds of an individual x and a stage y at time t and location l iff y is the slice of x at t and l (more later).

- (5.13) a. $\llbracket \text{John}_1 \rrbracket = \mathbf{john} : e$
 b. $\llbracket \text{John}_2 \rrbracket = \lambda x[\exists t[\exists l[R_{(t,l)}(\mathbf{john})(x)]]] : \langle e, t \rangle$

It is important for the theory to guarantee an equivalence between the denotations in (5.13a) and (5.13b), since both are possible meanings of *John* and *John* is in no sense “ambiguous.” To the extent that a formal distinction between ambiguity and polysemy exists, the proposal is that names are polysemous, not that they are ambiguous. Thus we guarantee that an equivalence relation \equiv should hold between the two meanings of *John*. Making this distinction precise requires formalization of some requisite notions. We begin with the spatio-temporal extent of an individual, a notion from which the appropriate kind of equivalence relation can be built:⁶

- (5.14) **Definiton:** The *spatio-temporal extent* ε_x of a physical individual x whose existence is n units of time long is the set $\varepsilon_x = \{\langle t_1, l_1 \rangle, \dots, \langle t_n, l_n \rangle\}$ where for each $i \leq n$, l_i is the location of x at time t_i .

In symbols:

$$\varepsilon_x =_{\text{df}} \{ \langle t, l \rangle \mid \exists y [R_{(t,l)}(x)(y)] \} \quad (\text{where } y \text{ is an i-stage})$$

Times are ordered by the precedence relation $<$ so that for any t_1, t_2, t_3 , $<$ is anti-symmetric ($t_1 < t_2 \not\Rightarrow t_2 < t_1$), transitive ($t_1 < t_2 < t_3 \Rightarrow t_1 < t_3$), irreflexive ($t_i \not< t_i$ for all i), and dense ($t_1 < t_2 \Rightarrow \exists t [t_1 < t < t_2]$). Furthermore, every temporal extent is upper bounded: there is always a t_n such that for all t_i with $i \neq n$, $t_n \not< t_i$. And extents of different individuals need not match; i.e. x 's terminal (or initial) point t_n need not be

⁶ The extent of an individual can be seen as parallel to the notion of an event's runtime, a construct made crucial use of events-based semantic frameworks (some examples are Champollion 2010 and Krifka 1989). And runtimes are closely related to the derivative notion of *temporal extent of x* , which could be defined here as $\{t \mid \exists l [\langle t, l \rangle \in \varepsilon_x]\}$.

the same as a distinct individual y 's terminal (or initial) point.

Finally, the collection of all i -stages of an individual x , called x 's *stage-set*, is defined as in 5.15. The initial target concept of stage-equivalence can then be defined as in 5.16.

(5.15) **Definition:** The *stage-set of an individual x* is the set $\{y|\exists\langle t,l\rangle \in \mathcal{E}_x[\mathbf{R}_{(t,l)}(x)(y)]\}$.

(5.16) **Definition:** If S is the stage-set of x , then x and S are *stage-equivalent*, written $x \approx S$.

(example: $x \approx \{y|\exists\langle t,l\rangle \in \mathcal{E}_x[\mathbf{R}_{(t,l)}(x)(y)]\}$)

With this terminology in mind, the nature of the polysemy posited in 5.13 can be characterized precisely: names are polysemous between referring to an individual x and referring to (the characteristic function of) x 's stage-set. And since \approx is an equivalence relation (this can be proved), the denotation of a name contains an equal amount of truth-conditional information on each reading.

The four-place predicate R in (5.13b)-5.16 is a primitive relation inspired by Carlson's (1977b) realization relation between natural kinds and individuals. In fact, we can view Carlson's treatment of kind-terms and the present analysis of names as instances of a single, more general analysis: kinds in Carlson's (1977b) theory are the individuals of this theory, and individuals in Carlson's theory are stages here. Carlson's (1977b) use of reified natural kinds in natural language ontology has inspired parallel analyses in other domains as well; see Wilkinson 1995 and Zamparelli 1998 on the semantics of the noun *kind* and Scontras 2014 on amounts, degrees, and their relationship to one another.

With these ingredients in hand, we can compositionally derive interpretations for the contrasting sentences in 5.9. Simplifying 5.9 to the form in 5.17, lexical entries for the relevant items are listed in (5.17a)-(5.17d). Recall the maximality condition in the definition of ι from Chapter 2:§3.5, which guarantees that $\iota x[P(x)]$ denotes the sum of all (relevant) individuals satisfying P .⁷

(5.17) The/#A tall John will be a star.

- a. *Determiners:* $\llbracket \mathbf{a(n)} \rrbracket = \lambda P[\lambda Q[\exists x[P_{@}(x) \wedge Q_{@}(x)]]]$ $\llbracket \mathbf{the} \rrbracket = \lambda P[\iota x[P_{@}(x)]]$
- b. *Modifier:* $\llbracket \mathbf{tall} \rrbracket = \lambda x[\mathbf{tall}(x)]$
- c. *Name:* $\llbracket \mathbf{John}_2 \rrbracket = \lambda x[\exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]]$ (= (5.13b))
- d. *Predicate:* $\llbracket \mathbf{will\ be\ a\ star} \rrbracket = \lambda x[\exists t' > @[\mathbf{star}_{t'}(x)]]$

⁷ For example if I have three dogs in context c —named Sparky, Rags, and Fido—then *the dogs* will denote a three-part sum-individual consisting of these dogs: $\llbracket [\text{DP the dogs}] \rrbracket^c = \mathbf{sparky} \oplus \mathbf{rags} \oplus \mathbf{fido}$.

Beginning with (5.9a), *The incredibly tall John Smith will be a basketball star someday*, we have

$$\begin{aligned}
 (5.18) \quad & \mathbf{PM}(\llbracket \text{tall} \rrbracket)(\llbracket \text{John}_2 \rrbracket) = \lambda x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]] \\
 & \mathbf{FA}(\llbracket \text{the} \rrbracket)(\lambda x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]]) \\
 & \quad = \iota x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]] \\
 & \mathbf{FA}(\llbracket \text{will be a star} \rrbracket)(\iota x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]]) \\
 & \quad = \exists t' > @[\mathbf{star}_{t'}(\iota x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]])]
 \end{aligned}$$

“At some point in the future, the following individual will be a star: the sum of all tall x such that x instantiates John at some spatio-temporal point.”

Since being tall is a permanent property, the sum denoted by *the tall John* will be the sum of all of John’s i -stages. Again, assuming that the sum of all i -stages of an individual is equivalent to that individual, then what 5.18 says is that someday, John will be a star. This corresponds to what (5.9a) intuitively means.

So what accounts for the deviance of sentence (5.9b), *An incredibly tall John Smith will be a basketball star someday*? Intuitively the reason is that *will be a basketball star someday* is not the kind of property that can hold of a temporally and spatially anchored stage; instead it must hold of an individual—whether qua individual or qua maximal sum of i -stages—with a non-degenerate temporal extent. The derivation in 5.19 predicates future stardom of a single stage, thus accounting for the deviance of (5.9b) as desired.

$$\begin{aligned}
 (5.19) \quad & \mathbf{PM}(\llbracket \text{tall} \rrbracket)(\llbracket \text{John}_2 \rrbracket) = \lambda x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]] \\
 & \mathbf{FA}(\llbracket \text{a} \rrbracket)(\lambda x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]]]) \\
 & \quad = \lambda Q[\exists x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]] \wedge Q(x)]] \\
 & \mathbf{FA}(\lambda Q[\exists x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]] \wedge Q(x)]](\llbracket \text{will be a star} \rrbracket)) \\
 & \quad = \exists x[\mathbf{tall}(x) \wedge \exists t[\exists l[\mathbf{R}_{(t,l)}(\mathbf{john})(x)]] \wedge \exists t' > @[\mathbf{star}_{t'}(x)]]
 \end{aligned}$$

#“There is a spatio-temporal slice of John in which he is tall, and that slice will be a star.”

To conclude this discussion: allowing names to denote predicates of spatio-temporal individual-stages explains the effect of determiner choice (indefinite versus definite) on the interpretation of modified proper names. Names without determiners refer directly to individuals, without mediation of a stage-set.

5.1.3 Some putative cousins of non-restrictive modification

In this brief discussion I identify a handful of constructions which, on the face of it, resemble or seem related to non-restrictive modification. For each case, I provide only a short description and in some cases

preliminary suggestions about how the internal composition might work.

Close appositives are noun phrases consisting of two juxtaposed noun phrases. Some examples are *John the doctor, the eagle the bird (not the symbol)*, and perhaps even archaic proper names like *Alexander The Great*. Despite their resemblance to garden variety appositives, close apposition structures differ intonationally (no prosodic breaks between the noun phrases), as well as syntactically and semantically: while appositives can be indefinite (*John, a doctor*), close appositives cannot (**John a doctor*). Exactly what the relationship is between these two constructions is unclear, but the common thread seems to be that both involve adjunction of some phrase to a full DP (unlike other modification constructions) whose referential content is the same as the unmodified DP. Thus, it is likely that the correct analysis of close apposition will be similar in shape to the correct analysis of ordinary appositives.

A related kind of noun phrase appears in Greek—the **polydefinite construction** (sometimes referred to as “determiner spreading”). Polydefinite DPs contain multiple instances of the definite article, with attributive adjectives appearing to introduce their own articles. Interestingly, attributive adjectives are prenominal in Greek (shown in 5.20), but in the polydefinite construction, the adjective can appear postnominally if accompanied by an additional definite article (shown in 5.21). Alexiadou (2001:232) states that “pre-adjective determiner[s] [are] obligatory for post-nominal adjectives,” as illustrated in 5.21.

(5.20) Alexiadou (2001:223)

- a. to meghalo spiti
the big house
'the big house'
- b. *to spiti meghalo

(5.21) Alexiadou (2001:232)

- a. to vivlio *(to) kokkino
the book the red
'the red book'
- b. to kokkino to vivlio (Greek)

The polydefinite construction can potentially tell us much about the nature of modification, as argued extensively in Alexiadou 2001; Alexiadou et al. 2007; Alexiadou & Wilder 1998; and related contributions. Further investigation is however beyond the scope of this dissertation. The take-away point is that polydefinites provide a nice domain for testing theories of modification, especially syntactically-oriented ones.

Expressive attributive adjectives like *damn* in *the damn dog* were taken by Potts (2005,2007) to involve a kind of non-restrictive modification. Roughly, Potts argued that *the damn dog* refers to the dog, but introduces a side-issue implication to the effect that the speaker has a negative attitude towards something relating to dog(s) at the time of utterance. Should expressives receive the same treatment as other non-restrictive adjectives? The answer is unclear, but one interesting difference is that expressive adjectives appear to be incapable of restricting: not only are they ungrammatical as predicates (**The dog is damn.*), but they do not appear to admit of contrastive uses either: compare (5.22a) with (5.22b).

- (5.22) a. Go chase the big dog, not the small one.
b. #Go chase the damn dog, not the good one.

Determining the semantic scope of an expressive modifier is not trivial. An analysis is not attempted here; the point of this discussion is merely to identify the common core of expressives and other non-restrictive modifiers.

Finally, **attributive quantity modifiers**, including numerals and DP-internal *many* and *few*, have a syntax and semantics very much like that of non-restrictive adjectives, and have in fact been argued by Solt (2009) to involve the same interpretive processes as ordinary non-restrictive modifiers. Consider a sentence like (5.23a), along with a natural paraphrase of its content.

- (5.23) a. John's three brothers are tall.
b. John's brothers are tall (there are three of them).

Loosely speaking, the contribution of *three* is not to restrict the reference of the noun *brothers*—because *John's* is sufficient to do that—but rather to introduce side-issue information about the cardinality of the full DP's referent. These semantics clearly resemble non-restrictive adjectives; cf. *my sick mother*. The reader is referred to Solt 2009 for more information about this connection.

5.2 Subtriggering: modifier position and discourse coherence

In a series of papers, Dayal (1995,1998,2005) analyzed an interaction called “subtriggering,” the clearest example of which holds between the negative polarity item *any* and certain types of nominal modifiers. Subtriggering, or licensing by modification, occurs when an ungrammatical structure is rendered grammatical by the addition of a modifier, or when an additional reading of a sentence is made available by the addition

of a modifier (LeGrand 1975; Carlson 1981). Here I will argue for the existence of a novel constraint on subtriggering: the triggering modifier must establish a discourse coherence relation with some salient entailed proposition (often the one expressed by the sentence in which the modifier occurs). This added constraint is compatible with Dayal’s situation-parameter theory of subtriggering, but explains some restrictions not explained by her proposal alone. I also introduce some observations about subtriggering with implicit relative adjectives, and show why they are difficult for any existing proposal to explain.

5.2.1 *Any* in positive contexts

Any-subtriggering occurs when a normally ungrammatical *any*-headed DP is “rescued” by the addition of a relative clause. Because it is not in the scope of a modal, negation, or the habitual, *any* is ungrammatical throughout 5.24. But the relative clauses in 5.25 render these structures grammatical:

- (5.24) a. *John talked to any woman.
b. *Any man didn’t eat dinner.
c. *Any woman contributed to the fund.

- (5.25) a. John talked to any woman who would listen.
b. Any man who found a fly in his soup didn’t eat dinner.
c. Any woman who heard the news contributed to the fund.

Dayal (1995) observed that not just any modifier licenses *any* in positive contexts. To see this, observe first that restrictive/integrated relative clauses are often “ambiguous”⁸ in an interesting way. For instance the relative in 5.26 can have either of the indicated interpretations, termed by Dayal as the accidental and essential readings.

- (5.26) Every student who’s taking Mary’s class is writing a paper on polarity items.

Accidental: ‘They happened to all write a paper on polarity.’

Essential: ‘In virtue of being in her class they wrote on polarity.’

The putative grammatical difference between the two readings cannot be directly observed with the determiner *every*, but it can with *any*. To borrow a scenario from Dayal 1995, suppose professor Mary is teaching a Field Methods class and a semantics class, and all of the graduate students (and only them) are in both

⁸ Scare quotes indicate the unclear status of the opposing readings: do they differ truth-conditionally or merely pragmatically? Only in the former case could this be considered a genuine ambiguity.

classes. In the semantics curriculum, writing a paper on polarity items is obligatory. In the Field Methods curriculum, it is not. In this situation, both of the following sentences are true:

- (5.27) a. Every student who's taking Mary's semantics class is working on polarity.
b. Every student who's taking Mary's F.M. class is working on polarity.

The relative clause in (5.27a) most naturally has the essential reading whereas the relative clause in (5.27b) most naturally has the accidental reading. Both sentences, though, are grammatical, felicitous, and true. But now consider a slight modification to 5.27, in which *every* is replaced with *any*. There is a strong contrast in felicity between (5.28a) and (5.28b) when both are evaluated with respect to the above context. This contrast suggests that only relative clauses with an essential reading can subtrigger *any* (Dayal 1995).

- (5.28) a. Any student who's taking Mary's semantics class works on polarity items.
b. #Any student who's taking Mary's Field Methods class works on polarity items.

There are also structural constraints on subtriggering. Attributive adjectives are apparently unable to subtrigger in the way that relative clauses can (Carlson 1981). Compare:

- (5.29) *Context:* There was a party last night. Mary is cute and recently single, so
a. *Any guy talked to her. / Every guy talked to her.
b. *Any single guy talked to her. / Every single guy talked to her.
c. Any guy who was single talked to her. / Every guy who was single talked to her.

With respect to the licensing of *any* in positive contexts, more important than the syntactic category of a modifier appears to be its structural position: it is not only relative clauses that can subtrigger, but apparently also (certain) postnominal PP modifiers, as shown in 5.30 (Carlson 1981).

- (5.30) a. The Nazis harassed any man who wore women's clothing.
b. The Nazis harassed any man with women's clothing on.

The point Dayal takes (5.28b) to make can now be made more precise, assuming 5.29 represents a general pattern: only postnominal modifiers with essential readings subtrigger *any*. In the two subsections that follow, I explicate Dayal's (1995) analysis of these distributional facts and present an alternative interpretation of them.

5.2.2 Ambiguity versus discourse coherence in subtrigging

Dayal (1995) analyzes the accidental/essential distinction as a genuine semantic ambiguity in relative clauses. She argues that accidental readings correspond to “individual-loaded” relative clauses, while essential readings correspond to “property-loaded” relatives. The formal distinction she draws between these two kinds of relatives is that individual-loaded ones are semantically simple properties of the form $\lambda x[P(x)]$, while property-loaded relatives are predicativized nominalized properties of the form ${}^{\cup}({}^{\cap}\lambda x[P(x)])$ (where ${}^{\cap}$ is nominalization and ${}^{\cup}$ is predicativization; see Chapter 3:§3.5 for details). While any property P is extensionally equivalent to ${}^{\cup\cap}P$ (Chierchia 1984), Dayal states that the two differ in “mode of presentation.” As I understand it, the distinction can be understood as a difference in how a statement is verified: to see whether property P holds of individual x is presumably a simpler process than seeing whether ${}^{\cup\cap}P$ holds of x . The core of the analysis is then that *any* requires a property-loaded relative clause in order to be subtrigged (hence 5.28), whereas *every* can combine with both property- and individual-loaded relatives. The two readings of 5.31 according to this view are the following:

(5.31) Every student in Mary’s class works on polarity.

- a. $\forall x[[\text{student}(x) \wedge \text{in-Marys-class}(x)] \rightarrow \text{w.o.-polarity}(x)]$ (*ind-loaded*)
- b. $\forall x[[\text{student}(x) \wedge {}^{\cup\cap}(\lambda z[\text{in-Marys-class}(z)](x))] \rightarrow \text{w.o.-polarity}(x)]$ (*prop-loaded*)

While this analysis succeeds in identifying a crucial difference between *every* and *any*—namely that the latter but not the former requires a property-loaded modifier—notice that it does not address the question of *why* essential but not accidental relative clauses are capable of subtrigging *any*. I propose that the reason has to do with discourse coherence.

The basic idea that I would like to propose is that there is no semantic difference between the accidental and essential construals of relative clauses. Instead, one simply may or may not infer an explanatory-type coherence relation between the content of a matrix clause and of a relative clause contained in it. As argued in Chapter 4:§2.2 and also by Rohde et al. (2011), this kind of process occurs in the most natural interpretations of sentences like *John detests the co-workers who are arrogant and rude*: there is a relation of (perhaps) EXPLANATION that holds between the matrix implication “John detests a subset x of arrogant and rude co-workers” and the relative clause presupposition “there is a subset x of arrogant and rude co-workers of John.”

Suppose Mary devised six different test forms, to ensure that her students weren't able to cheat on the test by looking at their neighbors' answers. Suppose further that she made an error on the test forms such that everyone who received version B failed, and everyone else got at least a B+. Finally, assume that neither Mary nor the students were aware of the error during Mary's office hours, the time when disgruntled students get a chance to plead their cases to her. Sentence (5.32a) but not sentence (5.32b) is a natural description of this situation.

- (5.32) a. Mary talked to any student who was angry.
 b. #Mary talked to any student who received version B. (# in specified context)

The problem posed by these data is the following: it is in virtue of having been given test form B that some set of students talked to Mary, just as it is in virtue of being angry that they talked to her. This means that both relative clauses should admit of essential readings, which on Dayal's theory implies that both should be capable of subtriggering *any*. And yet, we see that only (5.32a) is an acceptable sentence given the contextual information. What is wrong with (5.32b), intuitively, is that *from both Mary's and the students' perspective*, having version B is no reason to complain during office hours. In other words, having been given version B is not the explanation for why a student might have gone to talk to Mary. According to this approach, the difference between (5.32a) and (5.32b) boils down to the fact that if we infer an explanatory relation between the implications of each sentence, we get the true statement that anger explains the student-Mary talking event for (5.32a) (hence its felicity), and the false statement that version B explains the student-Mary talking event for (5.32b) (hence its incompatibility with the specified context). If this approach has the right shape, then the difference between *every* and *any* can be precisely stated as

- (5.33) **Subtriggering:** For a modifier to subtrigger *any*, a coherence relation must be inferred between the content of the modifier and the assertion of the sentence in which it occurs.⁹

This way of stating the restriction on subtriggering also allows us to maintain the more theoretically conservative position that relative clauses are not genuinely ambiguous. For example, instead of positing forms like (5.31a) and (5.31b) for the two interpretations of *Every student who's in Mary's class works on polarity*, we have only a single interpretation—reflecting the intuition that this sentence is semantically unambiguous—and on top of that, we optionally infer a relation of explanation between “there is a set of students *x* in

⁹ Note crucially that inferring a coherence relation is not assumed to be a sufficient condition for the availability of subtriggering, but only a necessary condition.

Mary’s class” (existential import presupposition of *every*) and “all of *x* works on polarity” (the assertion). In case one does make this inference, we get Dayal’s property-loaded reading (5.31b); in case one does not, we get the individual-loaded reading (5.31a). This analysis eliminates the phantom operations \cap and \cup from the semantics of relativization, posits no systematic ambiguity in restrictive relatives, and has a certain intuitive appeal: different “modes of presentation” are pragmatic icing on top of the literal content of an utterance—just like conversational implicature in general.

5.2.3 A counterexample to the extra-parameter analysis (Dayal 1995,1998)

In this final subsection about subtrigging, I first summarize Dayal’s (1998) theory of the pre- versus postnominal contrast in modifiers’ ability to subtrig *any*. Following that exposition, I introduce a set of observations that cannot be accounted for on this approach. Without reaching firm conclusions, I suggest some potentially useful ways in which the data can be thought about.

Dayal (1995,1998) argues that *any* introduces universal quantifiers over individuals and situations (or situation-individual pairs), so it in a sense quantifies possible individuals. Her analysis of ungrammatical positive *any*-statements is that in, e.g. *Mary talked to any student*, the matrix predication can only hold of actual individuals, while the quantifiers introduced by *any* can take as values any possible individual; Dayal refers to this situation as a “presuppositional clash.”

Simplifying slightly, Dayal’s (2005:223-224) proposed explanation for the pre- versus postnominal contrast in subtrigging is as follows: *angry student* restricts the quantifiers introduced by *any* to situation-individual pairs consisting of situations *s* containing angry students *x*, and a sentence like *Mary talked to any angry student* asserts that Mary talked to every such *x*. In other words, *Mary talked to any angry student* means that for any relevant possibility involving an angry student *x* (where *x* may or may not be an actual individual), Mary (actually) talked to *x*; and this statement requires Mary to have actually talked to non-actual individuals. *Student who was angry*, on the other hand, restricts *any*’s domain to situation-individual pairs consisting of situations *s* containing a student, and individuals *x* that were angry in an *actual* past situation.¹⁰ The crucial difference between relative clauses and attributive adjectives that is responsible for the contrast, according to Dayal (1995,1998), is that relative clauses introduce a spatio-temporal parameter in a way that

¹⁰ In Dayal’s (2005:223) words, “A sentence like [(5.35a)] is unacceptable because there is a clash between the presupposition that the domain of quantification ranges over possible individuals and a predication that can only apply to actual individuals...Sentence [(5.36a)], however, avoids this by restricting the domain of quantification to those possible individuals who also fulfill the property in the relative clause, including the property of being in the temporal slice denoted by the temporal specification of the relative clause.”

prenominal modifiers do not. For example, the semantics of *angry* might look something like (5.34a), while that of *who was angry* would be more like (5.34b). This explanation is built upon Sadler & Arnold’s (1994) proposals outlined in Chapter 3:§4.3.

- (5.34) a. $\llbracket [\text{AP } \textit{angry}] \rrbracket = \lambda s[\lambda x[\mathbf{angry}_s(x)]]$
 b. $\llbracket [\text{CP } \textit{who was angry}] \rrbracket = \lambda s[\lambda x[\exists t < s[\mathbf{angry}_t(x)]]]$

The following formulas illustrate the truth-conditional difference crucial to the idea behind Dayal’s theory, although for clarity these formulas are simplified from those Dayal actually posits.¹¹

- (5.35) a. *Mary talked to any angry student.
 b. $\forall s[\forall x[(\mathbf{student}_s(x) \wedge \mathbf{angry}_s(x)) \rightarrow \exists t' < @[\mathbf{talk}_{t'}(x)(\mathbf{mary})]]]$
 ‘for any possible situation-individual pair s - x , if x is an angry student in s , then Mary talked to x in some actual past situation.’

- (5.36) a. Mary talked to any student who was angry.
 b. $\forall s[\forall x[(\mathbf{student}_s(x) \wedge \exists t < @[\mathbf{angry}_t(x)]) \rightarrow \exists t' < @[\mathbf{talk}_{t'}(x)(m)]]]$
 ‘for any possible situation-individual pair s - x , if x is a student in s and x is angry in some actual past situation, then Mary talked to x in some (possibly distinct) actual past situation.’

To repeat, the crux of the analysis is that relative clauses and other postnominal modifiers introduce a bindable situation variable in a way that attributive adjectives do not; *any* must bind a situation variable in its restrictor in order to avoid the presuppositional clash noted above; therefore, only (a subclass of) postnominal modifiers can license *any*.

Here I present a counterexample to this conclusion: postnominal but not prenominal implicit relative adjectives can subtrig *any*. To see this, first recall that in universal DPs, *possible* can receive a “direct” or “implicit relative” reading prenominally, but only the latter postnominally, as illustrated in 5.37. This has been taken as some to indicate that prenominal *possible* can have the structure of a reduced relative clause (which corresponds to the implicit relative reading).

¹¹ From Dayal 1998:454:

- (i) a. *Mary talked to any angry student.
 $\forall s[\forall x[(\mathbf{student}_s(x) \wedge \mathbf{angry}_s(x) \wedge C(s)) \rightarrow \exists s' > s[\mathbf{talk}_{s'}(x)(\mathbf{mary})]]]$
 b. Mary talked to any student who was angry.
 $\forall s[\forall x[(\mathbf{student}_s(x) \wedge C(s) \wedge \exists s'' > s[\text{PAST}(s'') \wedge \mathbf{angry}_{s''}(x)]) \rightarrow \exists s' > s[\mathbf{talk}_{s'}(x)(\mathbf{mary})]]]$

- (5.37) a. Mary talked to every possible bachelor.
 b. Mary talked to every bachelor possible.

Here is the crucial observation: implicit relative adjectives such as *possible* and *necessary* can apparently subtrig *any* when occurring in a postnominal position, but *not* when occurring in a prenominal position.¹² The evidence for this generalization comes from contrasts such as the following:

- (5.38) *Context: Mary is desperate to get married.*
 a. (?)At the party, Mary talked to any possible bachelor.
 b. At the party, Mary talked to any bachelor possible.
- (5.39) *Context: Mary gets better tips at work if she talks to certain bachelors.*
 a. *To make rent this month, Mary talked to any necessary bachelor.
 b. To make rent this month, Mary talked to any bachelor necessary.

Importantly, the unavailability of subtrigging with prenominal *possible* or *necessary* cannot be a confound of the direct/implicit relative ambiguity from Chapter 4. The postnominal position of *necessary* in (5.39b) makes its direct reading unavailable. And (5.39a) is degraded regardless of how *necessary* is interpreted.

What can the contrasts in 5.38 and 5.39 tell us about the phenomenon of subtrigging? One natural conclusion to draw is that relative clause structure just can't be the licenser of *any* in subtrigging. Here is why. Assume for the sake of argument that a situation variable is present in all relative clauses, and that following Dayal (1998), it is the presence of this variable—more precisely, the binding of this variable by GEN—that licenses *any*. If this is correct, then the contrast between pre- and postnominal implicit relative adjectives in ability to subtrig *any* constitutes strong evidence against the reduced relative analysis of prenominal indirect *possible/necessary*. However: we have seen in Chapter 5 extensive arguments that all implicit relative readings involve a covert relative clause. Therefore, by contraposition, since prenominal implicit relative adjectives *do* contain a covert relative clause, one of two things must be true: the first possibility is that the kind of relative clause contained in *possible* or *necessary* lacks a free situation variable (hence the deviance of (5.39a)). However, if this were true, then it would follow that (5.39b) should be ungrammatical for the same reason as is (5.39a), given that the syntactic position of an implicit relative adjective does not affect its internal structure (namely, absence of a situation variable). This prediction, of

¹² Perhaps marginally for *possible*, more research and judgments needed.

course, contradicts the observed fact that (5.39b) contrasts sharply with (5.39a) in grammaticality. Another possibility is simply that relative clause structure and/or the presence of a situation variable is completely orthogonal to the phenomenon of licensing by modification.

If the above argument is correct and relative clause structure is not the licenser of subtriggering, then what is? If not (reduced) relative clause structure, what is the defining property of postnominal modifiers in English? According to the conclusions of previous chapters, outer syntactic position cannot be the right property either. I leave these questions for future research.

5.3 Structural implicatures and non-restrictive modifiers

This section explores theoretical issues relating to the following observation: non-restrictive adjectives fail to trigger a specific kind of manner implicature in downward-entailing contexts, in contrast to restrictive modifiers. For example the restrictive modifier *harmful* in (5.40a) generates the implicature in (5.40b) because the alternative to (5.40a) *Every substance will...* is both briefer and logically stronger than (5.40a). The negation of this alternative, namely (5.40b), is therefore inferred as an implicature by a Gricean principle of manner/brevity (ignoring Sauerland's (2004) primary/secondary implicature distinction). But no corresponding implicature is likely to be inferred from an utterance of (5.41a), which contains a *non*-restrictive occurrence of *harmful*.

- (5.40) a. Every harmful chemical will be eliminated by this product.
b. \rightsquigarrow Not every chemical will be eliminated by this product.
- (5.41) a. Every harmful toxin will be eliminated by this product.
b. $\not\rightsquigarrow$ Not every toxin will be eliminated by this product.

Before attempting any explanations of this contrast, we first explicate in some detail why it is interesting from a theoretical point of view in the first place.

5.3.1 Scalar implicatures, manner implicatures, and structural alternatives

An utterance of sentence (5.42a) often invites the addressee to infer something like (5.42b).

- (5.42) a. John ate some of the beans.
b. John did not eat all of the beans.

This inference is called a scalar implicature, a specific kind of conversational inference that can arise from sentences containing scalar terms such as numerals (e.g. *three*), quantifiers (e.g. *some*), gradable adjectives (e.g. *warm*), and certain connectives (e.g. *or*). Crudely speaking, $\neg\psi$ is a scalar implicature of ϕ if ψ is a (relevant) stronger utterance than ϕ and differs only in what scalar terms it contains. What distinguishes scalar implicatures from garden variety conversational implicatures is that they require consideration of determinate, fixed sets of alternative utterances (“alternatives”), generated on the basis of entailment-ordered lexical scales—sometimes called “Horn scales”—such as $\langle \textit{some}, \textit{all} \rangle$ and $\langle \textit{warm}, \textit{hot} \rangle$, or $\langle \textit{one}, \textit{two}, \textit{three} \dots \rangle$. So part of the reason that (5.42b) is an implicature of (5.42a) is that $\langle \textit{some}, \textit{all} \rangle$ is a Horn scale, and therefore uttering (5.42a) forces the addressee (if relevant) to consider the scalar alternative *John ate all of the beans*. A language user then uses a set of Gricean principles (say) to reason about the logical relationship between the utterance and its alternatives, ultimately reaching conclusions about speaker beliefs based upon what a speaker chose *not* to utter.

Importantly for our present discussion, the monotonicity of a scalar term’s environment affects a sentence’s implicatures. Since negation creates a downward-entailing context, (5.43a) is stronger than (5.43b). Therefore if (5.43b) is uttered, this will usually give rise to the primary implicature that (5.43a) is not assertable—because if it were then it would have been more informative—and possibly also the stronger (“secondary”) implicature that (5.43a) is false (which is equivalent to (5.42a)). Put another way: negation “reverses” the ordering of logical strength between *some* and *all*-sentences, as a result of which implicature patterns are also reversed.

- (5.43) a. John didn’t eat some of the beans.
 b. John didn’t eat all of the beans.

A similar phenomenon exists in the domain of *manner* implicatures—inferences made upon the basis of *how* and *how much* information is linguistically encoded. For example, the adjective *harmful* in (5.44a) is most naturally interpreted restrictively. Therefore, the alternative utterance (5.44b) is a logically stronger assertion than (5.44a) (the former asymmetrically entails the latter given that harmful chemicals form a proper subset of chemicals). A reasoner might then conclude that if the speaker had believed the more informative (5.44b), she would have uttered that instead. By contraposition, the reasoner concludes that the briefer, stronger form (5.44b) is unassertable or false. In other words, an utterance of (5.44a) often implicates (5.44c).

- (5.44) a. Every harmful chemical will be eliminated by this product.
 b. Every chemical will be eliminated by this product.
 c. Not every chemical will be eliminated by this product.

There are (at least) two ways of understanding this phenomenon theoretically. First, one might attribute the pattern in 5.44 to a Gricean principle of manner, according to which redundant or extraneous words should not be used (Simons 2013; Schlenker 2004,2012; Schlenker 2008; a.o.). Such a principle would lead one to conclude that *harmful* in (5.44a) is not redundant (i.e. is restrictive), and as a result that (5.44b) is stronger than (5.44a). There are a variety of ways that this kind of principle could be stated; see Schlenker 2004,2008 and Katzir 2007 for some proposals.

Alternatively, manner implicatures of the variety in 5.44 could be viewed as a special case of *scalar* implicatures: just as we have entailment-ordered lexical scales like $\langle \textit{some}, \textit{all} \rangle$ and $\langle \textit{warm}, \textit{hot} \rangle$, a modified noun $\mu(v)$ can be viewed as forming a scale with the noun v alone.¹³ Crucially, the ordering $\langle [_{NP} v], [_{NP} \mu(v)] \rangle$ is logical in the exact same way that lexical scales' orderings are: *tall man* implies *man* just like *hot* implies *warm*. And since the complement of *every* is a downward monotonic environment, the unmodified form (5.44b) is logically stronger than the modified form (5.44a) in the same way that the *some*-sentence in (5.43a) is stronger than its *all*-alternative in (5.43b).

A theory of the second variety was advanced by Katzir (2007). His proposal—the theory of structural alternatives—was designed to subsume scalar cases like 5.42 while avoiding the symmetry problem (see Fox & Katzir 2011 for a clear formulation) as well as modificational cases such as 5.44. The alternatives of a sentence are determined by the presence of scalar items (as in Horn-scale based approaches; see Sauerland 2004 and references therein) but also by syntactic structure: given a sentence φ , sentence ψ is an alternative to φ , written $\psi \in Alt(\varphi)$, iff ψ can be obtained from φ by a finite series of (i) deletions of constituents, (ii) contractions (remove tree edge and identify end nodes), and (iii) substitutions of terminal elements for other elements of the same syntactic category, which are either lexical (in the case of scalar terms) or are subtrees of φ . For example, (5.42b) is an alternative of (5.42a) in virtue of (iii): just substitute the D *all* for the D *some*. And (5.44b) is an alternative to (5.44a) in virtue of deletion and contraction: just delete the item *harmful* from (5.44a), identify its mother node with its grandmother node, and then we have

¹³ The main difference between these two kinds of “scales” is that an utterance containing an unmodified noun will not evoke (infinitely many) alternatives containing modified forms, while an utterance containing a scalar term will always evoke an alternative containing a distinct scalemate. This difference will be shown to follow from the approach here, namely the theory of structural alternatives of Katzir (2007).

sentence (5.44b). Once alternatives are specified, their relative logical strengths can be reasoned about and implicatures can be computed on the basis of some algorithm (of which a great many have been proposed; see Sauerland 2012 for a recent survey).

A *better-than* relation \prec can be imposed on a fixed set of alternatives $Alt(\varphi)$ (including the trivial alternative φ itself). For $\psi \in Alt(\varphi)$, ψ is *at least as good as* φ , written $\psi \preceq \varphi$, iff ψ entails φ . If $\psi \preceq \varphi$ but $\varphi \not\preceq \psi$, then ψ is *strictly better* than φ , written as $\psi \prec \varphi$. Equivalently,

(5.45) ψ is better than φ , written $\psi \prec \varphi$, iff

- a. $\psi \in Alt(\varphi)$; (ψ is an alternative to φ)
- b. $\psi \Rightarrow \varphi$; and either (ψ entails φ)
- c. (i) $\varphi \notin Alt(\psi)$; or (φ is not an alternative to ψ)
(ii) $\varphi \not\Rightarrow \psi$. (φ does not entail ψ)

Notice that the definition of \prec does the same work as would a Gricean principle of manner (brevity), since relative structural simplicity factors into the definition of Alt and hence \prec .

One question left unaddressed by the theory of structural alternatives is exactly what notion of “entailment” is relevant for the computation of \preceq . Two candidate notions are specified in 5.46 (from Magri 2009).

(5.46) Where $[\varphi], [\psi], C \subseteq W$ and $C = \bigcap \{p \in \wp(W) \mid p \text{ is mutual belief}\}$,

- a. ψ *logically entails* φ , written $\psi \Rightarrow \varphi$, iff $[\psi] \subseteq [\varphi]$.
- b. ψ *contextually entails* φ in C , written $\psi \Rightarrow_C \varphi$, iff $([\psi] \cap C) \subseteq [\varphi]$.

Magri (2009,2011) argues that (5.46a)—*logical* entailment—is the notion of entailment relevant for computing scalar implicatures. This is called the “blindness hypothesis” because it says that scalar implicatures are computed without access (“blind”) to world knowledge. Blindness is crucial in capturing the oddness of sentences like #*Some Italians come from a warm country*: this sentence triggers the blind implicature that not all Italians come from a warm country, even though world knowledge guarantees that if some Italians come from a warm country, then all do (and conversely). In other words, the oddness of the sentence follows if (5.46a) is the relevant notion of entailment, but is unexpected on the context-relativized notion (5.46b).

Because Katzir’s (2007) theory treats both scalar and manner implicatures, there is the additional question whether something like blindness also holds for the latter. Here is an argument that strength should

indeed be computed blindly for manner implicatures. Consider the sentences in 5.47, relative to the specified context:

- (5.47) *Context:* At the party, we are surprised to see that Mary decided to show up. So I say
- a. Look, Mary came to the party too!
 - b. Look, Mary came to the party!

Syntactically, the only difference between (5.47a) and (5.47b) is the presence of *too* in (5.47a). Semantically, *given the specified contextual information*, the two sentences are equivalent (since the speaker is at the party, the presupposition of *too* is satisfied). In the given context, uttering (5.47a) would clearly however *not* give rise to the implicature that (5.47b) is not assertable (from which a contradiction would immediately follow). But if entailment is defined in terms of (5.46b), then it follows that (5.47b) \prec (5.47a) (because (5.47b) is a simpler contextually equivalent structural alternative to (5.47a)). Therefore, according to Katzir's (2007) theory and a context-sensitive notion of alternative-strength, an utterance of (5.47a) will implicate at least that (5.47b) is not assertable, and potentially also that it is false. This contradicts the strong intuition that no such implicature could in principle be generated by (5.47a).

This result can be generalized: whenever *too* introduces a presupposition that is already mutual belief, a sentence containing *too* is predicted to be strictly worse than its *too*-less counterpart. Viewed another way, if contextual entailment is used to compute manner implicatures, then any lexical item whose sole function is to introduce a presupposition (e.g. *also*, *too*, *even*) should always be infelicitous if its presupposition is known to be true by all conversational participants. This is obviously an undesirable result.

Now suppose the notion of *logical* entailment (= (5.46a)) is used in 5.45. In this case, (5.47b) is *not* predicted to be strictly better than (5.47a). The following sequence of steps shows why:

- (5.48) a. (5.47b) is a structural alternative to (5.47a) (delete *too* from (5.47a) and contract its branch);
- b. but (5.47b) does not logically entail (5.47a), since there are logically consistent scenarios (inconsistent with the context) in which Mary came to the party and no one else did.

Hence a context-insensitive definition of entailment predicts that an utterance of (5.47a) will *not* give rise to the implicature that (5.47b) is unassertable. This is the desired result, and applies to any purely presupposi-

tional lexical item.¹⁴

An interim summary: scalar and manner implicatures can be unified with the theory of structural alternatives, and the relative strengths of alternatives should be computed context-insensitively for both varieties.

5.3.2 The puzzle: Manner implicatures of (non-)restrictive modifiers

With a concrete theory of manner and scalar implicatures in hand, let us now revisit the critical data from 5.40-5.41 above:

- 5.40 a. Every harmful chemical will be eliminated by this product.
b. \rightsquigarrow Not every chemical will be eliminated by this product.
- 5.41 a. Every harmful toxin will be eliminated by this product.
b. $\not\rightsquigarrow$ Not every toxin will be eliminated by this product.

What does the theory of structural alternatives with blindness predict about the implicatures of sentences (5.40a) versus 5.41? Take case (5.40a) first. Each of the following statements verifies one of the three conditions in 5.45, thus predicting the observed pattern in 5.40.

- (5.49) a. $Alt((5.40a)) = \{\dots, \text{Every chemical will be eliminated by this product (}=\text{df } \varphi') , \dots\}$
b. Every chemical is eliminated \Rightarrow Every harmful chemical is eliminated
c. (5.40a) can't be obtained from φ' via substitution, deletion, and contraction, so (5.40a) $\notin Alt(\varphi')$.

Jointly, these statements establish that φ' is better than (5.40a). Therefore, (5.40a) is predicted to have the implicature that φ' is unassertable and under the right conditions—which for simplicity we assume obtain—that φ' is false. And the denial of φ' is equivalent to the observed implicature in (5.40b), exactly as desired.

¹⁴ A slightly stronger point can be made: it is not just logical entailment that is required here, but rather logical entailment *relative to local contexts* in the sense of Schlenker 2009 (at each clause should be sufficient). If entailment were computed globally, then (ia) should be dispreferred to (ib), since (ia) has no global presupposition.

- (i) a. John came to the party, and Mary came to the party too.
b. John came to the party, and Mary came to the party.

But if excludability is determined via (5.46a) relative to local contexts, no problem arises. This is because the second clause of (ia) *does* have a presupposition—(ib) does not entail (ia) if entailment is defined relative to local contexts. Even though (ib) is a strictly simpler structural alternative to (ia), the second clause of (ib) does not logically entail the second clause of (ia) viewed in isolation.

Consider now 5.41, which differs minimally from 5.40 in using the noun *toxins*; and assume that toxins' being harmful is not a logical truth, i.e. that there are logically consistent states in W (but presumably not in any reasonable C) such that $\llbracket[\text{NP harmful toxins}]\rrbracket \neq \llbracket[\text{NP toxins}]\rrbracket$. What does the theory of structural alternatives plus blindness predict in this case?

- (5.50) a. $Alt((5.41a)) = \{\dots, \text{Every toxin will be eliminated by this product } (=_{df} \varphi'), \dots\}$
 b. Every toxin is eliminated \Rightarrow Every harmful toxin is eliminated
 c. (5.40a) can't be obtained from φ' via substitution, deletion, and contraction, so (5.40a) $\notin Alt(\varphi')$.

The conditions of 5.45 are satisfied by the set of statements in 5.50, exactly as they are in the preceding case of 5.49. In other words, nothing in the theory as defined above blocks the implicature of (5.41b) from (5.41a). Furthermore, using a notion of contextual entailment for the computation of strength would then make incorrect predictions for the *too*-case presented in 5.47 as well as for Magri's (2009) data cited above. So under a naive theory of non-restrictive modification, this appears to be a genuine puzzle. In what remains, I will sketch three possible solutions to this problem. The first solution is based crucially upon a specific proposal of this dissertation, namely that non-restrictive modification structures generate presuppositions via **TMAP**.

5.3.3 Three analysis sketches

There are a number of ways that the undesired inference pattern of (5.41a) \rightsquigarrow (5.41b) could conceivably be blocked. In this discussion, three possibilities are considered.

Here is one theory: nouns can denote kinds, and adjectives can denote predicates of kinds (see Chapter 3:§5 for motivation and for details of how 5.51 is derived). In non-restrictive modification, **TMAP** combines an adjective and a noun to form a constituent that presupposes the adjective property holds of the noun's referent, and which has the same referent as the unmodified noun. In other words, the proposals of Chapters 3 and 4 generate the following semantics as one interpretation of *every harmful toxin*:

$$(5.51) \quad \llbracket[\text{DP every harmful toxin}]\rrbracket = \lambda P[\text{harmful}(\lambda w[\text{toxins}_w(x)])]. \forall x[\text{toxins}_@(x) \rightarrow P(x)]$$

The crucial bit of 5.51 is that via presupposition, it guarantees that the alternative *Every toxin will...* does not logically entail *Every harmful toxin will...*, because the latter has a presupposition that the former lacks. Hence, as with the *too*-case, we get the desired blocking.

A second approach, advanced in Leffel 2011, crucially relies on a presuppositional treatment of non-restrictive adjectives as well. The primary divergence of this theory is that it assumes the two-domains based syntax of Cinque (2010) for restrictive versus non-restrictive adjectival modification. In this theory, the presupposition in 5.51 is generated syntactically, via the (postulated) functional head that introduces non-restrictive direct modifiers. Specifically, this functional item is defined as in 5.52. Composing 5.52 with simple property-type entries for the noun and the adjective results in the term in (5.53b) as the denotation for the structure shown in (5.53a).

$$(5.52) \quad \text{GEN}_{\text{NR}} : \lambda P \lambda Q \lambda y : \Gamma x [P(x)] [Q(x)]. P(y)$$

$$(5.53) \quad \text{a.} \quad [_{\text{XP}} [_{\text{AP}} \text{harmful}] [_{\text{X}'} [_{\text{X}} \text{GEN}_{\text{NR}}] [_{\text{NP}} \text{toxin}]]]$$

$$\text{b.} \quad \lambda y: \Gamma x [\text{toxin}(x)] [\text{harmful}(x)]. \text{toxin}(y)$$

“the set of y such that y is a toxin (defined iff toxins are generally harmful)”

Because the mechanism by which the implicature is blocked is the same as in the **TMAP** theory, I will not comment further upon the difference between these approaches.

The third kind of analysis does not depend upon a grammatical treatment of non-restrictive adjectival modification. In other words, this analysis assumes that in non-definite nominals at least, non-restrictive readings of attributive adjectives arise via reasoning about belief states of the speaker (and potentially hearer as well), and not via genuine presupposition. The advantages of such an approach, were it sufficiently formalized, would include its parsimony and elegance, though further assumptions would need to be made if this theory is to capture the *too*-data in 5.47 above.

Essentially, a pragmatic analysis of the contrast between 5.40 and 5.41 would run as follows. First, we assign labels to the relevant sentences:

$$(5.54) \quad \varphi \quad \text{Every harmful chemical will be eliminated by this product.}$$

$$\varphi' \quad \text{Every chemical will be eliminated by this product.}$$

$$\neg\varphi' \quad \text{Not every chemical will be eliminated by this product.}$$

$$(5.55) \quad \psi \quad \text{Every harmful toxin will be eliminated by this product.}$$

$$\psi' \quad \text{Every toxin will be eliminated by this product.}$$

$$\neg\psi' \quad \text{Not every toxin will be eliminated by this product.}$$

The implicature of $\neg\phi'$ from ϕ arises just as it does on the other theories: (i) ϕ' is a briefer alternative of ϕ that entails ϕ ; and (ii) ϕ does not entail ϕ' . Therefore, ϕ has the implicature that ϕ' is not assertable, and under the right conditions this is strengthened to the secondary implicature $\neg\phi'$ (i.e. that ϕ' is false).

The divergence of the pragmatic theory comes in the analysis of 5.55. Here are the steps: ψ is uttered, instead of the briefer ψ' . The hearer notes that the briefer alternative ψ' entails ψ . The hearer then asks whether the speaker *could have* meant to implicate $\neg\psi'$, concluding that the speaker could not have meant this, since mutual belief guarantees that $\neg\psi'$ contradicts ψ (specifically, mutual belief that all toxins are harmful).

Alternatively, the hearer may reason that since ψ and ψ' are contextually equivalent but the speaker used the more complex ψ anyway, *harmful* must contribute some pragmatic effect (e.g. establishing a coherence relation between *toxins are harmful* and *this product gets rid of toxins*). This inference is licensed because without a non-trivial pragmatic effect, the use of *harmful* in ψ would violate the maxim of manner (brevity). Since the utility of *harmful* is pragmatic, and since the universal harmfulness of toxins is mutually known, no implicature to the effect of $\neg\psi'$ is computed.

5.3.4 A final puzzle: Implicatures of pre- versus postnominal modifiers

I close this chapter by returning to a critical contrast discussed in Chapter 4: in English, postnominal modifiers—including integrated relative clauses, locative and other PPs, and attributive adjectives when syntactically available—systematically lack non-restrictive interpretations. What is surprising about this fact, of course, is that prenominal attributive adjectives easily admit of such interpretations (given the right conditions of lexical choice, context, etc.). This generalization can be illustrated in a number of more-or-less equivalent ways. For example, one can intuit that of these three sentences, only (5.56a) could be felicitously stated in a scenario in which all words are assumed to be unsuitable:

- (5.56) a. Every unsuitable word was deleted.
b. Every word unsuitable was deleted.
c. Every word that was unsuitable was deleted.

Alternatively, one can examine sentences irrespective of context and consider what implications are likely to arise. For example both (5.57b) and (5.57c) seem to suggest that there is more than one president under consideration, while (5.57a) does not necessarily imply a multitude of presidents.

- (5.57) a. The Texan president will return to Crawford.
 b. The president from Texas will return to Crawford.
 c. The president who's from Texas will return to Crawford.

Finally, one can construct minimal pairs such that one sentence feels pragmatically or semantically deviant/infelicitous, while the other does not. In evaluating 5.58, world knowledge tells us that all flowers are colorful, and therefore that any modifier that denotes the property of being colorful will be non-restrictive with respect to the noun *flowers*. Thus, the infelicity of (5.58b) when compared to (5.58a) indicates that an integrated relative clause cannot be felicitously used non-restrictively, while an attributive adjective with roughly the same descriptive content can.

- (5.58) a. In Anna's garden there are colorful flowers.
 b. #In Anna's garden there are flowers that are colorful.

Regardless of how the contrast between pre- and postnominal modifiers is established, the generalization that subsumes the cases in 5.56-5.58 appears to be something of the following shape (we require that nouns have non-empty extensions so as to avoid cases involving talk about fictional entities like unicorns):

- (5.59) Where v is a noun with $[[[N v]]^C] \neq \emptyset$ and μ is a modifier, an utterance containing the phrase $[_{NP} [_N v] \mu]$ is felicitous only if $[[[_{NP} [_N v] \mu]]^C] \subset [[[_N v]]^C]$

The restriction in 5.59 should ideally follow from a principle of manner constraining redundant material, combined with some kind of explanation of the relevance of pre- versus postnominal position. Schlenker (2004) proposes the following principle of manner, which applies specifically to definite descriptions.

- (5.60) *Minimize Restrictors!* (Schlenker 2004:7) A noun phrase $[_{DP} \text{the } \mu [_N v]]$ (where the order of μ vs. v is irrelevant) is deviant if μ is redundant, i.e. if both of the following hold:
- a. *Referential irrelevance:* $[_{DP} \text{the } [_N v]]$ is grammatical and has the same denotation as $[_{DP} \text{the } \mu [_N v]]$; and
- b. *Pragmatic irrelevance:* μ does not serve another purpose.

While this principle does suitably constrain a specific kind of redundant material, it does not reflect the kind of sensitivity to linear order that appears to exist in the distribution of non-restrictive modifiers. So what is interesting about the data in 5.56-5.58 is that they show that a highly intuitively plausible commonsense

principle like 5.60 cannot be fully correct. In particular, the linear order of μ and ν is apparently *not* irrelevant. Even more surprising is *how relevant* linear order is: the pragmatic relevance of *Texan* in (5.57a) is apparently enough to license this referentially irrelevant/redundant modifier. By contrast, the postnominal modifier *from Texas* in (5.57b) has exactly the same kind of pragmatic relevance, but nevertheless causes infelicity (under the assumption that *from Texas* is referentially irrelevant).

A fourth way that the pre- versus postnominal contrast in (non-)restrictiveness can be illustrated is in terms of manner implicatures of the variety discussed in this section. The inference patterns in 5.61 contrast, for example, for precisely the same reason that (5.57c) but not (5.57a) implies the existence of more than one president—which in turn is the same reason that (5.58b) but not (5.58a) is infelicitous, and so on. The point can be strengthened still: (5.62a) has the expected implicature, but since non-restrictive interpretations are unavailable for relative clauses, (5.62b) necessarily has a corresponding implicature. Since this implicature relies upon an intermediate inference that contradicts world knowledge—that there exist non-harmful toxins—the sentence #*Every toxin that is harmful is eliminated* is judged as infelicitous. A theory of modification should explain how and why this intermediate inference arises.

- (5.61) a. Every harmful chemical is eliminated. \rightsquigarrow Not every chemical is eliminated.
 b. Every harmful toxin is eliminated. $\not\rightsquigarrow$ Every toxin is eliminated.
- (5.62) a. Every chemical that is harmful is eliminated. \rightsquigarrow Not every chemical is eliminated.
 b. #Every toxin that is harmful is eliminated. \rightsquigarrow Not every toxin is eliminated.

With this datum in mind, the problem can be stated thus: in English, a postnominal modifier is always required to be referentially relevant, while a prenominal modifier need not be. This way of thinking about the paradigm allows us to reduce the original problem to a simpler one. To explain the contrasts in 5.56-5.62, we need only answer the question of *why* generalization 5.63 holds of English; i.e. why prenominal modifiers are special in admitting non-restricting interpretations.

- (5.63) If $[[[_{NP} \nu]]^C \neq \emptyset$, then
- a. an utterance containing $[_{NP} [_N \nu] \mu]$ in C implies $[[[_{NP} [_N \nu] \mu]]^C \subset [[[_N \nu]]^C]$; but
- b. an utterance containing $[_{NP} \mu [_N \nu]]$ in C need not have this implication.

Why should linear order be so important in determining whether non-restrictive modification is possible/felicitous?

5.3.5 Predictions of the three analyses for the linear order puzzle

The three analyses sketched in §5.3.3 make divergent predictions about the pre- versus postnominal asymmetry. The second theory—the syntactically generated presupposition approach (Leffel 2011)—actually makes some sense of the pattern. At least in the case of non-definite nominals, non-restrictive adjectives must appear in a direct modification position, and must furthermore be introduced by the non-restrictive functional head (5.53a). Because the position of this head is fixed—and because it is immediately prenominal in English—the unavailability of non-restrictive interpretations postnominally is an immediate consequence. So despite the somewhat stipulative nature of this analysis (and despite its reliance upon a syntactic theory that I have argued here is incorrect), it does provide a crisp, straightforward solution to the problematic data in 5.56-5.62.

What about the first theory, according to which non-restrictive adjectives compose via **TMAP** while restrictive modifiers do not? This theory also captures the contrast under discussion under the assumption that **TMAP** can apply only to prenominal modification structures for syntactic reasons (setting aside appositives). In other words, the arguments in Chapter 3:§3,5 are sufficient to guarantee that only head-adjunction structures can involve **TMAP**; and since head-adjunction is only possible for prenominal modifiers in English (Chapter 3:§5), only prenominal modifiers should in principle be capable of receiving non-restrictive interpretations. And this result alone is enough to derive the pattern in question.

Because the third, pragmatic theory sketched above is not very precise, it does not make concrete predictions about the puzzle at hand. Here I supplement this approach with some auxiliary assumptions in order to give it some bite. The resulting theory explains the asymmetry between pre- and postnominal modifiers in terms of an interaction between the semantics of focus, and default syntax-prosody correspondences in English. Before presenting the details, I motivate the focus-based approach with some empirical data.

Focusing a modifier evokes alternatives of a certain kind: (5.64a) evokes the alternative in (5.64b), the idea being that there are both harmful and non-harmful chemicals, and the speaker wants to contrast the harmful ones with something else.

- (5.64) a. Cigarettes contain [HARMful]_F chemicals.
b. Cigarettes contain harmless chemicals.

The classic example in (5.65a) becomes unambiguous if *unsuitable* is focused, precisely because it evokes the alternative in (5.65b), suggesting that there do exist suitable words.

- (5.65) a. Every [unsuitable]_F word was deleted.
b. Every suitable word was deleted.

Now consider an adjective-noun pair in which things in the extension of the noun always have the adjective property (these modification structures can be used non-restrictively but not restrictively). If the adjective is focused in this configuration, the result is infelicitous: 5.66 suggests that there are non-harmful carcinogens, which contradicts our world knowledge.

- (5.66) #Cigarettes contain [HARMful]_F carcinogens.

The generalization suggested by 5.66, combined with the assumption that postnominal modifiers are in some sense “inherently focused” or “inherently contrastive,” jointly provide an explanation for why postnominal modifiers cannot be non-restrictive: since they are (noun) phrase-final, they are assigned main stress in the phrase. This prosodic status is semantically interpreted as focus, and alternatives are therefore evoked.

- (5.67) [DP Every word [unsuitable]_F] will be deleted. (only restrictive)

Further evidence for the idea that postnominal modifiers are in some sense “inherently focused” comes from the syntactic distribution of postnominal adjectives in noun phrases containing the focus particle *only*: adjectives that are generally prohibited from postnominal position can appear there if *only* is structurally above them and inside the noun phrase. This is significant because *only* typically requires an overtly focused element in its scope: a contrast between *[DP the [NP N AP]] and [DP the only [NP N AP]] could indicate that normally ungrammatical APs become acceptable if the focus alternatives they generate are actually operated on (used) by *only* or another focus-sensitive particle.

- (5.68) a. *the person victorious b. ?the only person victorious
c. *the dog afraid d. ?the only dog afraid

Albeit not entirely clear what 5.68 shows, it is at least suggestive that there is some kind of focus-related phenomenon going on here.

With these observations in mind, we now formalize the focus-based theory. First, assume that fo-

cus evokes alternatives in the sense of Rooth (1985,1992). Formally, this means that there is a distinction between the ordinary and focus-semantic values of each expression. If a sentence α contains an F-marked (e.g. prosodically) subconstituent β , then “the focus semantic value for $[\alpha]$ is the set of propositions obtainable from the ordinary semantic value by making a substitution in the position [of β]” (Rooth 1992:76). For example, the focus value of *John saw [Mary]_F*, written $\llbracket \text{John saw [Mary] }_F \rrbracket^f$, is equal to $\{\mathbf{saw}(x)(\mathbf{john}) \mid x \in D_e\}$ —the set of propositions of the form “John saw x .” Similarly the focus value of *[John]_F saw Mary* is $\{\mathbf{saw}(\mathbf{mary})(x) \mid x \in D_e\}$, the set of propositions of the form “ x saw Mary.”

Following Rooth, we assume that the alternatives evoked by focus must have certain properties: in the case of *John saw [Mary]_F*, we require that there is a contextually supplied set C of propositions satisfying three constraints: (i) that $C \subseteq \{\mathbf{saw}(x)(\mathbf{john}) \mid x \in D_e\}$; (ii) that the ordinary value of the sentence is in C ($\mathbf{saw}(\mathbf{mary})(\mathbf{john}) \in C$); and finally, that there is a proposition distinct from the assertion in C ($\exists \phi \neq \mathbf{saw}(\mathbf{mary})(\mathbf{john}) [\phi \in C]$). In plain English, where Rooth’s (1992) squiggle operator \sim indicates the position at which focus is interpreted, we assume Rooth’s *Focus Interpretation Principle*:

(5.69) *Focus interpretation principle, set case (Rooth 1992:93):* $\phi \sim C$ presupposes that C is a subset of the focus semantic value for ϕ and contains both the ordinary semantic value of ϕ and an element distinct from the ordinary semantic value of ϕ .

Under the assumption that postnominal adjectives are reduced relatives, then they would have the structure of full clauses, i.e. CPs. I adopt this assumption here for the sake of argument, showing that it affords a natural analysis of the puzzle using only existing concepts and constraints. Second, assume that CP (the clause) is a prosodic domain in which at least one constituent must be F-marked. Third, assume an adapted version of Chomsky & Halle’s (1968) nuclear stress rule: default prosodic prominence falls on the rightmost constituent in a prosodic domain. Finally, assume that the phonological feature of being the most prominent in a domain implies the semantic feature of bearing an F-marker.

With these assumptions in hand, we are in a position to sketch a focus-based analysis of the asymmetry between pre- and postnominal modifiers with respect to non-restrictive readings. For concreteness, the goal is to derive the following generalization (where for simplicity modifiers and nouns are assumed to uniformly have type $\langle e, t \rangle$):

(5.70) **Generalizations about English:**

a. *Postnominal modifiers must restrict noun denotations*

If $\llbracket N \rrbracket \neq \emptyset$, then a nominal $[_{DP} D [_{NP} N XP]]$ implies $\llbracket [_{NP} N XP] \rrbracket \subset \llbracket [_{NP} N] \rrbracket$.

b. *Prenominal modifiers need not restrict noun denotations*

There are felicitous nominals $[_{DP} D [_{NP} XP N]]$ such that

$\llbracket [_{NP} XP N] \rrbracket = \llbracket [_{NP} N] \rrbracket \neq \emptyset$.

Now consider an occurrence of the NP *chemicals that are harmful* with default stress. To the extent that prosodic prominence can be intuited, highest prominence will fall on the stressed syllable of *harmful*—the initial one. Therefore the relative clause *that are harmful* is F-marked via focus projection (see Selkirk 1996), and we have the configuration in 5.71.

(5.71) $[_{NP_2} [_{NP_1} [_{N} \text{carcinogens}] [_{CP_F} \text{that are hármful}]]] [\sim C]$

(where \acute{v} marks main stress and \grave{v} secondary stress)

By the focus interpretation principle 5.69, the constituent NP_2 must satisfy the presuppositions in 5.72. Assuming these are satisfied, the value of NP_2 will simply be that of NP_1 .

(5.72) $\llbracket NP_2 \rrbracket^o$ is defined iff for a salient set C ,

a. $C \subseteq \{\lambda x[\text{carcinogen}(x) \wedge P(x) | P \in D_{\langle e,t \rangle}]\}$;

b. $\lambda x[\text{carcinogen}(x) \wedge \text{harmful}(x)] \in C$; and

c. $\exists \varphi \neq \lambda x[\text{carcinogen}(x) \wedge \text{harmful}(x)] [\varphi \in C]$

In other words, in order for 5.71 to be felicitous, there must be a non-trivial set of alternatives to 5.71 that have the same syntactic shape as *carcinogens that are harmful*. So because the modifier is F-marked, we generate the presupposition that there are contextually salient properties of the form $\lambda x[\text{carcinogen}(x) \wedge P(x)]$. Crucially, the existence of such properties is incompatible with a non-restrictive construal of the relative clause! Presumably the property $\lambda x[\text{carcinogen}(x) \wedge \text{harmless}(x)]$ would be an element of C , and if so, then an utterance of *carcinogens that are harmful* will force us to consider this alternative. In other words, an utterance of *carcinogens that are harmful* presupposes via focus that there is a set of harmless carcinogens. This, I propose, is the explanation for generalization (5.70a). Viewed from a different angle: the oddness of e.g. *Cigarettes contain carcinogens that are harmful* makes sense because it would be quite odd for a speaker to invite an addressee to consider an alternative utterance that refers to a property that is

mutually known to have an empty extension.

Finally, we consider a prenominal variant of 5.71:

(5.73) $[_{NP_2} [_{NP_1} [_{AP} \text{harmful}] [_{NF} \text{carcinogens}]] [\sim C]]$

In this case, since the noun *carcinogens* bears default stress in virtue of its phrase-final position, the kind of alternatives to be considered are not properties of the form $\lambda x[\mathbf{carcinogen}(x) \wedge P(x)]$ but rather properties of the form $\lambda x[P(x) \wedge \mathbf{harmful}(x)]$. As a result, the contextually specified set of alternatives *C* will never contain the property $\lambda x[\mathbf{carcinogen}(x) \wedge \mathbf{harmless}(x)]$, and therefore an utterance of *harmful carcinogens* with default prosody will never require that there are non-harmful carcinogens. From this, generalization (5.70b) follows.

This approach makes many specific predictions, but I will point out only one here: if focus is shifted to the modifier, non-restrictive readings will no longer be available. In other words, this theory predicts that modified nouns like $[_{NP} \text{harmful}_F \text{ nouns}]$ are incompatible with a non-restrictive construal of *harmful*. The discussion surrounding examples 5.64-5.67 should be sufficient to confirm the accuracy of this prediction.

A final, concluding remark: of the three analytical strategies, the pragmatic account is arguably the most theoretically conservative, the most intuitive, and also the most elegant. But despite its advantages, it must be conceded that the proposal runs into serious problems when faced with the considerations of §3.1-3.2, or when one looks at languages whose prosodic properties are different from those of English. But concerning cross-linguistic variation, the pragmatic theory also makes very interesting typological predictions, whose investigation could constitute very strong evidence for or against it (depending on results of said investigation). To mention just one: we might expect the pre- versus postnominal contrast in implicatures to be reversed in languages that allow both pre- and postnominal modification, but differ from English in having default phrase-initial stress instead of phrase-final stress.

Chapter 6

Conclusions

In this dissertation, I have introduced a general syntax and semantics for nominal modification. Like the influential two-domains model, I argue for a fundamental asymmetry among attributive adjectives, which can be informally characterized as “inner” versus “outer.” However, unlike most previous approaches, I locate the interpretive properties associated with each modifier type in the semantics of nouns and noun phrases. The most important and innovative features of my theory are the following:

- Nouns lexically denote taxonomic sub-kind predicates, expressions of type $\langle\langle s, e \rangle, t \rangle$.
- Inflectional morphology on nouns is semantically checked at the lowest NP node, and has two effects:
 - it converts the noun’s denotation into an $\langle s, \langle e, t \rangle \rangle$ -type intensional predicate of individuals; and
 - it operates on the mereological structure of the noun’s extension; e.g. singular noun denotations consist solely of atomic entities, while plurals contain atoms and sums.
- Attributive adjectives, relative clauses, and other non-appositive modifiers are adjoined either to N (head-adjunction) or to NP (phrasal adjunction). As a result,
 - head-adjoined modifiers constrain the type/kind of entity that the noun applies to; while
 - phrasally adjoined modifiers constrain the set of individuals that the NP predicate is true of.
- Language-specific parameter settings determine whether
 - argument nominals can have category NP, or must have category DP;
 - head-adjunction is leftward only, rightward only, or bi-directional; and
 - phrasal adjunction is leftward only, rightward only, or bi-directional.

This theory is applied to three Bolinger contrasts; in each case the basic distribution of readings in English and in Italian as reported in the literature is derived (with minor auxiliary assumptions). Further empirical patterns are also explained, which I show cannot be accounted for on previous approaches. A fourth topic,

modified proper names, is also shown to have a natural analysis on my theory of modification.

Chapter 2§5.2: stage-level versus individual-level interpretation of attributive adjectives

Chapter 3§4: restrictive versus non-restrictive interpretations of attributive adjectives

Chapter 4§4: direct versus implicit relative interpretations of certain modal adjectives

Chapter 5§1: proper names containing modifiers, articles, and both

Aside from the empirical patterns the proposal sheds light on, it also makes precise Larson's intuition that inner (here, head-adjoined) adjectives often have generic interpretations, as well as Bolinger's intuition that inner adjectives are "reference-modifying" while outer adjectives are "referent-modifying." On my theory, reference-modification amounts to narrowing kind reference, while referent-modification amounts to narrowing individual reference.

The ultimate test for the proposals made here is cross-linguistic application: while I have sketched an account of how the mirror-image distribution of adjective interpretations in English versus Romance should be analyzed, I have not examined in any detail

- languages in which all nominal modifiers occur prenominaly;
- languages in which all nominal modifiers occur postnominally;
- languages in which there is no inflectional morphology on DP-internal elements;
- languages without determiners; or
- the host of semantically independent Bolinger contrasts, as partially catalogued in Cinque 2010:5-36.

The proposals of Chapter 2 can be viewed as a framework for the cross-linguistic analysis of modification, which builds on and (I argue) improves the two-domains theories of Larson, Cinque, and others. The merits of the theory will be determined by the range of languages—and the range of additional Bolinger contrasts—it can be successfully applied to. Some interesting problems are likely to arise: for example, in languages without the singular/plural distinction, it may be necessary to posit systematic polysemy in the lexical semantics of nouns, whereby they can denote subkind predicates or individual predicates. Alternative strategies might include positing an invisible syntactic layer in which kind-reference is converted to individual-reference.

In addition to the syntactic and semantic subject matter of this dissertation, I have also argued that discourse coherence relations play a crucial role in certain aspects of the pragmatics of modification. I find this conclusion of interest primarily because noun modification is an inherently clause-internal grammatical process, while coherence relations have for the most part been used as tools for analyzing inter-sentential

pragmatic phenomena like narrative structure, long-distance anaphora, or ellipsis resolution. The idea that coherence relations are operative at a sub-sentential level—if true—therefore has potential implications for both the organization of grammar and for the modularity (or lack thereof) of syntactic, semantic, and pragmatic processing systems. In particular, this idea predicts that pragmatic processes may be temporally interleaved with syntactic and semantic ones—a notion that conflicts with traditional idealized views of language as a system of discrete, independent, and linear combinatory subsystems (i.e. roughly the idea that sentences are constructed from lexical arrays, which are then shipped off to semantic and phonetic interpretation, and that the semantic interpretation is then fed to the general-purpose pragmatic reasoning system to compute implicatures and other conversational inferences). Instead, sub-sentential coherence effects suggest that pragmatic reasoning might interact with or even guide aspects of syntactic structure building and semantic interpretation. I have advanced analyses in terms of discourse coherence relations for the felicity conditions of non-restrictive modification (Chapter 3§2.4); and for the licensing condition that modifiers must satisfy in order to subtrig *any* in positive contexts (Chapter 5§2). Finally, the discussion in Chapter 5§3 suggests that in the domain of manner implicatures, world-knowledge does not factor into the relative strength of structural alternatives. This along with some assumptions about the prosody-syntax and prosody-semantics interfaces affords an interesting (albeit tentative) explanation for an asymmetry in the implicatures of restrictive versus non-restrictive modifiers.

My main hope is that this dissertation's framework for the study of nominal structure and modification is applied to a diverse collection of languages. Such studies may well call for adjustments to some of the basic assumptions I have introduced here, but such adjustments will hopefully improve the generality and empirical application of the basic theory I have defined in this dissertation.

Bibliography

- Abney, Steven. 1987. *The English noun phrase in its sentential aspect*: MIT dissertation.
- Alexiadou, Artemis. 2001. Adjective syntax and noun raising: Word order asymmetries in the DP as the result of adjective distribution. *Studia Linguistica* 55(3). 217–248.
- Alexiadou, Artemis, Liliane Haegeman & Melita Stavrou. 2007. *Noun phrase in the generative perspective*. Mouton de Gruyter.
- Alexiadou, Artemis & Chris Wilder. 1998. Adjectival modification and multiple determiners. In *Possessors, predicates and movement in the DP*, 303–332. John Benjamins.
- Asher, Nicholas. 1993. *Reference to abstract objects in discourse*. Springer.
- Asher, Nicholas. 2000. Truth conditional discourse semantics for parentheticals. *Journal of Semantics* 17(1). 31–50.
- Asher, Nicholas, Daniel Hardt & Joan Busquets. 2001. Discourse parallelism, ellipsis, and ambiguity. *Journal of Semantics* 18(1). 1–25.
- Asher, Nicholas & Alex Lascarides. 1998. The semantics and pragmatics of presupposition. *Journal of Semantics* 15(3). 239–300.
- Asher, Nicholas & Alex Lascarides. 2003. *Logics of conversation*. Cambridge University Press.
- Asher, Nicholas & Alex Lascarides. 2013. Strategic conversation. *Semantics and Pragmatics* 6(2). 1–62.
- Asher, Nicholas & Laure Vieu. 2005. Subordinating and coordinating discourse relations. *Lingua* 115(4). 591–610.
- Barker, Chris. 1992. Group terms in English: Representing groups as atoms. *Journal of Semantics* 9(1). 69–93.

- Barker, Chris, Raffaella Bernardi & Chung-chieh Shan. 2011. Principles of interdimensional meaning interaction. In *Proceedings of SALT 20*, 109–127.
- Barker, Chris & Geoffrey K Pullum. 1990. A theory of command relations. *Linguistics and Philosophy* 13(1). 1–34.
- Barker, Chris & Chung-chieh Shan. 2008. Donkey anaphora is in-scope binding. *Semantics and Pragmatics* 1(1). 1–46.
- Bemis, Douglas & Liina Pylkkänen. 2011. Simple composition: A magnetoencephalography investigation into the comprehension of minimal linguistic phrases. *The Journal of Neuroscience* 31(8). 2801–2814.
- Bernstein, Judy. 1993. *Topics in the syntax of nominal structure across Romance*: CUNY dissertation.
- Bierwisch, Manfred. 1989. The semantics of gradation. In *Dimensional adjectives: Grammatical structure and conceptual interpretation*, 71–261. Springer-Verlag.
- Bolinger, Dwight. 1967. Adjectives in English: Attribution and predication. *Lingua* 18. 1–34.
- Brugger, Gerhard. 1993. Generic interpretations and expletive determiner [*sic*]. *University of Venice Working Papers in Linguistics* 3(1). 1–30.
- Büring, Daniel. 2001. Discourse trees and dynamic updates. Manuscript, UCLA.
- Büring, Daniel. 2003. On D-trees, beans, and B-accents. *Linguistics and Philosophy* 26(5). 511–545.
- Bylinina, Lisa. 2013. Degree infinitival clauses. In *Proceedings of SALT 23*, 394–411.
- Carlson, Greg. 1977a. Amount relatives. *Language* 53(3). 520–542.
- Carlson, Greg. 1977b. *Reference to kinds in English*: University of Massachusetts, Amherst dissertation.
- Carlson, Greg. 1981. Distribution of free-choice *any*. In *Chicago Linguistic Society*, 8–23.
- Carlson, Greg. 2003. Interpretive asymmetries in major phrases. In *Asymmetry in grammar: Syntax and semantics*, 301–314. John Benjamins.
- Champollion, Lucas. 2010. *Parts of a whole: Distributivity as a bridge between aspect and measurement*: University of Pennsylvania dissertation.

- Chierchia, Gennaro. 1984. *Topics in the syntax and semantics of infinitives and gerunds*: University of Massachusetts, Amherst dissertation.
- Chierchia, Gennaro. 1995. Individual-level predicates as inherent generics. In *The generic book*, 176–224. The University of Chicago Press.
- Chierchia, Gennaro. 1998. Reference to kinds across languages. *Natural Language Semantics* 6(4). 339–405.
- Chierchia, Gennaro. 2009. *Dynamics of meaning: Anaphora, presupposition, and the theory of grammar*. University of Chicago Press.
- Chierchia, Gennaro, Danny Fox & Benjamin Spector. 2011. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In *Handbook of semantics*, De Gruyter.
- Chomsky, Noam. 1955. *Syntactic structures*. Walter de Gruyter.
- Chomsky, Noam & Morris Halle. 1968. *The sound pattern of English*. MIT Press.
- Cinque, Guglielmo. 1994. On the evidence for partial N-movement in the Romance DP. In *Paths toward Universal Grammar: Studies in honor of R. Kayne*, 85–110. Georgetown University Press.
- Cinque, Guglielmo. 1999. *Adverbs and functional heads: A cross-linguistic perspective*. Oxford University Press.
- Cinque, Guglielmo. 2003. The dual source of adjectives and XP- vs. N-raising in the Romance DP. Handout from NELS 34.
- 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* 25(1). 1–30.
- Corver, Norbert. 1997. Much-support as a last resort. *Linguistic Inquiry* 28(1). 119–164.
- Dayal, Veneeta. 1992. The singular-plural distinction in Hindi generics. In *Proceedings of SALT 2*, 39–58.
- Dayal, Veneeta. 1995. Licensing *any* in non-negative/non-modal contexts. In *Proceedings of SALT 5*, 72–93.

- Dayal, Veneeta. 1998. *Any* as inherently modal. *Linguistics and Philosophy* 21(5). 433–476.
- Dayal, Veneeta. 2004. Number marking and (in)definiteness in kind terms. *Linguistics and Philosophy* 27(4). 393–450.
- Dayal, Veneeta. 2005. Licensing by modification. *Ilha do Desterro* 47. 217–238.
- Demonte, Violeta. 2008. Meaning-form correlations and adjective position in Spanish. In *Adjectives and Adverbs: Syntax, Semantics, and Discourse*, 71–100. Oxford University Press.
- Déprez, Viviane. 2005. Morphological number, semantic number and bare nouns. *Lingua* 115(6). 857–883.
- Dowty, David. 2003. The dual analysis of adjuncts/complements in categorial grammar. In *Modifying adjuncts*, 33–66. Walter de Gruyter.
- Dowty, David R. 1979. *Word meaning and Montague Grammar: The semantics of verbs and times in generative semantics and in Montague's PTQ*. D. Reidel.
- Dubinsky, Stanley. 1998. Easy clauses to mistake as relatives: The syntax of English postnominal infinitives. In *Proceedings of WECOL 27*, 108–119.
- Engelhardt, Paul, Karl Bailey & Fernanda Ferreira. 2006. Do speakers and listeners observe the Gricean maxim of quantity? *Journal of Memory and Language* 54(4). 554–573.
- Ernst, Thomas. 2002. *The syntax of adjuncts*. Cambridge University Press.
- Ernst, Thomas. 2007. On the role of semantics in a theory of adverb syntax. *Lingua* 117(6). 1008–1033.
- Fleisher, Nicholas. 2008. A crack at a hard nut: Attributive-adjective modality and infinitival relatives. In *Proceedings of WCCFL 26*, 163–171.
- Fleisher, Nicholas. 2011. Attributive adjectives, infinitival relatives, and the semantics of inappropriateness. *Journal of Linguistics* 47(2). 341–380.
- Fox, Danny & Roni Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19(1). 87–107.
- Gallin, Daniel. 1975. *Intensional and higher-order modal logic: With applications to Montague Semantics*. Elsevier.

- del Gobbo, Francesca. 2003. *Appositives at the interface*: UC Irvine dissertation.
- Grosu, Alexander & Manfred Krifka. 2007. The gifted mathematician that you claim to be: Equational intensional ‘reconstruction’ relatives. *Linguistics and Philosophy* 30(4). 445–485.
- Grosu, Alexander & Fred Landman. 1998. Strange relatives of the third kind. *Natural Language Semantics* 6(2). 125–170.
- Hardt, Daniel & Maribel Romero. 2004. Ellipsis and the structure of discourse. *Journal of Semantics* 21(4). 375–414.
- Harris, Jesse. 2012. Domain adjectives. In *Proceedings of WCCFL 30*, 162–172.
- Heim, Irene. 1987. Where does the definiteness restriction apply? Evidence from the definiteness of variables. In *The representation of (in)definiteness*, 21–42. MIT Press.
- Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9(3). 183–221.
- Heim, Irene. 1999. *Notes on superlatives*. Manuscript, MIT.
- Heim, Irene. 2000. Degree operators and scope. In *Proceedings of SALT 10*, 40–64.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Wiley-Blackwell.
- Henkin, Leon. 1963. A theory of propositional types. *Fundamenta Mathematicae* 52. 323–334.
- Herdan, Simona. 2008. A superlative theory of amount relatives. In *Proceedings of WCCFL 26*, 234–242.
- Hobbs, Jerry. 1982. Towards an understanding of coherence in discourse. In *Strategies for natural language processing*, 223–243. L. Erlbaum Associates.
- Hobbs, Jerry. 1990. An integrated abductive framework for discourse interpretation. In *Working Notes of the Spring Symposium on Automated Abduction*, 10–12.
- Hobbs, Jerry & Andrew Kehler. 1997. A theory of parallelism and the case of VP ellipsis. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, 394–401.
- Hobbs, Jerry R. 1985. *On the coherence and structure of discourse*. CSLI.

- Huettig, Falk, Joost Rommers & Antje Meyer. 2011. Using the visual world paradigm to study language processing: A review and critical evaluation. *Acta psychologica* 137(2). 151–171.
- Jackendoff, Ray. 1977. *X-bar syntax*. MIT Press.
- Jacobson, Pauline. 1999. Towards a variable-free semantics. *Linguistics and Philosophy* 22(2). 117–185.
- Kamp, Hans. 1975. Two theories about adjectives. In *Formal semantics of natural language*, 123–155. Cambridge University Press.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In *Formal semantics: The essential readings*, 189–222. Blackwell.
- Kamp, Hans & Uwe Reyle. 1993. *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation*. Springer.
- Katzir, R. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30(6). 669–690.
- Kayne, Richard. 1994. *The antisymmetry of syntax*. Cambridge: MIT Press.
- Kehler, Andrew. 1993. The effect of establishing coherence in ellipsis and anaphora resolution. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, 62–69.
- Kehler, Andrew. 2000. Coherence and the resolution of ellipsis. *Linguistics and Philosophy* 23(6). 533–575.
- Kehler, Andrew. 2002. *Coherence, reference, and the theory of grammar*. CSLI.
- Kennedy, Christopher. 1999. *Projecting the adjective: The syntax and semantics of gradability and comparison*. Garland.
- Keshet, Ezra. 2012. *Embedded coherence relations*. Handout from GLOW 2012 presentation, Potsdam.
- Kim, Young-Kook. 1997. Agreement phrases in DP. *University College London working papers in linguistics* 9. 1–24.
- Kratzer, Angelika. 1981. The notional category of modality. In *Words, worlds, and contexts*, 38–74. Walter de Gruyter.

- Krifka, Manfred. 1989. Nominal reference, temporal constitution and quantification in event semantics. In *Semantics and contextual expression*, 75–115. Foris.
- Krifka, Manfred. 2003. Bare NPs: Kind-referring, indefinites, both, or neither? In *Proceedings of SALT 13*, 180–203.
- Krifka, Manfred, Jeffrey Pelletier, Greg Carlson, Alice ter Meulen, Godehard Link & Gennaro Chierchia. 1995. Genericity: An introduction. In *The generic book*, 1–124. The University of Chicago Press.
- Kripke, Saul. 2009. Presupposition and anaphora: Remarks on the formulation of the projection problem. *Linguistic Inquiry* 40(3). 367–386.
- Kubota, Yusuke & Wataru Uegaki. 2009. Continuation-based semantics for Conventional Implicatures: The case of Japanese benefactives. In *Proceedings of SALT 19*, 306–323.
- Larson, Richard. 1988. On the double object construction. *Linguistic Inquiry* 19. 335–391.
- Larson, Richard. 1991. The projection of DP (and DegP). Manuscript, Stony Brook University.
- Larson, Richard. 1998. Events and modification in nominals. In *Proceedings of SALT 8*, 145–168.
- Larson, Richard. 1999. *Semantics of adjectival 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. 2000a. ACD in AP? In *Proceedings of WCCFL 19*, 1–15.
- Larson, Richard. 2000b. Temporal modification in nominals. Handout from International Roundtable on the Syntax of Tense and Aspect, Paris.
- Larson, Richard & Sungeun Cho. 2003. Temporal adjectives and the structure of possessive DPs. *Natural Language Semantics* 11(3). 217–247.
- Larson, Richard & Franc Marušič. 2004. On indefinite pronoun structures with APs: Reply to Kishimoto. *Linguistic Inquiry* 35(2). 268–287.
- Larson, Richard & Naoko Takahashi. 2007. Order and interpretation in prenominal relative clauses. In *Proceedings of WAFL 2*, 101–120.

- Lascarides, Alex & Nicholas Asher. 1993. Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy* 16(5). 437–493.
- Leffel, Timothy. 2011. Nonrestrictive adjectives and the theory of scalar implicature. Manuscript, NYU.
- Leffel, Timothy, Miriam Lauter, Masha Westerlund & Liina Pykkänen. 2014. Restrictive versus non-restrictive composition: An MEG study. *Language, Cognition, and Neuroscience* (in press). doi:10.1080/23273798.2014.956765.
- LeGrand, Jean Ehrenkranz. 1975. 'Or' and 'Any': *The semantics and syntax of two logical operators*: University of Chicago dissertation.
- Leu, Thomas. 2008. *The internal syntax of determiners*: NYU dissertation.
- 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.
- Magri, Giorgio. 2009. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural Language Semantics* 17(3). 245–297.
- Magri, Giorgio. 2011. Another argument for embedded scalar implicatures based on oddness in downward entailing contexts. In *Proceedings of SALT 20*, 564–581.
- Marušič, Franc & Rok Žaucer. 2009. Dual source for indefinite pronoun-adjective combinations. In *Proceedings of nels 38*, 135–148.
- McCawley, James. 1981. The syntax and semantics of English relative clauses. *Lingua* 53. 99–149.
- McNally, Louise. 2004. Bare plurals in Spanish are interpreted as properties. *Catalan Journal of Linguistics* 3. 115–133.
- McNally, Louise. 2008. DP-internal 'only', amount relatives, and relatives out of existentials. *Linguistic Inquiry* 39(1). 161–169.

- McNally, Louise & Gemma Boleda. 2004. Relational adjectives as properties of kinds. In *Empirical issues in syntax and semantics 5*, 179–196.
- Mikkelsen, Line. 2002. Two types of definite description subjects. In *Proceedings of the 14th ESSLLI Student Session*, 141–153.
- Montague, Richard. 1974. *Formal philosophy: Selected papers of Richard Montague*. Yale University Press.
- Morzycki, Marcin. 2005. *Mediated modification: Functional structure and the interpretation of modifier position*. University of Massachusetts, Amherst dissertation.
- Morzycki, Marcin. 2008. Nonrestrictive modifiers in nonparenthetical positions. In *Adjectives and adverbs: Syntax, semantics, and discourse*, 101–122. Oxford University Press.
- Morzycki, Marcin. 2014. *Modification*. to appear with Cambridge University Press.
- Partee, Barbara. 1987. Noun phrase interpretation and type-shifting principles. In *Studies in discourse representation theory and the theory of generalized quantifiers*, 115–143. Foris.
- Partee, Barbara & Mats Rooth. 1983. Generalized conjunction and type ambiguity. In *Formal semantics: The essential readings (2002)*, 334–356. Blackwell.
- Partee, Barbara, Alice Ter Meulen & Robert Wall. 1990. *Mathematical methods in linguistics*. Springer.
- Piñón, Christopher. 2005. Comments on Morzycki & Katz. Manuscript, Université Lille 3.
- Polanyi, Livia. 1988. A formal model of the structure of discourse. *Journal of Pragmatics* 12(5). 601–638.
- Pollard, Carl. 2008. Hyperintensions. *Journal of Logic and Computation* 18(2). 257–282.
- Postal, Paul. 1969. On so-called pronouns in English. In *Modern studies in English: Readings in Transformational Grammar*, 201–224. Prentice-Hall.
- Potts, Christopher. 2002. The lexical semantics of parenthetical-as and appositive-which. *Syntax* 5(1). 55–88.
- Potts, Christopher. 2003. Expressive content as conventional implicature. In *Proceedings of NELS 33*, 303–322.
- Potts, Christopher. 2005. *The logic of conventional implicature*. Oxford University Press.

- Potts, Christopher. 2007. The expressive dimension. *Theoretical Linguistics* 33(2). 165–198.
- Rizzi, Luigi. 1997. The fine structure of the left periphery. In *Elements of grammar: Handbook in generative syntax*, 281–337. Kluwer.
- Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Ohio State Working Papers in Linguistics* 91–136.
- Rohde, Hannah, Roger Levy & Andrew Kehler. 2011. Anticipating explanations in relative clause processing. *Cognition* 118(3). 339–358.
- Romero, Maribel. 2011. Modal superlatives and 3-place vs. 2-place ‘-est’. *The Baltic International Yearbook of Cognition, Logic and Communication* 6.
- Romero, Maribel. 2013. Modal superlatives: A compositional analysis. *Natural Language Semantics* 21. 79–110.
- Rooth, Mats. 1985. *Association with focus*: University of Massachusetts, Amherst dissertation.
- Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1(1). 75–116.
- Rubin, Edward. 2002. *The structure of modifiers*. Manuscript, University of Utah.
- Sadler, Louisa & Douglas J. Arnold. 1994. Prenominal adjectives and the phrasal/lexical distinction. *Journal of Linguistics* 30(1). 187–226.
- Saeboe, Kjell Johan. 1996. Anaphoric presuppositions and zero anaphora. *Linguistics and Philosophy* 19(2). 187–209.
- Van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9(4). 333–377.
- Sauerland, Uli. 2003. A new semantics for number. In *Proceedings of SALT 13*, 258–275.
- Sauerland, Uli. 2004. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27(3). 367–391.
- Sauerland, Uli. 2012. The computation of scalar implicatures: Pragmatic, lexical or grammatical? *Language and Linguistics Compass* 6(1). 36–49.

- Sauerland, Uli, Jan Anderssen & Kazuko Yatsushiro. 2005. The plural is semantically unmarked. In *Linguistic evidence*, 413–434. Mouton de Gruyter.
- Scha, Remko & Livia Polanyi. 1988. An augmented context free grammar for discourse. In *Proceedings of the 12th Annual Meeting of the Association for Computational Linguistics*, 573–577.
- Schlenker, Philippe. 2004. Minimize Restrictors! (Notes on definite descriptions, Condition C and epithets). In *Proceedings of Sinn und Bedeutung*, 385–416.
- Schlenker, Philippe. 2007. Expressive presuppositions. *Theoretical Linguistics* 33(2). 237–245.
- Schlenker, Philippe. 2008. Be articulate: A pragmatic theory of presupposition projection. *Theoretical Linguistics* 34(3). 157–212.
- Schlenker, Philippe. 2009. Local contexts. *Semantics and Pragmatics* 2(3). 1–78.
- Schlenker, Philippe. 2012. Maximize presupposition and gricean reasoning. *Natural Language Semantics* 20(4). 391–429.
- Schlenker, Philippe. 2013. Supplements without Bidimensionalism. Manuscript, NYU/Institut Jean Nicod.
- Schwarz, Bernhard. 2005. Modal superlatives. In *Proceedings of SALT 15*, 187–204.
- Schwarzschild, Roger. 1996. *Pluralities*. Springer.
- Scontras, Gregory. 2014. Amount semantics. In *Proceedings of SALT 24*, .
- Sedivy, Julie, Michael Tanenhaus, Craig Chambers & Greg Carlson. 1999. Achieving incremental semantic interpretation through contextual representation. *Cognition* 71(2). 109–147.
- Selkirk, Elisabeth. 1996. Sentence prosody: Intonation, stress, and phrasing. In J.A. Goldsmith (ed.), *The handbook of phonological theory*, 550–569. Wiley-Blackwell.
- Seuren, Pieter. 1973. The comparative. In *Generative grammar in Europe*, 528–564. Springer.
- Shaer, Benjamin. 2000. Syntactic position and the readings of ‘manner’ adverbs. *ZAS Papers in Linguistics* 17. 265–286.
- Siegel, Muffy. 1976. *Capturing the adjective*: University of Massachusetts, Amherst dissertation.

- Simons, Mandy. 2013. On the conversational basis of some presuppositions. In *Perspectives on linguistic pragmatics*, 329–348. Springer.
- Solt, Stephanie. 2009. Attributive quantity words as nonrestrictive modifiers. In *Proceedings of NELS 39*, 1–14.
- Sproat, Richard & Chilin Shih. 1988. Prenominal adjectival ordering in English and Mandarin. In *Proceedings of NELS 18*, 465–489.
- Stalnaker, Robert. 1979. Assertion. In *Syntax and Semantics 9: Pragmatics*, 315–332. Academic Press.
- Stalnaker, Robert. 1998. On the representation of context. *Journal of Logic, Language and Information* 7(1). 3–19.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25(5). 701–721.
- 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.
- Szabolcsi, Anna. 1983. The possessor that ran away from home. *The Linguistic Review* 3(1). 89–102.
- Szabolcsi, Anna. 1986. Comparative superlatives. *MIT Working papers in Linguistics* 8. 245–265.
- Szabolcsi, Anna. 1987. Functional categories in the noun phrase. *Approaches to Hungarian 2* 167–190.
- Tanenhaus, Michael, Michael Spivey-Knowlton, Kathleen Eberhard & Julie Sedivy. 1995. Integration of visual and linguistic information in spoken language comprehension. *Science* 268(5217). 1632–1634.
- Thomason, Richmond & Jerry Hobbs. 1997. Interrelating interpretation and generation in an abductive framework. In *AAAI Fall Symposium on Communicative Action*, .
- Tonhauser, Judith & Paul Kiparsky. 2011. Semantics of inflection. *Semantics Handbook* .
- Txurruka, Isabel Gómez. 2003. The natural language conjunction ‘and’. *Linguistics and Philosophy* 26(3). 255–285.
- Umbach, Carla. 2006. Non-restrictive modification and backgrounding. In *Proceedings of the Ninth Symposium on Logic and Language*, 152–159.

- Vendler, Zeno. 1968. *Adjectives and nominalizations*. Walter De Gruyter.
- Vergnaud, Jean-Roger & Maria Luisa Zubizarreta. 1992. The definite determiner and the inalienable constructions in French and in English. *Linguistic Inquiry* 23(4). 595–652.
- Wilkinson, Karina. 1995. The semantics of the common noun ‘kind’. *The generic book* 383–397.
- Zamparelli, Roberto. 1998. A theory of kinds, partitives and OF/Z possessives. *Linguistics Today: Possessors, predicates and movement in the determiner phrase* 22. 1–28.
- Zamparelli, Roberto. 2000. *Layers in the determiner phrase*. Routledge.
- Zeevat, Henk. 1992. Presupposition and accommodation in update semantics. *Journal of Semantics* 9(4). 379–412.
- Zweig, Eytan. 2008. *Dependent plurals and plural meaning*: New York University dissertation.
- Zweig, Eytan. 2009. Number-neutral bare plurals and the multiplicity implicature. *Linguistics and Philosophy* 32(4). 353–407.