

UCLA

UCLA Electronic Theses and Dissertations

Title

The Meaning and Semantics of Singular Definite Noun Phrases

Permalink

<https://escholarship.org/uc/item/1x58d6vd>

Author

Nichols, Paul

Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

The Meaning and Semantics of Singular Definite Noun Phrases

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Philosophy

by

Paul Nichols

2014

ABSTRACT OF THE DISSERTATION

The Meaning and Semantics of Singular Definite Noun Phrases

by

Paul Nichols

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2014

Professor John Samuel Cumming, Co-Chair

Professor Joseph Almog, Co-Chair

In my dissertation I develop a uniform, purely-referential account of all singular definite noun phrases. The first chapter argues that names, pronouns and definite description are used to accomplish a single communicative function: namely, to convey the identity of the entity to which they refer. The differences between these three types of singular definite noun phrase is attributed to the different communicative contexts in which they are used. The second chapter extends the account of the first chapter to generic, predicational, and so-called narrow-scope definite descriptions. The view defended is that generic, narrow-scope and predicational uses of definite descriptions refer to kinds, which have as members the individuals referred to by specific uses. The third chapter applies the theory of definite noun phrases developed in the first two chapters to the semantics of indirect discourse. Once the semantics of definite noun phrases in simple sentences is properly understood, indirect discourse poses no new obstacles to a purely referential semantics.

The dissertation of Paul Nichols is approved.

David Kaplan

Edward Keenan

John Samuel Cumming, Committee Co-Chair

Joseph Almog, Committee Co-Chair

University of California, Los Angeles

2014

Contents

1	Reference and Communication	1
1.1	Introduction	1
1.2	The Use of Noun Phrases to Identify Individuals	5
1.3	Meaning and the Semantics of Noun Phrases	19
1.4	Incomplete Referring Expressions	24
2	Singular Reference and Generality	32
2.1	Introduction	32
2.2	Descriptions as Referring Expressions	38
2.3	No Syntax Required	54
3	Direct Reference in Indirect Discourse	79
3.1	Introduction	79
3.2	Basic Semantics of Indirect Discourse	81
3.3	Point of View in Embedded Clauses	82
3.4	Substitution Failures	88
3.5	Identificational Predications and the Relation of Identity	95
3.6	Conclusion	101
A	Distributively Determined Noun Phrases	102

Acknowledgments

First and foremost, I would like to thank the UCLA Philosophy Department for taking a chance on me, and then nurturing my philosophical development for the past seven years.

I would like to thank Joseph Almog for pushing me to develop my own ideas and for his unflagging support. Joseph has been my teacher, boss, adviser, dissertation chair, co-author, and mentor. There is never any doubt that Joseph will come through when you need him.

I would like to thank David Kaplan for teaching me how to understand my own ideas and to convey them clearly. David has been extremely generous with his time and wisdom as a teacher, a boss, and a dissertation committee member, and in the process of (re)writing my proposition.

I would like to thank Samuel Cumming for guiding me through the process of writing this dissertation, and for making UCLA an even better place to be a student. Sam's thoughtfulness, idealism, integrity, and ability to consider other people's point of view is evident in his teaching, mentoring, and advising.

I would like to thank Andrea Bianchi for his support and mentoring. An early version of the second chapter was written in Andrea's apartment in Bologna, under his guidance (and strict supervision!). Andrea has been extremely generous with his time reading drafts and discussing philosophical problems. Appendix A is a response to a problem raised by Andrea for the theory of definite descriptions of Chapter 2.

I would like to thank my wife, Kristina Gehrman, with whom I have worked closely on this dissertation. I have discussed every aspect of the material contained herein with Kristina, and she has edited draft after draft of the entire manuscript. Chapter 2 is a response to Kristina's work on the role of characterizing claims in normative judgments.

Finally, I would like to thank the Higgs boson, without which none of this would have been possible.

VITA

Education

- University of California, Los Angeles (September 2006 – December 2014)
 - Ph.D. Candidate in Philosophy, April 2010 – December 2014
 - M.A., Philosophy, 2008

- Reed College (1995 – 2001)
 - B.A., Philosophy, 2001
 - Thesis: “In Defense of the Obvious: Leibniz’s Law”

Chapter 1

Reference and Communication

1.1 Introduction

The logical languages typically used as models for natural languages have the structures and constituents that they have so that the logical relationships between sentences will be reflected in the syntax of those sentences – logical languages were designed specifically in order to perform logical operations formally, that is, syntactically. From the point of view of logic, the shortcoming of the language of propositional logic is not that there are states of affairs that it is too clumsy to represent; the problem with propositional logic is that it has no way of *formally* representing many of the logical relationships that hold between propositions.

In contrast, the fundamental reason why natural languages have the structures and constituents that they have is so that people will be able to understand what other people say. The reason why speakers of English do not rely on simple symbols (like “P” and “Q”) to represent complex states of affairs is because it would drastically limit what, how much, and to whom they could communicate. The reason why English sentences are articulated into structured complexes of meaningful parts is so that they can be understood. If you want to know why a natural language utterance contains a given expression, the fundamental

question to ask is: how does that expression contribute to the audience's ability to interpret the speaker's message?

Taking interpretability as fundamental suggests that rather than trying to discover how the conventional meanings of the *ultimate* constituents of a complex representation combine to *produce* the truth-conditional semantic content of the whole, we should consider how the *immediate* constituents *contribute to the understanding* of the semantic content of the whole. From this perspective, the semantic content of a constituent of a sentence is not necessarily the same as its meaning, nor does its having a meaning imply that it has semantic content.

Meaning, in contrast to semantic content, is a characteristic of the words and grammatical structures of a language, and is something that people learn in the process of learning to speak and understand the language. One way to think of the meaning of an expression is as the conventions that govern its use as a linguistic sign. The word "black" can be used to describe something as being white, but it is not then being used according to its meaning. Linguistic expressions have the potential to function as signs because part of the process of acculturation in a language community involves coming to have culturally shared knowledge of how the expressions of the language are used (by that community). Meaning is what allows people to use linguistic representations in communication with others and with themselves over time.

On the other hand, knowledge of the semantic content of an utterance is not the product of acculturation, but is instead the result of understanding a particular product of linguistic behavior. The semantic content of an utterance is what it represents in the particular circumstances in which it is being used as a representation. First and foremost, it is complete utterances that have semantic content; although complete utterances typically have meaningful parts, they are the minimal units of representation.¹ Except in special circumstances content words like "mountain", "odd" and "sing" are not used as representations on their

¹This is not to say that complete linguistic representations are always complete sentences by the standards of grammar school teachers. That a one-word answer to a question can be a complete utterance does not mean that the semantic content of that utterance is the conventional meaning of that word.

own – they are typically used as parts of representations. We can define the semantic content of a constituent of a complete utterance as being that part of the content of the whole that its presence is used to communicate.

My thesis is that the semantics of singular definite noun phrases is uniform and purely referential. In what follows here, I focus only on noun phrases that can be understood as representing a particular individual, that is, “specific” uses. The account developed herein is applied to non-specific uses and to indirect discourse reports in Chapter 2 and Chapter 3, respectively.

A theory of singular definite noun phrases must be able to account for the full diversity of their contributions to the semantic content of complete utterances. Nonetheless, there are two reasons that justify an initial focus on specific uses. First, the vast majority of uses of singular definite noun phrases are specific. Though philosophers and linguists have tended towards focusing on the most challenging cases, I suggest that a clear account of typical examples will provide a solid platform from which we can go on to account for deviations from the norm. Second, it is in specific uses that what singular definite noun phrases have in common is most easily seen. Singular noun phrases form a syntactic paradigm, and by looking at what they have in common as tools of communication we can best perceive the semantic characteristics of the category as a whole.² Although in the end the reason for the differences in the kinds of constructions the different types of noun phrases are apt to appear in requires an explanation, it is natural to first look for what they have in common by examining their common behavior.

When we look systematically at the factors that influence the choice of words to refer to a given object we see that names, pronouns and definite descriptions are in complementary distribution with requirements for communicating the identity of their referents. The choice amongst co-referential expressions is mainly determined by the speaker’s beliefs about how

²The same line of reasoning suggests that singular indefinite noun phrases ought to be taken into consideration as well. In a work in progress I am extending the pure-reference theory to singular indefinites and also to quantified noun phrases. See also Pepp et al. (2015) for a related effort.

she can best help her audience to identify the thing she is talking about. She will use expressions with different meanings to convey the same semantic content in different situations, depending on what it will take *in that situation* for the content to be understood. Once we properly understand the difference between (i) the contribution of an expression to the audience’s ability to understand an utterance and (ii) its contribution to the semantic content of that utterance, we are in a position to solve some of the major problems of semantics. For example, the idea that meaning determines reference (or denotation) is the source of the problem of incomplete definite descriptions. The idea that semantic content explains communication is the source of puzzles that have preoccupied direct reference theorists for the last 35 years.³ Neither of these issues arise once the different roles of the meaning and the semantic content of noun phrases are properly distinguished.

1.1.1 Specific Uses

We can categorize singular noun phrases by the contribution they make to the message conveyed by an utterance.⁴

(1.1) The whale₁ hit the ship₂.

(1.2) The whale₃ is a mammal₄.

(1.3) Napoleon₅ was the greatest French soldier₆.

(1.4) In the 1930s, the President of the United States₇ had few staff, most of them based in the U.S. Capitol₈, where the President₉ has always had an office₁₀.

(1.5) Every girl who received a prize thinks she deserves it.

³See for instance Kaplan (1990); Kripke (1979); Cumming (2013); Fine (2007).

⁴Examples (1.1) - (1.3) are from Strawson (1950). Sentence (1.4) is from Wikipedia (2014b). My “specific uses” are roughly equivalent to Strawson’s “uniquely referring uses”. Although my understanding of noun phrases is strongly influenced by Strawson’s work, I use a different terminology partly to avoid the issue of fidelity, but also because it is my view that the uses that Strawson contrasts with “uniquely referring uses” are also referential.

Of the underlined noun phrases in sentences (1.1)-(1.5), 1, 2, 5 and 8 are specific uses – each of them is used to identify an individual. Phrases 4 and 6 are predicational, they are used to describe something that has already been identified (by “the whale” and “Napoleon” respectively). Phrases 3, 7, 8 and 9 are like specific uses in that they are used to identify something, but they differ from specific uses in that the something they identify is a kind, not an individual – sentence (1.2) describes a kind of animal, not an individual whale; and sentence (1.4) does not express a statement about an individual president, but rather about what the people who are presidents have in common (the role of being president). Finally, I’m not going to say anything at all about the expressions underlined in sentence (1.5) until Appendix A.⁵

1.2 The Use of Noun Phrases to Identify Individuals

One of the basic communicative functions of definite noun phrases is to convey the identity of specific individuals. Whether you want to request a certain book, describe your opinion of it, or demand that it be removed from your sight, you are likely to use a definite noun phrase to identify the book you are requesting, describing, or making a demand about.

(1.6) Could you hand me *Moby Dick*?

(1.7) The book you are reading is one of my favorites.

(1.8) Get it out of my sight!

A fundamental consideration in choosing the words (and syntax) used to communicate any message is the audience’s ability to understand what the words are being used to convey. Given that one of the basic communicative functions of noun phrases is to identify individuals, we should expect that a fundamental consideration governing the choice of a noun phrase (for a specific use) is the utility of the expression for helping the audience identify the individual to which the expression refers.

⁵See Pepp et al. (2015) for an account of such sentences roughly consistent with the present work.

Singular definite noun phrases include pronouns, names and definite descriptions, each of which are used to identify individuals. Any individual can be identified by expressions of any of the three types of definite noun phrase. The syntactic type of the expression used to refer to something is not constrained by the characteristics of the referent. Rather, the type of expression used to refer to a given individual depends primarily on which type will allow the audience to identify that individual most readily. An example will illustrate the point.

Suppose I have just learned some gossip and now I want to tell people about it. Let's say that I have discovered that the chair of the UCLA philosophy department just bought a Ferrari. The words I would choose to convey this tidbit will depend on who I am talking to. If my audience happens to know the chair of my department, or at least know of him by name, I might say "Gavin Lawrence just bought a Ferrari." But to members of my department or other people who know Gavin well, I would more likely use only his first name – unless, that is, he were in the room, or we had just been talking about him, in which case I would probably say, "He just bought a Ferrari." On the other hand, to people who had not just been discussing Gavin Lawrence, don't know of him by name and were not in the same room with him, I would use a definite description and say "The chair of my department just bought a Ferrari" – unless the conversation was already about the UCLA philosophy department, in which case I would be less specific in my description and say "The chair just bought a Ferrari."

The point I want to emphasize about this example is that the choice of expression, whether to use "he", "Gavin", "Gavin Lawrence", "the chair", or "the chair of my department" varies according to what it would take for the audience to know who I am talking about, as opposed to varying according to what I want to tell them. In some sense, the same job is being done by these different expressions (with their different meanings). I use each of the different utterances, with their different subject expressions, to say the same

thing to different people.⁶

Examples like this show that there is a substantive distinction between two roles of noun phrases: *designating* an individual, and *identifying* the individual designated. What a specific use of a definite noun phrase refers to contributes to what is represented by the utterance of which it is a constituent. How a definite noun phrase identifies its referent contributes to the communication of the identity of its referent to an audience (and thus to the communication of the content of the utterance as a whole). The independence of these two roles that noun phrases play in communication is evident in the fact that what is designated and how it is designated are sensitive to different aspects of communicative interactions – namely, the content of the intended message and the relationship of the audience to the content of the message.

The three types of singular definite noun phrases are in complementary distribution according to the requirements of the communicative situation in which they are used. On any given occasion, only one of the three types is (i) best for identification, (ii) heard as natural, and (iii) likely to actually be used by the speaker. In order to make this claim more precise, it is necessary to first describe how noun phrases perform their identificatory function, and also to identify some relevant characteristics of communicative situations.

1.2.1 Identification of Referents

Noun phrases can be used to identify individuals because their use as referring expressions is governed by conventional rules. There are two mechanisms through which conventional rules contribute to interpretation. First, indirectly, conventions of use contribute to interpretation via the psychological mechanism of priming. Due to a general psychological mechanism, hearing (or seeing) a word (like “cat”) that is consistently used in representations of things of a particular kind (like cats) makes things of that kind come to mind more easily.⁷ The

⁶The same pattern holds when the definite noun phrase is in a clause subordinated by a verb, as in, “Paul thinks that Gavin Lawrence just bought a Ferrari”. See Chapter 3.

⁷See Raffraya et al. (2007); Tabossi and Johnson-Laird (1980).

second mechanism is more specific to language use. The fact that the speaker chose a particular expression serves as the basis for the audience to make inferences about the content of the message the speaker is trying to convey because they will expect the speaker to be using language “properly”.

In very general terms, the use of a noun phrase conveys the information that it is being used to designate something that the speaker chose that expression to stand for in light of the rules governing its use within the shared language of the speaker and the audience. Other things being equal, the use of a noun phrase conveys the information that the speaker believes (and believes that the audience believes) that its referent is something that the expression can be used conventionally to stand for.⁸ There are two kinds of conventions governing the use of noun phrases. There are conventions limiting the use of particular expressions to particular kinds of things – “the cat” should be used to refer to cats but not to dogs. And there are also conventions that limit the use of particular kinds of expressions to things that are related in the right way to the audience. Definite noun phrases should not be used unless they will allow the audience to identify their referents. For example, “Gavin Lawrence” should not be used to refer to Gavin Lawrence unless its use will allow the audience to identify him as its referent.

Corresponding to the two kinds of conventions, there are two kinds of information conveyed by noun phrases. First, there is qualitative information about the designated individual. The use of a name to refer to an individual conveys the information that the individual referred to is so-named; the use of a pronoun conveys information about gender and animacy; the use of a definite description conveys the information that its referent is an instance of one of the kinds of things it is conventionally used to represent (for example, “the chair” conveys the information that its referent is a chair – though not whether it is a

⁸Donnellan (1966) discusses exceptions to this generalization. In Donnellan’s example “The king is on his throne”, both the speaker and the audience are aware that the speaker has a special motive for violating the conventional limitations on the use of “the king”. There would, in fact, be no point in using “the king” to refer to the pretender if not for the fact that the conventional limitations on its use are such that, normally, its use indicates that the person to whom it refers is a king.

piece of furniture or the holder of a position in an institution).

There is a difference between the kind of qualitative information conveyed by definite descriptions as compared to names. Names convey information that is relative to the symbolic system being employed; whereas, in general, descriptions convey information that is independent of the symbolic system. That Barack Obama is named “Barack Obama” and that cats are called “cats” are both facts of language and both conventional. But the fact that Felix is a member of the kind of thing that are called “cats” is a language-independent fact about Felix. Thus, by using “the cat”, the speaker conveys the information that the individual she is referring to is a member of a language-independent kind. The use of a proper name on the other hand conveys only language-dependent information.⁹

Pronouns fall somewhere in between names and descriptions with respect to the kind of information they convey. To the extent that grammatical gender and number are conventional, the information conveyed by pronouns is like that conveyed by names – relative to the symbolic system. Nonetheless, knowing something’s grammatical gender depends on knowing that it is one of the kinds of things that have that gender conventionally. The (partially) conventional assignment of gender and number to kinds can only be applied to individuals via language-independent kind membership. For instance, the use of feminine pronouns in English is generally limited to female animals (including humans) and certain vehicles, especially boats. Thus one makes use of the knowledge that a potential referent is a member of one of those categories when interpreting a feminine pronoun.

The significance of the difference between language-dependent and language-independent information is particularly relevant to the analysis of predicational uses of noun phrases.¹⁰ However, while we are focused on specific uses, the important point is that *both* kinds of information are *used in the same way* in identifying the referent of a noun phrase. Both

⁹This is a simplification. In fact there are various conventions governing naming that have the effect of attaching non-conventional qualitative information to the use of a name. There are names for males, and different names for females; family names are used for people of the same family (according to further conventions); there are names that are used by certain religions and not others; etc. These conventions are not considered grammatical, so in some sense they are not part of conventional linguistic meaning.

¹⁰See Chapter 2 and Chapter 3.

the language-dependent information conveyed by names and pronouns and the language-independent information conveyed by definite descriptions allow language users to rely on a combination of their knowledge of language and world knowledge to connect the use of an expression to an individual via properties of the individual that match the properties that license the conventional use of that expression.

The second kind of information noun phrases convey concerns the relationship of their referents to the audience. More particularly, the use of a given noun phrase conveys information about the speaker's beliefs about the audience's ability to identify the referent. The use of a pronoun indicates that the speaker expects the audience to be able to identify the referent (at the moment of interpretation of the expression) on the basis of its grammatical gender and number. The use of a name indicates that the referent can be identified by the audience in virtue of being so-named. And the use of a definite description indicates that the referent can be identified on the basis of its being the kind of thing the description is used to designate. In general, the use of a noun phrase conveys the information that it conveys enough information for the audience to identify its referent (at the moment of interpretation). But the efficiency with which people are typically able to identify noun phrase referents would not be possible if their use was not further constrained.

Identification works as it does because expressions are chosen, not just according to whether they convey *sufficient* information, but more specifically, because they convey *the minimum* of qualitative information sufficient for identification – just enough and no more. There is a scale of (qualitative) informativeness with pronouns at the bottom and descriptions at the top. The use of any one of the three types of definite noun phrase indicates that neither of the other two would have been suitable. The use of a description indicates, other things being equal, that neither giving the name nor the gender of the referent would ensure identification by the intended audience. The use of a name indicates both that a description would have been unnecessary and also that the referent could not have been identified by gender. The use of a description with an adjective modifying the head noun

indicates that the noun alone would have left open more than one possible referent, and also that no additional modifiers were needed to get down to just one.

The specificity of qualitative information conveyed increases with the number of potential referents. More information is given when more is needed to distinguish the referent from other things that the audience can be expected to consider as possible referents. If the intended referent is female, and there are no other females that the audience will consider to be at least as plausible candidates for being the referent, then she will likely be referred to using a pronoun. If she will not be seen as the most likely female referent and the speaker knows that the audience knows her by name, and there isn't another person with the same name who might just as easily be chosen as the referent, then she will likely be referred to by name. If and only if neither of the first two conditions hold, she will likely be referred to using a description.

There are multiple factors that influence how likely of a referent a given entity will seem to an audience. A frequently cited factor is salience – which can be thought of as something like the inverse of the distance from the center of attentional focus. But salience alone does not predict which kind of noun phrase will be used. Suppose someone has just said, “My sister is coming to town next week”, when a large purple cow walks into the room. The attention of the conversants is drawn inexorably to the cow, but since they really don't know what to make of the situation, the speaker continues, “She will be living at my house for a month.” The use of “she” to refer to the speaker's sister was not inappropriate; although the cow was the most salient, and very publicly salient, female, she was not relevant to what the speaker was saying, and the audience could recognize that without having already resolved the reference of the pronoun. Rather than the most salient female, the audience is looking for the female mostly likely to be the person they are being told will be living at the speaker's house.¹¹

¹¹What cues the audience takes into account in cases such as this is a subject worthy of further research. In this case, it seems likely that the speaker would in some subtle way (perhaps in part with tone of voice) indicate that she is continuing with what she was saying before the cow walked into the room (Samuel Cumming, p.c.).

People are able to accomplish the astoundingly difficult job of communicating the identity of individual entities using definite noun phrases because the relationships between different expressions vastly increases the information conveyed by a speaker’s choice of a particular phrase. Much of what the use of a given phrase conveys is that it was chosen instead of the many alternative expressions the speaker could have used on that occasion. It is the *difference* between one phrase and all the other possible phrases that allows audiences to quickly identify the right entity out of the limitless possibilities of what could be being referred to.¹² The key to understanding the communicative function of noun phrases is to understand the structure of the system of referring expressions. The expression uttered tells the audience to eliminate from consideration things that do not meet a particular condition (for instance things that are not female, named “Linda”, or chairs). But, in most cases that will leave a vast number of alternatives from which to choose. That the expression uttered was chosen instead of the many available alternatives tells the audience that there should be no need to eliminate alternatives that do not meet other conditions; that is, that the actual referent meets a relevancy condition such that further qualitative conditions are unnecessary to reduce the alternatives down to one.¹³

In many cases a definite noun phrase can be interpreted without much reliance on the audience’s understanding of the rest of an utterance or the circumstances in which the utterance is made. When a definite follows closely after an indefinite with which it can be used co-referentially, it is often clear that second expression is meant to refer to the same

¹²In this regard, linguistic meaning fits with de Saussure (1983)’s picture: it is the differences between signs that convey their meanings.

¹³Another mechanism that contributes to interpretation is lexical entrainment. Participants in a conversation tend towards using the same phrase to refer to a given object (Brennan and Clark (1996)). Lexical entrainment can lead to violations of the generalization that the minimum sufficient information is used to convey identity. There is an apparent conflict between lexical entrainment and the use of more information to indicate the need to look further for the referent. The two mechanisms might be reconciled by recognizing that the use of qualitative information conveyed by conventional meaning is treated as informative by language users only insofar as they see its inclusion as an *effort* to distinguish the referent of the expression from distractors. In the same way that using a description rather than a pronoun involves greater effort at making something identifiable, using an unfamiliar description involves greater effort than using a familiar one, regardless of the number of modifiers it contains. Things close to mind are represented by phrases close to mind. Whether a phrase (like a face) is familiar or unfamiliar depends both on long-term and short-term experience.

things as the first. When someone begins a story with “Some guy on the bus said he was a millionaire” we can likely understand that “he” refers to the guy on the bus without relying on much world or cultural knowledge. Similarly, sometimes names can be interpreted just based on who is talking. When Bert says “Ernie” we are apt to know just who he is talking about. However, it is not uncommon for the interpretation of a noun phrase to require an understanding of the rest of the utterance or an understanding of the subject matter of the utterance or both.

(1.9) In order to make an official decision on this we need the chair.

(1.10) In order to reach the window latch we need the chair.

(1.11) I don’t think the man eating snake is real at all.¹⁴

(1.12) The rebels attacked the city under the command of General Smith.

In order to interpret “the chair” in sentences (1.9) and (1.10) the audience needs to understand the role that the referent of the expression is being described as having in the action represented, and also the world knowledge to know what kinds of things (furniture or professors) can, or are likely, to fill those roles. In order to interpret (1.11), the audience needs to use their understanding of the conversational context to even determine whether the “the man eating snake” refers to a man or to a snake.¹⁵

In order for the conventions of grammar to contribute to interpretation in examples like these, the conventions need to be accessible to high-level, global reasoning processes, which, in some cases, become conscious. The rules of grammar that underlie the use of definite noun phrases in communication are not part of an isolated language module in the brain, and the constituents they govern cannot be things of which people have no conscious awareness. In order for syntax to fit into an explanation of interpretation and communication, it has to be something to which people have appropriate cognitive access.

¹⁴Love (2014).

¹⁵One ought to hyphenate “man eating” if it is being used to attribute *being an eater of men*. However, most people seem to be unaware of this rule (including the author from whom this example was quoted).

1.2.2 Audience-Oriented

Speakers choose noun phrases that will allow their audience to identify the things to which they refer. It is common in practice for speakers to change their minds mid-utterance as to what it will take for their audience to identify a referent and consequently switch referring expressions, or add an additional phrase. Someone might begin with “Linda” and before going on with rest of the utterance add “my sister in law”. Or, the speaker might start with “my sister in law”, and then add “Linda” before going on. These changes are often (though not always) responses to signals of recognition, or lack thereof, by the audience.¹⁶

Another indication that noun phrase choice is specifically oriented towards the point of view of the audience is the kind of responses people have to the use of a phrase that does not conform to the conventions outlined above. Consider a scenario where, out of the blue (from the audience’s perspective), someone says, “I can’t believe she really did that.” Aside from wondering who did what, the audience is very likely to recognize, in the choice of “she” and “that”, the speaker’s failure to appreciate that the audience is not inside his head listening to his internal monologue, and so is not already focused on the person in question and her unbelievable action. The natural response is a recognition of the speaker’s failure to distinguish the audience’s point of view from his own.

Another type of case is a speaker who repeatedly uses unnecessarily specific expressions. In this case there is a recognition by the speaker of a difference between her point of view and the audience’s, but there is a failure to track the audience’s updated view as their experience or the conversational situation unfolds. Imagine a chemistry professor saying to a grad student advisee, “Please get me the tall skinny glass tube from the table.” With some luck, the student will avoid responding out loud, “Does she think I’m a freshman undergrad?”

The normal reactions to these kinds of mistakes are evidence of the fact that the conventions for choosing a noun phrase are oriented towards (conditioned on) the audience’s point

¹⁶A variety of cooperative strategies for securing reference identification are described in Clark and Wilkes-Gibbs (1986).

of view. The use of an inappropriate type of noun phrase signals a failure on the part of the speaker to recognize what the audience knows about the referent or its relationship to the current discourse, and, thus, a failure to judge correctly what it will take for the audience to identify the referent of the expression.

In light of the above observations concerning the use of noun phrases as devices of referent identification, we can form a core hypothesis about the use of singular definite noun phrases:

Core Hypothesis: The use of a singular definite noun phrase, NP, to refer to x indicates that x is the only individual that fits the relevance requirement of NP that also fits the qualitative requirement of NP.

As we will see shortly, there are factors aside from identification that influence noun phrase choice, and which will force a revision of the core hypothesis as a general rule. Nonetheless, the core hypothesis captures the fundamental requirement on the use of definite noun phrases in conventional linguistic communication: they must convey the identity of their referents to the intended audience.

1.2.3 What's Your Point?¹⁷

Definite noun phrases should always allow the audience to identify their referents, but in many cases they play additional communicative roles. In addition to communication of referent identity, noun phrases can also contribute to conveying the point or import of a message. An example will help illustrate the distinction between the two roles.

Suppose I see someone walking down the road in front of my house in rural Ohio who I am very surprised to see there. Naturally I want to share this interesting experience. When I relay the incident to my friends in Cote d'Ivoire, I say, "I saw the star of a very popular television show in front of my house this morning." When I tell my mother, I say, "I saw the host of The Tonight Show in front of my house this morning." When I tell my wife, I say, "I saw Jimmy Fallon in front of our house this morning." Aside from ensuring that

¹⁷This section is a development of ideas found in Hobbs (2010).

the people I talk to can identify who I am talking about, in each case I use a noun phrase that will allow my audience to grasp the unlikelihood of what I saw. If I identified Jimmy Fallon in some way that did not make it clear that it was a very surprising experience to see him, there would have been no reason to convey the information that I did see him. No one would be interested to hear that I saw a man with no known exceptional qualities on a street in Ohio.

It is noteworthy that the more familiar the audience is with an individual, the less need there is likely to be to make the point of an utterance about them perspicuous via noun phrase choice (or by any other means). Often there is no need to help the audience grasp the import of a message at all. For the same reasons, the speaker has less control over what significance the audience will take a message about a well-known individual to have. One can foreground a characteristic: “My brother John, the electrical engineer, can’t figure out how to use his smart phone.” But the import to the audience of John’s technological ineptitude still depends on how they think of John already. The speaker may intend the remark as evidence that smart phones are hard to use. His wife may take the message conveyed as further evidence that her brother-in-law is an idiot – all that education and he still doesn’t understand anything.

The degree to which particular features of an individual are crucial to the point of a message varies quite a bit. A famous person doing something commonplace is only interesting because the person is famous. On the other hand, an otherwise unremarkable person doing something bizarre is interesting regardless. If I see a man wearing a pink top hat, I may say, “That man is wearing a pink top hat!” There is no need for you to connect the man in any particular way to other things to grasp the point of the message.

In a variety of ways, conveying the point of a message is a less fundamental communicative role of singular definite noun phrases than identifying referents is. First, conveying the point of a message is only sometimes part of the communicative function of a singular definite noun phrase. Second, conveying the point of a message lacks determinate success

conditions. And third, the way that singular definite noun phrases help convey the point of a message is much less a product of conventional rules than is their use to identify referents. An utterance does not have a conventional point or import in the way that it has a semantic content.

If an utterance is part of a larger discourse, the point of the message it conveys is usually in large part its contribution to the larger discourse. The contribution of an utterance to a discourse can be theorized in various ways; for example, in terms of its relationship to a question under discussion, or in terms of its coherence relations.¹⁸ However, the phenomenon I am describing is broader than what either of these two models of discourse represent. First, there is no substantial sense in which an utterance has to be a part of a larger discourse or be responsive to a question under discussion in order to have a point, so the concept of the point of an utterance has wider application. The greater range of application is due to the broader inclusion of what an utterance can be responsive to. The point of an utterance is determined by how it is integrated within conversational participants' view of the world, which includes but is not limited to their view of the present discourse.¹⁹

1.2.4 Irony and Other Non-Semantic Functions

One's choice of words can communicate many things besides a conventionally expressed semantic content. Words can be used to insult, to be polite, to sound educated, to sound like other people, to be funny, ... etc. For the most part semantic theory need not account for these uses of language. I deal with them briefly here so as to more clearly distinguish the semantics of noun phrases from other communicative roles they play.

A sincere, conventionally conveyed message can be conveyed using noun phrases that perform semantically irrelevant functions. An interesting case is the sincere use of an ironic phrase. Suppose a teenager fancies herself to be more intelligent than her relatives. Suppose

¹⁸See Hobbs (1979); Roberts (2012).

¹⁹It might be fruitful to consider Donnellan (1966)'s distinction between referential and attributive uses in relation to whether the fact that the referent of a definite description fits the description is essential to the point that the speaker intends the utterance to make.

further that she is correct and everyone knows it (to keep the theoretical issues simpler). Her brothers may routinely refer to her with “the smartest person in the family”. If someone reports to mom, “The smartest person in the family got an A in algebra again”, the fact that the subject phrase is meant to be ironic does not undermine the conventional content conveyed by the utterance. Even if the smartest person in the family is not in fact the smartest person in the family, “the smartest person in the family” will work in the same way to convey messages about her – though such uses might arguably violate linguistic conventions. The same sorts of effects can occur with certain kinds of jokes, insults, etc.

1.2.5 Interaction Between NP Functions

As a rule, definite noun phrases used specifically should convey enough information so that the audience can identify their referents. Looked at in isolation from other influences on noun phrase choice, the core hypothesis (“The use of a singular definite NP to refer to x indicates that x is the only individual that fits the relevance requirement of NP that also fits the qualitative requirement of NP”), summarizes the function of the system of singular definite noun phrases. In many cases, there are no other influences on noun phrase choice, and in those cases, the core hypothesis is descriptively accurate. However, we have also seen that noun phrases can perform other communicative functions, which are associated with additional constraints on noun phrase choice. The need to make the point of a message explicit (or to be ironic, funny, polite, rude,...etc) can interfere with the default rule of using a phrase with the minimal required information for identification. We saw this in the example above of a man telling his wife, “My brother, the electrical engineer...” It was in order to get his point across that he included the otherwise unnecessary appositive “the electrical engineer”.

Whether information beyond the minimum required for identification is taken as confusing, as a sign of the speaker’s failure to track the audience’s perspective, as an indication of the point of the message, or as an ironic jab, depends on the audience’s ability to ascer-

tain the purpose of that information. And thus it is incumbent upon speakers to violate the core hypothesis only in ways that the audience will be able to recognize as purposeful violations (except where prior audience confusion is being accounted for). “My highly educated brother” indicates, other things being equal, that the speaker has more than one brother, and that he is talking about the only one of his brothers who is highly educated. In conversation with his wife, he may be confident that she will look for a motive aside from securing identification to explain his otherwise unnecessary specificity.

In summary: (i) definite noun phrases must convey the identity of their referent; (ii) the default is to identify the referent by satisfying the core hypothesis; (iii) violations of the core hypothesis are due to noun phrases filling more than one communicative function; and, (iv) violations do not impede the basic requirement of identification insofar as the audience can discern that a violation is occurring (and usually also what the explanation for the violation is). People interpret definite noun phrases using a combination of knowledge of the conventions governing the use of noun phrases, world knowledge, knowledge about what the speaker is doing (the action that includes producing the utterance), and their understanding of the rest of the utterance. Interpreting a definite noun phrase involves identifying the entity for which the choice by the speaker of that particular expression makes the most sense.²⁰

1.3 Meaning and the Semantics of Noun Phrases

In the preceding sections, we have seen that the information conveyed by a specific use of a definite noun phrase is fundamentally a device of identification, with secondary roles such as making the point of an utterance perspicuous. The choice of a definite noun phrase to refer to a given entity is narrowly constrained by the audience’s understanding of the relationship

²⁰An important corollary of the above account of definite noun phrase interpretation is that the process of linguistic interpretation is a global cognitive task: in order to *infer* the content of an utterance, its logical form has to be transparent. In other words, the audience needs to be able to discern its logical form with the same “part” of their mind as the part with which they consider whether Bill would most likely be talking about Sally Smith or Sally Brown.

of that entity (relative to other potential referents) to the ongoing discourse. The three basic types of definite noun phrase are in complementary distribution – which type will be used is predictable on the basis of the relationship of the referent to the discourse (from the point of view of the audience), the point the speaker is trying to make, and, in some cases, whatever secondary messages, like irony or insult, the speaker intends.

The complementary distribution of definite noun phrases is evidence of one of two things. Either the same semantic content cannot be communicated to audiences with different relationships to the subject matter, or the identificatory information conveyed by specific uses of definite noun phrases is not part of their semantic content – leaving only reference. The second option, that specific uses of definite noun phrases are purely referential, leaves us with a notion of semantic content according to which one can say the same thing to different people in different circumstances. More importantly, it allows for a fruitful (and tractable) theory of semantics that complements, rather than tries to incorporate, a broader theory of communication.

Dividing semantic content from its manner of transference separates in theory what is separate in function. The alternative fails to distinguish between the different contributions that noun phrases make to the content of messages. Specific uses of noun phrases are not used to communicate the fact their referents fit the conditions of their conventional use. The fact that names and indexical pronouns convey information about their referents, and do so as part of their communicative function, has been used as support for the claim that they are semantically descriptive.²¹ People identify the referents of names and indexicals in part via “descriptive” information. Nonetheless, it is a mistake to infer that such expressions have describing their referents as part of their *semantic* function. Names and pronouns are rarely used to describe things – they are not typically used to represent their referents as being so-named or as standing in a certain relationship to the utterance (author or audience).²²

²¹See for instance Fara (forth).

²²When names are used to describe, we can understand how they do so without recourse to a descriptivist semantics. For instance, consider “John is a real Rockefeller!” This can be used to make the statement that John is a genuine member of the Rockefeller family. It is very noteworthy that when names are used as

An important distinction is obscured if we are not careful to distinguish *what is being said* from *how people know what is being said* – even when people know what is being said because of the meaningful, conventional characteristics of linguistic utterances. Here lies the heart of the distinction between a Fregean descriptivist semantics and a referentialist semantics. Frege distinguished between what an expression stands for, its *Beduetung*, and the means by which language users know which thing it is, its *Sinn*. However, Frege also identifies the *Sinn* of a sentence, which comprises the *Sinne* of its constituents, as the thought it expresses, which, I take it, can be identified with the message the sentence is used to convey. This is exactly the point I am denying in this section.

Contra descriptivist views, the means of identification of referential semantic content is not (at least not typically) itself part of the message conveyed. With names and indexical pronouns, it is fairly easy to see that the information that the expressions themselves provide about the referent is not typically part of the message the expressions are used to convey. Saying “I am going to the store” is not typically a way of communicating that you are the person speaking, but simply a way of communicating that you are going to the store. The descriptive content of “I” is not part of the message. Rather, it is information that allows the audience to pick up on what the utterance is being used to convey, by allowing them to

predicates, they are typically preceded by an article (“a Rockefeller”, “three Rockefellers”) and in predicate position. If names just were predicates, it would not be clear why their predicative uses would require special indication.

There are two uses of “predicate” that should be clearly distinguished. First there is the traditional use of “predicate” to mean roughly “what is said of a subject”. Along the same lines, in traditional English grammar, a predicate is the verb phrase of a canonical clause. Thus in traditional grammar, names cannot be predicates because they are not verbs. They can be the complements of verbs, and they can form determined noun phrases (“a Rockefeller”) that can be the complements of verbs. (I follow the terminology of Huddleston and Pullum (2002) here. A determined noun phrase is a noun phrase with a determiner. “Rockefeller” is a noun phrase but not a determined noun phrase, whereas “a Rockefeller” and “the tycoon” are determined noun phrases.) On this understanding of what a predicate is, the relevant issue for the semantics of names is not how they can function as predicates, but, rather, what they contribute to the semantic content of the predicates of which they are constituents.

In predicate logic, a predicate is a function from individuals to truth values. So, $f(x)$ is a unary function that maps entities from the universe of discourse to either True or False. That natural languages such as English contain predicates in this sense is a hypothesis, albeit one that is currently widely accepted. In the semantic theory being developed herein, common nouns are not functions from individuals to truth values, so the fact that proper nouns sometimes occur where common nouns are more typical does not suggest that proper nouns are functions from individuals to truth values.

See Pepp et al. (2015) for a more in-depth consideration of this issue.

identify the referent. It is because the audience can *observe* that you are the person speaking that using “I” is an effective means of identifying yourself as the topic of the sentence.

The argument for a purely referential semantics of specific uses of definite descriptions is somewhat complicated by the fact that one of the typical uses of definite descriptions is to describe things. Nonetheless, the same considerations that support a purely referential account of names and pronouns apply likewise to definite descriptions. The point is obscured if we fail to adequately distinguish between the analysis of predicational uses as opposed to specific uses.

When a definite description is used to refer to a specific individual, its semantic role is not to describe that thing. Consider for example, “The one on the left is beautiful.” In a typical use of this sentence, the referent of the description is described as being beautiful, whereas the *previously available* information that it is on the left serves to identify the referent; the expression “the one on the left” lets the audience know which thing is being described as beautiful, precisely because the audience does *not* need to be told that its referent is the thing on the left. The sentence as a whole is not used to convey the information that the one on the left is the one on the left, nor that it is the referent of “the one on the left” – if the audience were not already in a position to know those things, then the speaker’s attempt at communication would fail.

On the other hand, when a definite description is used to describe something, it does not also do the job of referring to that thing. Consider for example, “Los Angeles is the second most populous city in the U.S.” This sentence would typically be used to describe Los Angeles as being the second most populous city in the U.S. As such, “Los Angeles” is not used to describe what it refers to, but simply to identify it as the topic of the statement. The expression “Los Angeles” lets the audience know what is being talked about, by relying on a standing ability to use the fact that Los Angeles is called “Los Angeles” to identify it as the referent. The predicate nominal “the second most populous city in the U.S.” is used to describe the thing that has already been referred to by the subject phrase. There is no

need for the predicate nominal to refer to the referent of the subject noun phrase: it has already been identified as the thing being described.

It is true that, in general, if you can use “the blank” to refer to something, you can describe that same thing correctly by saying of it either “It/he/she is a blank” or “It/he/she is the blank”. However, this observation does not imply that definite descriptions actually do describe the things they designate. Neither can we infer from this observation that a definite description that accurately describes something also refers to the thing it describes. What an expression can be used to do in one utterance does not tell us what it does do in another.²³

Being the subject of a verb, an object of a verb, the complement of a verb and the head of a determined noun phrase (e.g., “the three Rockefellers”) are semantic roles that noun phrases play only when they are in particular syntactic positions. There is no particular reason to expect that an expression in one syntactic position will make the same contribution to the semantic content of a complete utterance that it would make in another syntactic position. As the complement of a verb, a predicational noun phrase contributes first of all to the semantic content of a verb phrase. The semantic content of the verb phrase as a whole (what it contributes to the semantic content of a complete utterance) is only in part to be explained by the semantic contribution of any one of its sub-expressions. Thus what a verb phrase with a definite noun phrase complement contributes to the semantic content of a complete utterance should not be confused with what the complement of that phrase contributes.²⁴

²³Some authors seem not to appreciate this point. Delia Graff Fara uses this same observation (that a definite description can describe what it can refer to) as evidence that descriptions are never referential (e.g., Fara (2001)).

²⁴See Nichols (2014b, 2013) for a purely referential analysis of verb phrases with definite noun phrase complements.

1.4 Incomplete Referring Expressions

The pure-reference view gives a very clean solution to what has been a long-standing difficulty for semantic theories of definite descriptions. A definite description that can be used to describe more than one thing is considered incomplete. “The (current) president of the United States of America” is complete whereas more typical expressions like “the salt” and “the bathroom” are incomplete. Incomplete definite descriptions do not convey enough descriptive information to isolate a single individual from all of the things that could potentially be referred to (i.e., anything that can be thought of or even imagined). There are a great many chairs in the world, so whenever it is used, “the chair” will be an incomplete description: it does not describe what it designates in a way that distinguishes it from everything else, particularly from other chairs.

On any theory that holds that what a definite description designates is determined by the descriptive information it conveys, the question arises as to how incomplete descriptions are able to do their job: the descriptive content alone won’t do it. This is an issue equally for Russellian (quantificational) and neo-Strawsonian (presuppositional) theories. There is a standard hypothesis that some contextually-set parameter fills the gap, thus making apparently incomplete descriptions semantically complete – perhaps either the semantic content of definite descriptions is contextually enriched, or alternatively the domain in which potential referents are contained is contextually restricted. But there are real difficulties spelling out exactly how a contextual approach could actually work.²⁵

In stark contrast to the difficulties that arise for views on which the descriptive information determines reference, once definite descriptions are seen as just one part of a system of expressions used to convey the identities of their referents (see Section 1.2), the fact that the vast majority of definite descriptions are “incomplete” flows naturally from the way the system works. Definite noun phrases are not designed to *determine* the identity of their referents but to *communicate* it – and only in narrowly defined circumstances at that. The

²⁵See Stanley and Szabó (2000) for one of the most successful attempts at solving this problem.

role of the descriptive information is to distinguish the referent from amongst those things the audience is *actually likely to consider* as possible referents of a definite description *at that junction in the discourse*. It follows from this that descriptions will be incomplete in most cases.

It rarely takes uniquely identifying information to distinguish one potential referent of a definite description from others, and, most of the time, excess information will only mislead the audience. A description is generally only as specific as is required to eliminate from the audience's consideration things that the speaker thinks the audience might otherwise take to be the referent, and audiences understand this implicitly. Adding modifiers to a description ("the green chair" in place of "the chair") is an indication to the audience that the speaker believes that the shorter description would have been inadequate to convey her message. Because the appropriateness of a given referring expression depends on multiple factors, the audience may infer one of several things: (i) the speaker thinks that more than one chair is relevant enough to be believed by the audience to be the referent of "the chair"; (ii) the speaker has some reason aside from securing identification for emphasizing that the referent is green; or (iii) the speaker thinks that the intended referent is too far from the attentional focus of some members of the audience for them to identify it simply by its being a chair, and for some reason its being a green chair is more salient than its being a chair. It is factors like these; factors that will be taken into account by the audience for the purpose of interpreting an utterance, that determine how specific a definite description will be. How many things there are that actually fit a description is only relevant insofar as it influences the possibilities considered by the audience in the process of identification.

From the point of view of a purely referential theory of definite noun phrases, it would be rather a surprise for there to be a kind of definite noun phrase that is typically "complete" in contrast to definite descriptions, which are typically incomplete. And in fact almost all definite noun phrases are typically incomplete. Neither names nor pronouns convey uniquely

identifying characteristics of their referents.²⁶

That incompleteness is typical is an important fact to consider when assessing different theories of noun phrases. In the case of descriptivist theories, incompleteness tends to undermine the attempt to explain both reference determination and semantic content with the same mechanism. Historically the problem for descriptivist theories of the incompleteness of names and pronouns has been obscured insofar as it has not been considered a separate issue to the incompleteness of definite descriptions. If names and pronouns really were just definite descriptions, then there would be no independent fact of their being incomplete as well. Because it was hypothesized that there must be something that completes definite descriptions, the issue was not specifically addressed with regards to names and pronouns.²⁷

In the case of direct-reference theories, worries raised by incompleteness have been responded to with theories of reference determination and semantic content that aim either to downplay the prevalence of incompleteness, or to explain it away by positing logico-syntactic structures that distinguish between uses of the same name. A general method that has been employed is to identify some objective feature that distinguishes between the contribution to the truth-conditions of using a name for one person rather than another, and then hypothesizing that an abstract representation of that feature is part of what people understand when they interpret language.²⁸

The problem with this methodology is that it has failed to explain how the posited abstract representation actually contributes to communication. If people in fact pronounced or wrote subscripts on referring expressions to mark sameness of reference (or used some equivalent indication), then that would explain how people would be able to tell when expressions were used co-referentially – they would be able to hear and see the subscripts. Without an explanation of how the information is decoded, positing such representations

²⁶“I” comes the closest of any expression to being “complete”. However see Corazza et al. (2002) for evidence that it too falls short of completely determining its referent via conventionally conveyed meaning.

²⁷A related point that has been raised is the impossibility of replacing names with definite descriptions that do not themselves contain names. This issue was introduced to me by David Kaplan in his undergraduate course on Frege.

²⁸See Kaplan (1990); Fine (2007); Cumming (2013).

leaves us with the same question with which we began: How can language users recognize co-referential uses when they are not marked by observable syntactic features?

Generally the incompleteness of names and pronouns has been obscured by the focus on the metaphysics of the semantic properties of expressions as opposed to their use as a means of communication. The question of what makes it the case that when Ronald Reagan said, “Mr. Gorbachev, tear down this wall!”, his use of “Mr. Gorbachev” referred to Mikhail Gorbachev, is a question that can be answered in terms of the causal history of the utterance going back to Reagan’s interactions (both direct and indirect) with Gorbachev. But the question of how the causal histories of the utterances of names going back to the named individual relates to the use of names as public representations has not been adequately addressed. In his work on the metaphysics of names, Kaplan (1990) observes that somehow people are able to recognize the use of the same name (phonologically and lexicographically defined) as sometimes the name of one person and sometimes the name of another. The ability of Reagan’s audience to understand to whom he was referring depends, not on the causal history of expression uttered, but on their ability to make a *cognitive* connection between Reagan’s uttering it and its referent.

Part of what the pure-reference account of definite noun phrases explains is how people can tell which person’s name is being used on a given occasion. Names are used not only to stand for the things they are names of but crucially also to convey the identity of what they stand for. Whether or not we distinguish between same-sounding names on the basis of their bearers (or chains of transmission) or not, only what can be perceived can contribute to the communicative function of an expression. Names, like pronouns and definite descriptions are able to function communicatively only because their use conveys information about the relationship of their referents to the communicative situations in which they are used.

If instead of starting with the question of how words are related to the things they represent, we begin by looking at how speakers (and writers) use words to communicate messages, including conveying the identities of the things they talk (or write) about, meta-

physical worries about reference never arise. An expression can communicate the identity of a specific individual because people use referring expressions in a way that is tailored specifically to helping audiences identify the referent via preexisting connections they have to it, either from perception or memory of the referent, or else from perception or memory of something uniquely related to the referent – a story about it, or a television show in which it played a particular role.

The experience of an entity via the use of a definite noun phrase is either through preexisting channels to that very thing, or to something to which it has a known connection. People can use words to talk about things they have not had (and maybe cannot have) direct experience of, because they can experience things indirectly via communication with people who have had experience with those things (direct or indirect). It is because people can convey the identities of things using language that they can facilitate such indirect interactions. Words connect people to things because people connected to those things use them in a communicatively successful manner.

The concept of meaning used in contemporary semantic theories (namely, semantic content) is a binary relationship between linguistic expressions and the things they represent. Semantics abstracts away from the human involvement in linguistic communication when it describes meaning as a relationship between the words spoken and the things spoken about – the effect of the words being spoken on the audience is not contained within the relation of semantic meaning so understood.

When we then try to use the binary notion of meaning from semantic theory to explain the effects of linguistic communication – how it change people’s minds and their behavior – that is, the three-way interactions between words, people and the things the words represent, we are left with a mysterious gap. How does the connection between the words and the things make a connection between the people and the things? Knowing the meaning (in the sense of linguistic knowledge) of the words won’t do it. Alternatively, when we recognize semantic content as an *abstraction* from a three-way relation of meaning – words meaning

things to people – that is, words making people think of things – we can explain binary semantic characteristics of utterances from within communication theory.

So far I have argued that (i) the fundamental communicative role of the information conveyed by a specific use of a definite noun phrase is to identify its referent (not to contribute to the content of the message conveyed by the utterance); and, (ii) understanding (i) dissolves the problem of incomplete referring expressions. Now I want to argue in addition that syntactically complex definite noun phrases, such as definite descriptions, are not semantically complex.²⁹

Within standard contemporary semantic theories, the semantic value of a phrase is a function of the semantic values of its sub-expressions. In contrast, according to the view being developed here, the meanings of the constituents of a complex noun phrase constrain its conventional use, but do not determine what it refers to (its semantic value). “A/the chair” can only be conventionally used to refer to things that are chairs, but that leaves open many possibilities. The use of “the” adds the conventional constraint that the audience should be able identify the referent of the phrase.

If “the chair” is used conventionally, then if the speaker successfully communicates her message, there will be only one thing that met the conventional constraints imposed by “the” and “chair”. However, the sense in which these constraints determine the conventionally communicated referent of “the chair” on a given occasion, is quite different from the sense in which standard accounts hold that the semantic content of a complex singular definite noun phrase is the output of an algorithm, represented as a function from possible worlds (or situations) to individuals. To represent the meaning of a phrase like “the chair” as such a function is to obscure the role of flexible, global reasoning in the design and interpretation

²⁹It cannot be the case that *all* definite noun phrases are both semantically referential *and* also semantically atomic – definite noun phrases can have definite noun phrases as constituents (e.g., “(The man who ate (my lunch)) is playing tennis”. There seem to me to be reasons both for and against taking definite noun phrases that are constituents of definite noun phrases as semantically referential. They are certainly conventionally referential linguistic representations. But, the role they play in communicating the content of a complete utterance is to help identify the referent of a phrase rather than to contribute their own reference to the content of the message. With these considerations in mind, I don’t see a pressing need to have an official view one way or the other.

of definite descriptions. Under standard conditions, there will be *some* function from “the chair” to its referents; but, that is merely to say that the expression only refers to one thing at a time. There is, however, no such function that is part of the meaning of the expression. To know the meaning of “the chair” is not to know a function.

A function from worlds to individuals is *not* an abstract way of representing what could also be usefully described as the meaning of a definite noun phrase. The connection between the conventional meaning of a definite noun phrase and its referent is mediated by a flexible, global reasoning process. Thus any algorithmic approach to the semantic determination of reference of definite noun phrases is bound to fail.³⁰ And the use of a function to represent a flexible global reasoning process, while not inaccurate, fails to be explanatory.

Another reason that I take complex definite noun phrases to be semantically atomic is that the adjectives and nouns that are parts of a complete noun phrase do not themselves represent anything, at least not in the way that a complete phrase does. The word “chair” undoubtedly has a meaning, but in a quite different sense than a particular use of “the chair” has a referent as its meaning. There is no particular thing or things that “chair” represents when “the chair” is used to refer to something. There is nothing wrong with saying that “chair” represents chairs, but again, not in the same sense as “the chair” represented Gavin Lawrence in the hypothetical story above.

A clearer understanding of the different roles of content words and complete phrases is possible if we think of words like “chair” as being meaningful insofar as their inclusion in a phrase imposes a constraint on what the phrase can be used conventionally to represent. “Chair” means what it does because “the chair” is conventionally used to refer to chairs. Including both of these two senses of meaning within the technical concept “semantic value” can only undermine efforts to develop an explanatory theory of meaning.

A pure-reference account of singular definite noun phrases accounts for their contribution

³⁰In practice the successful use of definite noun phrases in real-time communication quite regularly relies on feedback from the audience. So, not only flexible, global reasoning, but also further communication and trial and error intervenes between conventional meanings and identification of referents. See Clark and Wilkes-Gibbs (1986).

to the semantic content of complete utterances while allowing for distinctions between the content of an utterance, the import of the message, and the means of identifying both the content and its import for the audience. It also allows for a distinction between the communicative functions of reference and description. Finally, it treats a syntactic paradigm (singular noun phrases) semantically uniformly.

Chapter 2

Singular Reference and Generality

2.1 Introduction

Most of the time, singular definite noun phrases are used in a way that can be naturally thought of as standing for particular individuals.

(2.1) I'd left the hospital, and was running late for the airport. But I wasn't running. I was sitting in traffic in Dar es Salaam – the wide, deep, unmoving kind of traffic...¹

Each of the underlined phrases represents a single entity: a woman, a hospital, an airport, the same woman again twice, a city, and a kind of traffic (the only one of the group that is not a concrete particular).

Just on the basis of these two lines, we might hypothesize that the semantic function of singular definite noun phrases is to represent single entities. However, we can also find examples that do not so clearly fit our hypothesis.

(2.2) If, at the time fixed for the beginning of the term of the President, the President elect shall have died, the Vice President elect shall become President.²

¹Reynolds (2014).

²U.S. Const. amend. XX.

What the underlined phrases in the second example represent is not completely obvious, but it is obvious that they do *not* represent a particular time and two particular people. The view defended in this chapter is that, just like the phrases in (2.1), these phrases refer to single entities. The entities they refer to are kinds.

Without committing to an explanation of the difference between the singular definite noun phrases in these two examples, we can call the underlined phrases in example (2.1) “specific uses” and those in example (2.2) “non-specific uses”.³ “Non-specific” is a catch-all for singular definite noun phrases that *cannot* be interpreted as representing particular individuals. Note that non-specific uses *can* represent single *entities*, just not those entities I am calling individuals.

(2.3) Be it known that, waiving all argument, I take the good old fashioned ground that the whale is a fish, and call upon holy Jonah to back me.

In this quote from Moby Dick, “the whale” cannot be understood as designating one individual animal, such as Moby Dick – it is a non-specific use. But we can interpret “the whale” as designating a single entity. We can gloss what Ishmael says with:

(2.4) The order *Cetacea* is properly classified as a type of fish.

Thus we can interpret “the whale” as designating the order *Cetacea* – a single entity, though not the same kind of entity as Moby Dick. The use of “the whale” to represent Moby Dick is a specific use, whereas its use to represent the kind of animal that includes all particular whales as members is non-specific. Likewise, the use of “the president” to designate Obama is a specific use, whereas its use to designate the role that Obama is the 44th person to occupy is non-specific.

Typically, the individuals referred to by specific uses of definite descriptions are not themselves also kinds. However, as I am using the term, “individual” designates a relative

³My use of “specific use” might be co-extensive with Strawson’s “uniquely referring use” (Strawson (1950)). Although my view is deeply influenced by his work, I use different terminology both in order to avoid the question of fidelity, and also because I argue that all singular definite noun phrases are referential, while Strawson’s terminology suggests otherwise.

characteristic. Thus it can happen that the same entity is an individual with respect to one referring expression, and a kind with respect to another. For example, the order *Cetacea* is an individual with respect to “the order *Cetacea*” and a kind with respect to “the whale”. Specific uses cannot be interpreted as designating an entity that has as members individuals that can be designated by the same expression. In (2.4), “the order *Cetacea*” is a specific use, though it designates the same thing as “the whale” in (2.3), because, in contrast to “the whale”, “the order *Cetacea*” cannot be used to designate things that are themselves members of the order *Cetacea*.⁴

Non-specific uses cannot be understood as designating individuals, but often they can be understood as generalizing over the individuals that could be designated by specific uses of the same expressions. Whereas sentence (2.1) tells us about a particular woman and particular buildings, (2.2) tells us about anyone who happens to become president or vice-president elect, and (2.3) tells us something about all whales.

There are various ways of trying to account for non-specific uses of singular definite noun phrases. We can divide these accounts into two groups. First, there are ontological theories that attribute the difference between specific and non-specific uses to differences between the *kinds of things represented* (for instance the difference between a species and its members). Second, there are syntactic theories that attribute the difference between specific and non-specific uses to differences in the structures that determine how the semantic constituents of a sentence combine to create the semantic content of the sentence as a whole. The approach I pursue here falls into the first category.

Over the last hundred years the syntactic approach to explaining non-specific uses has become standard, and is now rarely called into question.⁵ However, historically, what now

⁴We will focus on definite descriptions, but all types of singular definite noun phrases can be used non-specifically:

(a) Larry visits his mother on the 25th of December.

(b) Larry visits his mother on Christmas.

(c) Larry’s mother always looks forward to Christmas. It/That’s the only day she sees her son.

⁵Kripke (1977) lambastes the idea of trying to account for the semantics of definite descriptions without relying on the notion of scope. To deal adequately with the examples he uses in support of the need for scope would require a lengthy detour into the semantics of modals and indirect discourse (see Nichols (2013)

seems the naive approach to non-specifics was widely accepted. Aristotle took non-specific uses to represent the species and genera of which specific uses represent the members.⁶ In the middle ages there was considerable controversy concerning the nature of species and genera: realists, such as Scotus, held that they are universals in the category of substance, whereas nominalists like Occam (who did not believe that real entities could be universals) argued that species are actually mental symbols.⁷ Despite differences over what non-specific uses represent, there was agreement that they represent single entities that have some relevant relationship to the individuals represented by specific uses.

If non-specific uses of definite descriptions do not represent different kinds of entities than specific uses do, then syntax must explain the difference between specific and non-specific uses. If syntax explains the difference, then the syntactic structures of sentences containing definite descriptions must vary depending on whether they are specific or non-specific. This is what Russell proposed in “On Denoting”. Trampling over many details, we can say that Russell’s view is that definite descriptions are a type of existential quantifier. When used specifically, the quantifier (a sentential operator) is a modifier of the entire remaining formula (it takes widest scope). In contrast, when used non-specifically, the quantifier is part of a sub-sentence, which in turn is modified by at least one other sentential operator (it takes narrow scope).

Neo-Strawsonian views also appeal to syntax in attempting to explain the difference between specific and non-specific uses. In contrast to Russell’s theory of descriptions, however, neo-Strawsonian views hold that definite descriptions have a presuppositional meaning, and

for the latter). The thrust of Kripke’s argument is that a given definite description can be more than two-way ambiguous, so no binary distinction can account for the data. The distinction between a kind and its members is binary; however, from that it does not follow that a kind-reference account cannot handle cases of what Kripke describes as middle scope. Just as the same singular definite noun phrase can be used to refer to different individuals, so also the same phrase can be used to refer to different kinds.

Note that theories of generic noun phrases that treat certain definite descriptions as names of kinds are not exceptions to the current orthodoxy. Carlson (1977a) and those influenced by his work take generic noun phrases to be names, despite their un-namelike syntactic and morphological features. Thus on Carlson’s view, definite descriptions that designate kinds are thereby not really definite descriptions at all.

⁶Aristotle (1963).

⁷Spade (1994).

they designate the individuals that satisfy the expressed presuppositions. In general, a definite description expresses the presupposition that (i) the individual it represents has the properties represented by the description, and that (ii) that individual is the only thing that has those properties within the context in which the phrase is evaluated. The difference between specific and non-specific uses, on the neo-Strawsonian view, is supposed to result from a difference between the *contexts* within which the presupposition must be satisfied. Specific uses must have their presupposition satisfied within a single context (usually the context in which the expression is actually uttered). On the other hand, non-specific uses have their presuppositions satisfied within multiple contexts (which can be in the past, in the future, merely possible, or imagined); therefore, there is no single individual that they represent.

Crucially, in neo-Strawsonian theories, the same syntactic conditions that allow for a description to take narrow scope in Russellian theories are also responsible for supplying alternative contexts in which the presupposition of a definite description can be satisfied. That is, in the neo-Strawsonian framework, in order for a clause to have multiple contexts of evaluation, it must be within the scope of an appropriate operator (typically a quantifier).⁸ Thus despite their differences, Russellian and neo-Strawsonian views agree that non-specific uses can only occur within the scope of an appropriate operator.⁹ For this reason, the following argument applies to both Russellian and neo-Strawsonian theories of definite descriptions.

There are three fundamental considerations that bear on the choice between ontological and syntactic explanations of non-specific uses of definite descriptions. First, there is a question as to the semantic category within which non-specific uses of definite descriptions should be classified. If the two kinds of uses can be classified as two different kinds of

⁸See Stalnaker (1973); Heim (1990); Van der Sandt (1989); Schlenker (2008); Rothschild (2007).

⁹These include negation, subordinating verbs (such as “believes”), subordinating conjunctions, and quantifiers. From here forward, we will focus only on quantifiers and quantifier-like variable-binding operators. All of the phenomena we will consider occur in the absence of negation, subordinating verbs, and subordinating conjunctions.

expression, then the analysis of one type of use need not be related to the analysis of the other. This would narrow the range of cases for which syntactic theories would have to account. But in Sections 2.2.2.1 and 2.2.2.2, I show that whether a description is interpreted as specific or non-specific is determined by the same factors that influence which individual a specific use is interpreted as representing. The uniform way that specific and non-specific descriptions are interpreted is evidence that they are semantically uniform.

If it is true that specific and non-specific uses are semantically uniform, then the fact that the overwhelming majority of definite descriptions are used specifically is relevant to the analysis of non-specifics. Definite descriptions are singular, definite noun phrases, and they are typically used just as you would expect a singular definite noun phrase to be used, to represent a single entity that the speaker expects the audience to be able to identify. We should expect that non-specific uses are also used to represent single entities that the speaker expects the audience to be able to identify.

The second consideration is whether there is any evidence that non-specific uses represent things of some kind in the same manner that specific uses represent individuals. In Section 2.2, I extend the purely referential account developed for specific uses in Chapter 1 to non-specific uses. I show that the account of the communicational role of definite noun phrases defended there is able to predict when a definite description will be interpreted as non-specific just in case we consider both kinds and individuals as potential referents. From this we can infer that kinds are referents of non-specific uses of definite descriptions in just the same sense that individuals are the referents of specific uses.

The third fundamental consideration derives from an implication of syntactic accounts. Both Russellian and neo-Strawsonian theories of definite descriptions imply that non-specific interpretations can only occur under certain syntactic conditions – there has to be an appropriate operator within the scope of which the description is evaluated. From the point of view of traditional grammar, however, non-specific interpretations can occur in *any* syntactic environment in which singular noun phrases can occur, including as the subject of

the simplest clause types. Thus, simple sentences containing non-specific uses would appear to be counter-examples to syntactic accounts. Why haven't they been taken as counter-examples? Answering this question will lead us to the theory of generic sentences (the semantic theory of generalizations).

In Section 2.3, we will investigate whether the logical characteristics of generalizations expressed using non-specific definite descriptions constitute independent evidence for the kinds of syntactic structures implied by Russellian and neo-Strawsonian theories. We will find that there is no evidence of formal inference patterns of the sort that would be evidence of the existence of the phonologically unmanifest variable-binding operators needed for syntactic analyses of non-specific uses in simple sentences. Abstracting away from the specifics of my own proposal, I argue that non-specific uses cannot be accounted for syntactically. Whether non-specific uses refer to kinds, or to something else appropriately related to individuals, non-specific uses must designate something other than individuals. The syntax of English simply does not supply the right ingredients for a syntactic account of non-specific uses.

2.2 Descriptions as Referring Expressions

According to the picture that emerges from consideration of specific uses of definite noun phrases, there are no semantic rules that determine the referent of a referring expression. However, there are criteria for the appropriate use of a definite noun phrase, from which we can derive the conditions under which reference by a given expression to a given thing can be expected to occur.

Recall that the fundamental hypothesis from Chapter 1 gave us two fundamental criteria for the use of a definite noun phrase: the use of a singular definite noun phrase, NP, to refer to x indicates that x is the only individual that fits the relevance requirement of NP that also fits the qualitative requirement of NP.

In this chapter, the hypothesis we are considering is that when a definite description is used *non*-specifically, it refers to a kind. Thus we should expect the discourse relevance of kinds to be correlated with the interpretation of definite descriptions in the same way as the discourse relevance of individuals. The kind-reference hypothesis predicts that non-specific interpretations will occur when a kind is the most obvious choice amongst the potential referents of a definite description. We will test this hypothesis by manipulating the discourse relevance of potential referents including both kinds and individuals.¹⁰

2.2.1 Reference to Kinds

In order to systematically test the hypothesis that non-specific uses of definite descriptions refer to kinds, let us distinguish three types of non-specific uses: generic, narrow-scope and predicational.¹¹ The first two types, generic and narrow-scope, are defined according to an informal characterization of what they are used to talk about. On the other hand, predicational uses are distinguished by how they contribute to the message conveyed by an utterance. The two modes of categorization are conceptually independent and thus it might be possible for predicational uses to also be specific, generic or narrow scope. We will hold off on consideration of the interaction between the two dimensions of categorization

¹⁰As stated, the kind-reference theory implies that kinds, along with their members, have the characteristics that match the qualitative information conveyed by common nouns. Thus, for example, along with individual loons, the genus *Gavia* is in the extension of “loon”. On the theory being presented, the extension of a common noun is understood as the collection of things that the expression is conventionally used to represent. Thus from the fact that “the loon” can be used to refer to the genus *Gavia*, a candy loon, or a toy loon, it follows that the genus, the toys and the candy are all in the extension of “loon”. Note that it does not follow that there is a univocal interpretation of “is a loon” such that everything in the extension of “loon” is a loon.

Alternatively, one could limit the extension of “loon” to loons and rely on a mechanism of extended reference to explain how “the loon” is used to refer to things that are not in the extension of “loon”. Hobbs et al. (1993) discusses the sentence “The Boston office called” used to mean that someone from the Boston office called. It is common practice to refer to artifacts using the name of the person who either made or designed them. We are apt to interpret “There are three Rembrants in the museum” quite differently from “There are thee Rockefellers in the museum”. The notion of the extension of a noun would have to be stretched quite thin to account for these cases without some mechanism of extended reference. On the present view, the issue is not pressing; the extension of an expression plays no direct role in determining semantic content.

¹¹This is not an exhaustive typology. For example, consider, “Every boy removed the crust from his sandwich.” It does not seem that “the crust” can be understood as representing either an individual or a kind (Andrea Bianchi, p.c.). See Appendix A and Pepp et al. (2015) for analyses of cases like this one.

until Section 2.2.3.3, in which the predicational use will be analyzed. Before going into the analyses of the three types of uses it will be useful to look at some examples and make a few preliminary observations.

(2.5) The hotel is across the street from the train station. (specific)

(2.6) Tesla invented the alternating current motor. (generic)

(2.7) The president has always had an office in the capitol. (narrow-scope)

(2.8) *Titanic* is the most popular movie of all time. (predicational)

In (2.5) the subject, a specific use of “the hotel”, identifies the particular individual that is described as being across from the train station. Similarly, in (2.6), the generic “the alternating current motor” identifies the kind of thing that Tesla invented; and in (2.7), the narrow-scope “the president” identifies the political office whose holders have always had an office in the capitol.

In contrast to generic uses, narrow-scope uses pertain at most to one individual at a time (or in a given situation, possibility, etc...). (2.7) is naturally understood as saying something about all of the presidents of the United States, but only during the times when each was the one and only president; whereas (2.6) tells us something about all alternating current motors, including the millions of them that are currently in use. This difference corresponds to a difference in the nature of the two kinds *president of the United States* and *alternating current motor*: there can only be one president of the U.S. at a time but there can be many alternating current motors.

Predicational uses are distinguished not by what they designate but by how what they designate is incorporated into the semantic content of a complete utterance. In (2.8), “the most popular movie of all time” serves to describe the movie that has been identified by the subject phrase “*Titanic*”. Just as the sentence is not used to communicate that a certain movie is called “*Titanic*”, the predicate is not used to identify an individual. The sentence

is used to say something about the individual identified by “*Titanic*”, namely that it is the most popular movie of all time.¹²

2.2.2 Communication of Reference

The communicative role of a definite noun phrase is to convey the identity of its referent to the speaker’s intended audience. How much information the audience needs in order to be able to identify a given thing depends on both its perceived relevance to the conversation and its similarity to other relevant items. As the relevance of a given entity to a discourse becomes more apparent, less communicative work needs to be done to make it identifiable, and so a speaker will choose a less informative expression to refer to it. The audience expects there to be a reason for the inclusion of information in a referring expression (deviating further from using a pronoun), and so excess information will only lead them astray.

My hypothesis concerning non-specific uses of definite descriptions is that they refer to kinds. Whether a given definite description is interpreted as non-specific on a given occasion depends on three things: (i) the domain of entities the expression can be used conventionally to represent; (ii) the semantic content of the rest of the utterance to which the expression will contribute; and, (iii) the broader communicational context in which the utterance is interpreted. We will see that manipulating these three variables influences whether an expression is interpreted as non-specific, as predicted by the kind-reference hypothesis.

2.2.2.1 Generic Uses

Some noun phrases are more readily interpreted generically than others. In the following pair of sentences, the subjects are naturally understood to be generic uses.

¹²Note that not all predicate nominals (nominal complements of “is”) are predicational uses. See Higgins (1979), Doron (1988), and Section 2.2.3.3 below. This distinction is independent of the present claim about what “the most popular movie of all time” designates. Even if (2.8) is thought to express a logically symmetrical proposition such as an equation, or a statement to the effect that the subject nominal and the predicate nominal are co-referential, the predicational role of the “the most popular movie of all time” needs to be accounted for. The conventional communicative purpose of (2.8) is to characterize or describe something. “The most popular movie of all time” serves as a description; it tells about the subject.

(2.9) (a) The clownfish is mostly active at night.

(b) The common house cat is a menace to birds.

In contrast, out of the blue, with no stage setting, the subjects of the next pair of sentences are not easily interpreted generically.¹³

(2.10) (a) The fish is mostly active at night.

(b) The cat is a menace to birds.

Sentences (2.9 a) and (2.9 b) are more naturally understood as generalizations than (2.10 a) and (2.10 b), which seem to be about a particular fish and a particular cat. The most obvious difference between the subjects of the first and second of these pairs of sentences is that the descriptions in the latter pair are less specific. “Clownfish” is a word for a species, whereas “fish” is a word for a superclass. Both “cat” and “common house cat” can be used as words for the same kind of domesticated feline, but “cat” can also be used for the wider category that includes lions and tigers, and probably more relevantly, “common house cat” is an adjectivally modified phrase rather than a bare common noun.

In these out-of-the-blue examples, whether more or less specificity is required to get a generic interpretation depends on the standing degree of salience of a *kind* of thing as compared to the standing degree of salience of *members* of that kind. Individual cats are more cognitively salient relative to their kind as compared to the salience of individual clownfish, relative to the kind of creature they are. Most people know of several individual cats, and probably even know some of them by name. In contrast, I doubt that you, the reader, can identify a single individual clownfish. Thus when faced with a choice between a kind and a member of the kind, we lean further in the direction of individual cat than

¹³Sentences are always interpreted in a context. Within the framework we are considering, the significance of context is its influence on the expectations of the audience about what the speaker is trying to communicate. The context that you, as a reader, are in now has an influence on how you interpret these examples. Here an example is out-of-the-blue if it is offered without any background as to the context in which you should imagine it being used. Thus I am making claims as to how my audience will interpret these examples based on assumptions about what kinds of things they are likely to think about when given very little direction.

we do in the direction of individual clownfish. So, it takes more information to identify the kind *cat* than an individual cat, and less information to identify the kind *clownfish* than an individual clownfish.

A second influence on the interpretation of a definite description is its immediate linguistic context.

(2.11) (a) The clownfish is under the gray rock.

(b) The cat was first domesticated in Egypt some 10,000 years ago.

Sentence (2.11 a) has the same subject as (2.9 a), but unlike in (2.9 a), in (2.11 a) “the clownfish” is naturally interpreted as a specific use. On the other hand, (2.11 b) has the same subject as (2.10 b), but unlike in (2.10 b), in (2.11 b), “the cat” is naturally interpreted as a generic use. The relevant difference between (2.9 a) and (2.11 a) is that the predicate of (2.11 a) is something that doesn’t easily make sense as applied to a kind of animal, whereas it does make sense as applied to an individual animal. On the other hand, the predicate of (2.11 b) makes sense as applied to a kind of animal but not when applied to an individual cat. Here, linguistic context forces an interpretation that is the opposite of the default, out-of-the-blue interpretation.

Finally, it is when we begin to manipulate the larger context in which a complete sentence is interpreted that we can clearly observe that it is fundamentally the relevance to the discourse of potential referents that determines the interpretation of a definite noun phrase, whether it refers to an individual, or to a kind. By looking at how the interpretation of a sentence can be changed by embedding it within a larger discourse we can see that, just as introducing a particular individual into a discourse can lead to a specific reading with that individual as the referent, introducing a kind into a discourse can lead to a generic reading with that kind as the referent.

The following pair of sentences is used in the introduction to *The Generic Book* as evidence for the claim that only noun phrases associated with “well-established” kinds can

be interpreted generically.¹⁴

(2.12) (a) The Coke bottle has a narrow neck.

(b) The green Coke bottle has a narrow neck.

Sentence (2.12 a) is more naturally understood as a general statement than sentence (2.12 b). According to Krifka et al. (1995), the difference is explained by the fact that “the Coke bottle”, unlike “the green Coke bottle”, can function as a proper name of a kind. The idea is that because Coke bottles are thought of as constituting a “well-established” kind, with which “the Coke bottle” is associated, the kind can be named by that expression, which then admits of a generic interpretation. On the other hand, because Coke bottles that are green are not generally thought to constitute their own “well-established” kind, modifying the “Coke bottle” with “green” blocks its use as the name of a kind, and thus also blocks a generic interpretation.

In the next sentence, the subject seems to be too general to bear a generic interpretation.

(2.13) (a) The bottle is made of glass.

It looks like we have one description (“the green Coke bottle”) that is too specific to get a generic interpretation and one (“the bottle”) that is too general. But in fact, both of the examples that are hard to hear as generic out of the blue, can be forced into a generic interpretation by placing the sentence within a larger text that makes the appropriate kind relevant.

(2.14) Coca Cola uses different bottle shapes and colors for different beverages. In some cases there is more than one type of bottle for a given beverage. But, for a given beverage, the shape and color never vary independently. The green Coke bottle has a narrow neck; the clear one has a short, wide neck.

¹⁴My (2.12 a) and (2.12 b) are examples 24 a and b in Krifka et al. (1995). Originally from Carlson (1977a), via Barbara Partee.

(2.15) Cultures do not always come up with the same solutions to similar problems. All cultures have some way to store and transport liquids, but the type of container they use depends on the materials made available by other cultural practices. Cultures that eat meat tend to use skins to hold liquids. The bottle is made of glass. Only cultures that already used glass for some other purpose developed the use of bottles to store liquids.

In these two short texts, the descriptions that seemed either too specific or too general to get a generic interpretation in an out-of-the-blue context, are both naturally interpreted generically. Thus making a kind relevant to the discourse, and therefore making it more likely to be understood to be the referent of a definite noun phrase, can induce a generic interpretation.

The fact that generic interpretations are correlated with the discourse-relevance of kinds supports a referential account of non-specific uses of singular definite noun phrases. It is also evidence *against* the claim that being kind-referring is an exceptional trait of only a subclass of definite descriptions. The most influential theories of generic noun phrases hold that, properly speaking, definite descriptions never get generic interpretations; they only appear to because some expressions of the form “the CNP” function as proper names of kinds (which have failed to grow capital letters).¹⁵ Thus in Krifka et al. (1995), the semantic mechanism of specific uses is completely different from the semantic mechanism of generic uses – the former being generalized quantifiers and the later directly referential names.¹⁶ The uniform way that specific and non-specific definite descriptions are interpreted (via identification of an appropriate referent) is evidence against different underlying semantic mechanisms. Thus evidence that generic definite descriptions are referring expressions is also evidence that all definite descriptions are referring expressions.

¹⁵Carlson (1977a); Krifka et al. (1995).

¹⁶The view is a bit more complicated because it is supposed to be part of the semantics that they are names of kinds and thus that they have members (which comes into play in generic sentences).

2.2.2.2 Narrow Scope: Syntax Is Not Enough¹⁷

Now I want to start applying the lessons from the last section to sentences that have not traditionally been thought of as involving reference to kinds. We use the label “narrow scope” for the type of use that is the subject of this section because of Russell’s theory that narrow-scope uses are distinguished from specific uses by the relative scope of a quantifier associated with the description.

Scope is determined by syntax. Thus if a type of ambiguity is explained by scope, this implies that wherever the ambiguity arises the appropriate syntactic features are present.¹⁸ In this section we will see that syntactic conditions on their own do *not* give rise to so-called scope ambiguities. Instead, narrow-scope interpretations fit the same pattern as specific and generic interpretations: whether a definite description gets a narrow-scope interpretation is explained by the relevance of potential referents (both individuals and kinds).

Sentences that contain a definite description as well as a quantifier or a sentential modifier are often ambiguous. The following is a standard type of example.

(2.16) Every student got an A on the final exam.

This sentence could be used to state a fact involving a single exam (that every student got an A on it). Or, it could be used to state that every student got an A on whichever final exam he or she took. There are two well-known, competing explanations for this kind of phenomenon. According to Russellians, the ambiguity is due to there being two possible underlying syntactic structures. If definite descriptions are a kind of quantifier, then the specific reading, on which there is only one exam, can be explained as the result of the

¹⁷This section owes a lot to Strawson (1964), Doron (1988) and Rothschild (2007).

¹⁸The concept of scope derives from the language of predicate logic in which the syntactic features of a formula have determinate semantic consequences. Although the notion of scope applies originally to a syntactic relationship, the notion can be applied to languages in which the syntactic features required for scope in predicate logic are absent. Thus one could try to explain the different semantic possibilities of ambiguous sentences as resulting from different possible scopal relationships, while denying that for each scopal possibility there is a corresponding syntactic structure. However, such a theory would lack the explanatory power of a theory that associates semantic characteristics with observable features that can be used to facilitate communication. We can use the notion of scope as a description of semantic relationships, but then we cannot use the same notion to explain how those relationships come about, or how people know about them.

description raising above “every student” in the logical form of the sentence, thus taking wide scope. Leaving the description in place, below “every student”, yields the narrow-scope interpretation, on which there could be several different exams.¹⁹

Alternatively, neo-Strawsonians would explain the ambiguity as due to there being two possible domains in which the presupposition triggered by the definite description (that there is exactly one final exam) could be satisfied. If this presupposition is satisfied globally, within the domain of the discourse as a whole, then there will only be one exam designated by the expression. But, if the presupposition is satisfied locally, within the verb phrase, then, because the subject of the clause is a quantifier phrase, the presupposition needs to be satisfied for each of the individual students, possibly by different exams.²⁰

One thing that the Russellian and neo-Strawsonian views agree on is that the syntax of a sentence plays a crucial role in explaining definite-description-induced ambiguities. For the Russellian, the presence of quantifiers or other sentential modifiers (which are usually also given a variable-binding-operator analysis) gives the possibility of different quantifier scopes. Whereas for the neo-Strawsonian, the same variable-binding operators yield additional domains in which presuppositions can be satisfied.

Contrary to both Russellian and neo-Strawsonian views, we will see that syntax plays no essential role in how definite descriptions are interpreted. There are syntactic structures in which definite descriptions are rarely used non-specifically, but the role of syntax is always mediated by cognition and communication. The fact that the occurrence of these types of ambiguities is independent of the syntactic structure of the sentences in which they occur,

¹⁹Informally, the wide-scope logical form is something like “There is a unique final exam, and every student got an A on it.” And the narrow-scope version would be something like “Every student is such that there is a unique final exam on which he or she got an A.”

(a) $\exists x((\forall z \text{FinalExam}(z) \leftrightarrow z = x) \ \& \ \forall y(\text{Student}(y) \supset \text{Got_an_A_on}(y,x)))$

(b) $\forall y(\text{Student}(y) \supset \exists x(\forall z(\text{FinalExam}(y,z) \leftrightarrow z = x) \ \& \ \text{Got_an_A_on}(y,x)))$

Note that, unlike in (a), in (b) the uniqueness clause for the exam is relativized to the student. Thus the narrow-scope and wide-scope versions actually differ in more than quantifier scope.

²⁰The idea is that either the conversational participants presuppose that there is something in the situation represented by the discourse that fits the presupposition (global satisfaction), or they take it to be a presupposition of some situation that they are considering for the sake of interpreting a given clause of a sentence, which could be a counter-factual situation, a situation in the past or the future, or a more narrowly defined situation than that represented by the broader discourse.

is evidence of a syntax-independent mechanism at work.

With the following example, we will see that the syntactic ingredients that figure in both the Russellian and the neo-Strawsonian theories are not sufficient to predict when non-specific interpretations of definite descriptions will occur.

(2.17) For the last 5 years, the person by the window has been ignoring Susan.

Without any stage-setting, (2.17) is only naturally interpreted as stating that a particular person, the one by the window, has been ignoring Susan. The reason we don't get a narrow-scope interpretation is that "the person by the window" does not identify one of the many kinds of things that people usually think and talk about. However, by making things that fit the description "the person by the window" a conversationally relevant kind, we can induce a narrow-scope interpretation of the sentence.

(2.18) Susan went to the same bar every Friday, always hoping that some handsome man would turn away from the spectacular view of the bay and, seeing her face in the light of the setting sun, fall instantly in love with her. Sadly her dream remains only a dream. For the last 5 years, the person by the window has been ignoring Susan.

Well, there is always next Friday.

By embedding (2.17) in a story in which people who sit in a particular location have a special significance just in virtue of that attribute, a non-specific interpretation of "the person by the window" becomes natural. Thus narrow-scope interpretations, like specific and generic interpretations, are correlated with the conversational relevance of an appropriate referent. A referential account of definite descriptions provides a uniform analysis, without syntactic constraints.²¹

²¹Looked at in isolation, examples like (2.17) prove that syntactic conditions alone do not cause narrow-scope interpretations. Rothschild (2007) uses similar examples in his argument in favor of a neo-Strawsonian view. Standard neo-Strawsonian explanations of narrow-scope interpretations claim that the alternative relevant situations in which presuppositions can be satisfied are generated by a variable-binding operator within the scope of which the description is located. In Section 2.3, I will argue that the same type of ambiguity that is found in sentences with quantifiers (like (2.17)) can also be found in sentences with no quantifiers or other variable binding operators.

2.2.3 Kind-Reference Semantics

So far, we have seen that the way that definite descriptions are interpreted is best explained by the hypothesis that non-specific uses of definite descriptions refer to kinds. The goal of this section is show that pure reference can account for the semantic contributions of definite descriptions to complete clauses, without going beyond a schematic analysis of the contributions of other constituents. The first step towards a full semantic analysis of the contribution of an arbitrary definite description to the semantic content of an arbitrary sentence is an account of the contributions they make as subjects, objects, and complements of verbs.

2.2.3.1 Subjects

A subject and a predicate (a verb phrase) combine to form a clause. Thus for the case of definite descriptions in subject position, we need to explain how both individuals and kinds combine with the semantic content of predicates to form the propositional content of clauses.

(2.19) The common loon lays two eggs a year.

Sentence (2.19) can be used to express a general statement about the species *Gavia Immer*, or a statement about a member of the species. On the specific reading, stating the truth conditions using a purely referential analysis is simple: an utterance of the sentence is true if, and only if, the thing (the particular bird) referred to by “the common loon” lays two eggs a year.

Because the pure-reference view holds that the subject of (2.19) is also referential on the non-specific reading, it would seem that the truth conditions would be the same, with the exception of which thing does the egg laying. But it is not quite that simple: species of birds don’t lay eggs, at least not in the same way that individual birds do. Thus the pure reference view implies that “lays two eggs a year” sometimes designates one trait that holds of birds, and sometimes another trait that holds of species of birds.

In fact, we can find independent evidence that “lays two eggs a year”, along with many other predicates, can designate one characteristic that holds only of individuals, and also another that holds only of kinds.²² For instance, it is not difficult to find examples like the following:

(2.20) *Gavia immer* swims underwater to catch fish and propels itself with its feet.²³

The subject of (2.20) is the proper name of a species, and the verb is inflected for a singular subject. It looks very much like (2.20) ascribes a trait to the species *Gavia Immer*. If that is right, then “swims underwater to catch fish” must represent a trait of a species, not of a particular animal. The predicate ambiguity that is implied by a kind-reference analysis of non-specific definite descriptions is thus already implied by a referential analysis of proper names. As a general rule, when predicates that are typically used for individuals are used to describe kinds, we can understand them as attributing, to the kind, the property of having members that have the characteristic attributed to individuals by the same expression.

The kind-reference account of non-specific subjects is uniform no matter what the content of the predicate is.

(2.21) The dodo is extinct.

(2.22) The whale is warm blooded.

(2.23) The first date is awkward.

Each of these three sentences have the same basic subject-predicate surface form, and the kind-reference theory analyses them all as attributing the characteristic represented by the predicate to the entity represented by the subject. As we will see in Section 2.3, syntactic accounts of non-specific uses cannot treat sentences like these uniformly, and they cannot

²²An alternative possibility is that the trait designated is general enough that it can hold both of birds and species of birds. Either choice would be compatible with a pure-reference view of definite singular terms, though the ambiguity option fits better with the overall project of which this paper is a part. For a defense of the alternative path see Liebesman (2011).

²³<http://sanctuarysimon.org/species/gavia/immer/common-loon>.

account for their semantic content without introducing ad hoc constituents into their logical forms.

2.2.3.2 Objects

The kind-reference treatment of objects of verbs is exactly as you would expect. A simple example will suffice to make it clear.

(2.24) The Framers of the Constitution gave the President the power to veto acts of Congress.²⁴

On the most natural interpretation of sentence (2.24), “the President” is non-specific – it refers to the kind that has the current President of the United States as its member. Thus on the natural interpretation, the predicate describes the referent of the subject as having given the power to veto acts of Congress to the kind *President of the United States*. Again, the fact that the kind has the characteristic represented by “the power to veto acts of Congress” is intimately related to the fact that its members have the characteristic represented by the same expression on the less natural reading in which “the President” is specific.

2.2.3.3 Predicates

A predicative noun phrase combines with a verb to form a predicate. So for predicate nominals, we would like to know how the referent of the noun phrase contributes to the content of the predicate as a whole.

Definite descriptions in predicate position do not generally give rise to ambiguities that result in divergent truth conditions, but we can still distinguish between two uses. Consider the contrast between the following examples.

(2.25) (a) The guy in the corner is the mayor. But, the woman next to him could have been the mayor, if I had told a certain story to the media.

²⁴<http://www.archives.gov/legislative/resources/education/veto/>.

- (b) The guy in the corner is the person I was telling you about. But, the woman next to him could have been the person I was telling you about, if I had told you a different story.

The second sentence of (2.25 a) says that a particular person could have been something (mayor) that the first sentence says a different person actually is. Although (2.25 b) has the same structure as (2.25 a), given the natural interpretation of the first sentence, we cannot make sense of the second sentence. The difference is due to the fact that the predicate nominal of the first sentence of (2.25 a) is predicational, whereas in (2.25 b) it is specific. The person referred to by “the person I was telling you about”, would not have been a different person no matter which story was told. In contrast, in order to interpret the second sentence of (2.25 a) as a follow up of the first sentence, what matters is that “the mayor” refers to the same kind in both sentences, not to the same person.

Both specific and predicational predicatives form predicates that function as descriptions of a previously identified entity – namely, the referent of the subject phrase. In part, the difference between how predicational predicates and specific predicates describe the referent of the subject is due to a difference between the kinds of things they refer to. But there is more to it than that. The relationship that “is the person I was telling you about” describes the guy in the corner as having to the person I was telling you about is a *different relationship* than the relationship “is the mayor” describes the guy in the corner as having to the kind *mayor (of such and such)*. On the specific interpretation of (2.25 a), “is” means something like “is identical with”, whereas on the predicational interpretation of (2.25 b), “is” means something like “is a/the member of”.

Both the specific/predicational and the definite/indefinite predicational distinctions also apply when the subject refers to a kind.

(2.26) The triangle is the shape I was telling you about.

(2.27) The Toyota Camry is the best selling car.

In (2.26) “the shape I was telling you about” refers to a kind, but it is a specific use because it refers to the same kind as the subject “the triangle”. In (2.27) “the best selling car” refers to a kind of which the Toyota Camry is the sole member.

“Is” has to denote two different relations, if it is to conjoin both expressions that refer to things at the same taxonomic level (identificational), and expressions that refer to things at different taxonomic levels (predicational). Thus one might charge the kind-reference account of predicational uses with “multiplying meanings”. However, there is independent evidence for both of these meanings of “is”.

(2.28) (a) This is Linda.

(b) Linda is the kind of person who will never let you down.

Whatever the analysis of the distinction between (2.25 b) and (2.25 a), we need the “is” of identity for (2.28 a) and a different meaning of “is” for (2.28 b). (2.28 b) does not assert that Linda is identical with a kind of person, and (2.28 a) does not assert that the referent of “this” is a member of a kind, or that she has certain characteristics. This ambiguity in the meaning of “is” is just the long-recognized distinction between the “is” of identity and the “is” of predication.²⁵ This distinction is needed regardless of the semantics of definite descriptions.²⁶

²⁵See Aristotle (1963) for the *locus classicus* of both the ambiguity of “is”, and the type of ambiguity involved in predicating “lays two eggs a year” of both a bird and a species of bird. Aristotle points out that we mean something quite different when we say something *is a bird*, rather than *is in a tree* or *is white*. These differences are effaced by model-theoretic semantic theories that treat all predications essentially as claims about set membership. While this kind of generalization has its place, we should not lose sight of the fact that it involves abstracting away from significant differences between different kinds of predication. And the fact that it is possible to treat “is” as if it were univocal is not adequate grounds for treating NPs in whatever way such a univocal treatment of “is” requires. For arguments that “is” does not have a univocal semantics see Doron (1988).

²⁶This not an analysis of copular constructions; it is just a gloss of the distinction between two uses of *is*. See Higgins (1979); Doron (1988) for a great deal of evidence for there being such a distinction. See Nichols (2013) for a more in-depth look at the semantics of copular constructions as it relates to the semantics of noun phrases.

2.3 No Syntax Required

Both Russellian and neo-Strawsonian theories hold that non-specific uses of definite descriptions are (at least in part) the result of the description being within the scope of a variable-binding operator (or other sentential operator). Prima facie, there is a straightforward way to falsify these theories: show that definite descriptions can get non-specific interpretations in the absence of quantifiers, other variable binders, or sentential operators of any kind. Syntactic accounts predict that non-specific interpretations of definite descriptions will only occur in the right syntactic environments. However, non-specific uses can occur in *any* syntactic environment in which a singular noun phrase can occur.

In this section I focus on non-specific interpretations of the subjects of sentences consisting of a single canonical clause.²⁷ None of the following sentences contains a quantified noun phrase, an adverb, a sentential connective, clausal subordination, or any constituent that looks like a quantifier.

(2.29) The compressor vibrates.

(2.30) The goalie wears extra equipment.

(2.31) The 15th of March is unlucky.

(2.32) The president is becoming an absolute monarch.²⁸

(2.33) The Volkswagen Beetle floats.

(2.34) The desert tortoise hibernates.

(2.35) The OED is expensive.

(2.36) The Canada goose is nesting on Lake of the Woods.

²⁷Canonical clauses are not subordinate, not co-ordinate, positive (containing no negations), declarative, and have standard constituent order (Huddleston and Pullum (2002)).

²⁸Jerrold Nadler on the floor of the House of Representatives, 24 October 2014.

The subject phrases of each of (2.29) - (2.36) can naturally be interpreted as either specific or non-specific: the subjects of (2.29) - (2.32) are easily interpreted as narrow-scope uses, whereas the subjects of (2.33) - (2.36) are easily interpreted as generic uses. These sentences represent a variety of the simplest clause structures of English, and so, if there are any clausal constructions that do not contain variable-binding operators, it would be reasonable to suppose that at least one of these sentences had that structure.

Prima facie, (2.29) - (2.36) constitute decisive evidence against syntactic accounts of non-specific uses. Examples like these show that there is no correlation between non-specific uses of definite descriptions and the syntactic structures that figure in traditional grammatical descriptions of English. Be that as it may, in order to make an argument against standard contemporary semantic theories, it is necessary to take on the rather complicated task of proving that there is no underlying, un-pronounced, syntactic feature that is correlated with non-specific uses of definite descriptions. A thorough assessment of the syntactic approach to non-specific uses requires a rather in-depth look into contemporary theories of the logical form of sentences that express generalizations.

2.3.1 Limitations of the Syntactic Approach

Before turning to the analysis of generalizations, I would like to call attention to a limitation faced by syntactic accounts of non-specific uses of definite descriptions. There is a class of examples for which neither Russellian nor neo-Strawsonian theories seem to have an explanation: non-specific uses that cannot be interpreted as generalizing over specific uses. The following sentences attribute characteristics to kinds, but cannot be understood as distributed attributions of those characteristics to the individual members of those kinds.

(2.37) The peregrine falcon ranges across both the Old and New World.²⁹

(2.38) The steam engine gave way to the internal-combustion engine as a means of vehicle

²⁹Groombridge et al. (2002).

propulsion.³⁰

(2.39) The Camry comes in four intriguing trims.³¹

(2.40) In the 109th Congress, legislation was introduced to place the Secretary of Homeland Security into the line of succession after the Attorney General but that bill expired at the end of the 109th Congress and was not re-introduced.³²

The analyses of non-specifics given by Russellians and neo-Strawsonians do not seem to apply to sentences which, like (2.37) - (2.40), are true on a non-specific interpretation, but not because they are ever true on a specific interpretation.³³ According to syntactic accounts, sentences with non-specific subjects should be generalizations of open sentences derived from the predicate of the sentence.

(2.41) x ranges across both the Old and New World.

(2.42) In the 109th Congress, legislation was introduced to place x into the line of succession after y but that bill expired at the end of the 109th Congress and was not re-introduced.

Applying a syntactic explanation, we should evaluate the truth of (2.37) and (2.40) by evaluating (2.41) and (2.42) with respect to values of “ x ” and “ y ” that satisfy the presuppositions expressed by the phrases “the peregrine falcon”, “the Secretary of Homeland Security”, and “the Attorney General” – that is, with respect to peregrine falcons, Secretaries of Homeland Security and Attorneys General. But such an analysis just doesn’t work for examples like (2.37) - (2.40). (2.37) and (2.38) cannot be understood as a generalization of “This peregrine falcon ranges across both the Old and New World” and “This steam engine gave way to the internal-combustion engine as a means of vehicle propulsion” – no single falcon has

³⁰<http://www.britannica.com/EBchecked/topic/564472/steam-engine>.

³¹<http://www.newcars.com/toyota/camry/2010>.

³²Wikipedia (2014c).

³³The difficulty posed by examples like these for Russell’s theory of descriptions was recognized by Moore (1944); Geach (1950).

such a wide range, and no particular steam engine(s) gave way to the internal-combustion engine as a means of vehicle propulsion.

In his influential early work, Gregory Carlson concluded that there is no way to account for the semantic content of examples like these while maintaining that definite descriptions can only represent the individuals that satisfy the presuppositions posited by neo-Strawsonian accounts.³⁴ However, Carlson also concluded that because examples like these must be interpreted as referring to kinds, they are not actually definite descriptions. Despite the fact that they have the syntactic and morphological features of definite descriptions, the standard analysis (following Carlson) is that these phrases are *proper names* of kinds.

Insofar as I take names and definite descriptions to be semantically uniform, I agree with Carlson's analysis. However, it is a mistake to think that being kind-referring is an exceptional property of only a sub-class of definite descriptions. As we saw in Section 2.2.2.1, *any* definite description can refer to a kind: definite descriptions are used conventionally to refer both to individuals of a certain kind, and to the kind of which those individuals are members. Syntactic theories of non-specific uses cannot eliminate the need for kinds in the ontology, and there are no grounds for quarantining the trouble makers by reclassification – a semantic theory of definite descriptions must be able to account for examples like (2.37) - (2.40).

In sections 2.2.2.1 and 2.2.2.2, we saw that whether a definite description will be interpreted as specific or non-specific depends on the discourse-relevance of potential referents (including both kinds and individuals), and is not predictable on the basis of syntactic structure alone. We have now reached the further conclusion that some definite descriptions can only be interpreted as designating kinds. We now turn to the question of whether syntactic accounts can handle even those non-specific uses that can be understood as generalizations about individuals.

³⁴Carlson (1977a).

2.3.2 Argument by Analogy

Setting aside for the moment cases that can only be analyzed as designating kinds, we can think of the remaining cases as, in some intuitive sense, generalizations. One way to illustrate the generality of such sentences is by comparison with sentences that have an adjunct that explicitly indicates that the clause they modify expresses something general.

(2.43) (a) The CEO has the final say.

(b) The CEO always has the final say.

In (2.43 b) “always” indicates that the sentence expresses something general in that it represents a pattern of events. On their most natural, non-specific, interpretations, (2.43 a) and (2.43 b) mean more or less the same thing.³⁵ Therefore what “always” contributes explicitly to (2.43 b) is somehow implicitly expressed by (2.43 a). One might infer that if what “always” contributes to the semantic content of (2.43 b) is implicitly expressed by (2.43 a), then the semantic value of “always” is in fact part of the semantic content of (2.43 a), albeit implicitly.

David Lewis’s hypothesis that adverbs like “always”, “typically”, and “normally” are unselective variable-binding operators is widely accepted.³⁶ If Lewis’s analysis is correct, then the above argument by analogy leads to the further conclusion that, like (2.43 a), (2.43 b) contains a Lewisian unselective binder. However, in order for this line of reasoning to be relevant to the semantics of definite descriptions, we would also need to be able to apply it to all simple sentences in which definite descriptions are used to make generalizations.³⁷

As we will see, this cannot be done.

³⁵This is one of five (non-criterial) tests of genericity given by Krifka et al. (1995).

³⁶See Lewis (1975). A Lewisian analysis of “Cats usually drink milk”:

(a) Usually(If x is a cat & s is a possible milk-drinking situation & x is in s , then x is drinking milk in s)

In (a), “usually” is an unselective variable-binding operator – it binds all the free variables in its scope. “Usually(P)” is true if and only if “P” is true for most assignments of values to its free variables. So, (a) is true if and only if “If x is a cat & s is a possible milk-drinking situation & x is in s , then x is drinking milk in s ” is true for most assignments of values to “ x ” and “ s ”.

³⁷Though I focus on the fact that this line of thought cannot be generalized, even for friendly cases it is suspect. It is not clear why we should infer that the semantic value of “always” is part of the semantic content of (2.43 a) from the approximate synonymy of (2.43 a) and (2.43 b). Why shouldn’t we suppose that there are just two different ways (at the level of logical syntax) of expressing more or less the same thing?

The above argument by analogy assumes, crucially, that sentences (2.43 a) and (2.43 b) can be used to express more or less the same thing on both a specific and a non-specific interpretation of “the CEO”. When this is the case, we can consistently hypothesize that “always” is an implicit constituent on either interpretation, and explain the difference between the two as a result of scope. However, in some cases, there clearly is no implicit “always” (or other appropriate adverbial adjunct) in the specific interpretation.

(2.44) (a) The treatment kit comes with an instructional video.³⁸

(b) The treatment kit always comes with an instructional video.

Given a non-specific interpretation of “the treatment kit”, (2.44 a) and (2.44 b) mean more or less the same thing. Thus we might infer that what “always” contributes explicitly to (2.44 b) is implicitly part of the logical form of (2.44 a). However, on a specific interpretation of “the treatment kit”, the two sentences do not mean the same thing.³⁹ If you buy a treatment kit, it will come with an instructional video, but presumably that will happen only once per kit. There is no implicit “always” in the specific interpretation of (2.44 a), thus the difference in semantic content between the specific and non-specific interpretations of (2.44 a) cannot be the result of differences in scope with respect to “always”. Accounting for the non-specific interpretation in terms of scope would require positing the addition of a constituent that is not part of the logical form of the sentence on a specific interpretation of the subject. Evidence of such an additional constituent would have to come from beyond the theory of definite descriptions. This brings us to the theory of generics.

³⁸<http://superhairenergizer.net/products.html>.

³⁹For some readers, getting the specific interpretation of (2.44 a) may require some stage setting. Imagine a difficult customer asking, “Do they all come with an instructional video?” Sales clerk: “Yes, they all do.” Customer: “What about that one on the shelf?” Sales clerk: “Yes, that one comes with an instructional video.” Customer (to his wife, who has not been paying attention), pointing at the treatment kit on the shelf: “The treatment kit comes with an instructional video.”

With a creative imagination, it is possible to get further readings, including a specific reading of (2.44 a) that would not be greatly altered by the addition of “always”. Think of a treatment kit that carries an instructional video with it wherever it goes. However, this is not the specific version of the natural (and intended in the advertisement from which it is quoted) non-specific interpretation of the same sentence.

2.3.3 The Logic of Generalizations

There is a sub-field of semantics, the theory of generics, devoted to accounting for the truth conditions of the sort of sentences we are now focused on: sentences that can be understood as generalizations but which do not contain (phonologically manifest) constituents that are theorized to be variable-binding operators (such as “every”, “always” or “possibly”). Generic sentences can be either generalizations about events (including the actions of a single individual), or generalizations about the members of a kind.

(2.45) (a) Elmo talks about himself in the third person.

(b) Elmo talked about himself in the third person the first time he said anything in “Sesame Street 4-D”.

(2.46) (a) College football players are employees.

(b) Justin Worley is an employee.

(2.45 a) describes the manner in which Elmo talks, but not on any particular occasion. (2.46 a) describes college football players, but it does not represent any particular player. Informally we can characterize sentences (2.45 a) and (2.46 a) as generalizations of sentences like (2.45 b) and (2.46 b). What (2.45 a) and (2.46 a) describe in general, (2.45 b) and (2.46 b) represent particular cases of: a particular occurrence of Elmo talking about himself, and a particular college football player.

We can also consider the relationships between (2.45 a) and (2.45 b), and (2.46 a) and (2.46 b) from a logical point of view. Loosely, (2.46 a) implies (or is evidence for) (2.46 b), whereas a sufficient number of particular facts like that expressed by (2.45 b) implies (2.45 a). (As we will see, the transition to a less vague statement of the logical relationships involved is not trivial.)

A fundamental goal of the theory of generics is to give a semantic analysis of generic sentences that will disclose the logical relationships between them and syntactically related

non-generic sentences.⁴⁰ Within the current mainstream of semantics, the hypothesis that a semantic analysis of (2.45 a) will disclose its logical relationship to (2.45 b) amounts to the hypothesis that in (2.45 a), a variable-binding operator binds a variable in the logical form where (2.45 b) has a constant term (and similarly for (2.46 a)).

(2.47) GEN1(Elmo talked about himself on occasion x)

(2.48) GEN2(for x such that x is a football player, x is an employee)⁴¹

Both simple present tense verbal predicates and bare plural subjects are typically used to express generalizations. Thus, we might hypothesize that “GEN1” is part of the semantic content of simple present tense verbal predicates, and “GEN2” is part of the semantic content of bare plurals in subject position. However, note that if “GEN2” were part of the content of bare plural noun phrases, then the explanation of the generality of (2.46 a) would not carry over to sentences with singular definite subjects, and further, we would need some explanation of why bare plural subjects are not always generic.⁴² A given predicate can form a generic sentence with a variety of types of phrases as subject. Thus, if there is a generic variable-binding operator that figures in narrow-scope interpretations of singular definite noun phrases, it must be part of the content of the predicates of generic sentences.

Generic sentences back logical inferences to specific sentences. Sentences (2.49 a) and (2.50 a) are generic, and (in some sense) they back inferences to (2.49 b) and (2.50 b).

(2.49) (a) Ravens eat carrion.

(b) This raven eats carrion.

(2.50) (a) Ravens are intelligent.

(b) This raven is intelligent.

⁴⁰Krifka et al. (1995).

⁴¹Note that “GEN2” is a binary operator, relating two open sentences “ x is a football player” and “ x is an employee”.

⁴²There are also well known difficulties involved in analyzing bare plurals as quantifiers. See Carlson (1977b).

Not all sentences with bare plural subjects are generic. Sentences (2.51 a) and (2.52 a) are not generic, which can be seen from the fact that they do not imply specific counter-parts.

(2.51) (a) Ravens are picking through the dumpster.

(b) This raven is picking through the dumpster.

(2.52) (a) Ravens are available (at the pet store).

(b) This raven is available (at the pet store).

Sentences with kind-referring subjects that do not back logical inferences to non-generic counter-parts are not generic either.

(2.53) The mini skirt is making a bold comeback.⁴³

(2.54) *Circaetus beaudouini* occupies a relatively narrow band of sub-Saharan Africa.⁴⁴

(2.55) This species occurs naturally throughout the mountains and into the upper Piedmont.⁴⁵

If there is in fact a difference at the level of logical form between the predicates of generic and non-generic sentences, then whatever explains the difference between generic and non-generic sentences with bare-plural, definite description and kind-name subjects should also be found in sentences with individual-referring proper name subjects.

Sentences (2.56 a) and (2.57 a) have predicates that form generic sentences with bare plural subjects, so they should include the generic operator that backs inferences to non-generic sentences.

(2.56) (a) Pat eats fish.

(b) Pat is eating fish.

⁴³<http://www.mummyofstyleandsubstance.com.au/autumn-styling-with-birdsnest/>.

⁴⁴<http://www.birdlife.org/datazone/speciesfactsheet.php?id=31639>.

⁴⁵<http://nc-forestry.stores.yahoo.net/whitpinim2yr.html>.

(2.57) (a) Pat is intelligent.

(b) This Pat is intelligent now.

When the subject of a generic sentence refers to an individual there cannot be a non-generic analogue along the dimension of individual/kind. However, thinking of simple present verbal predicates as a model, we can take such sentences as having non-generic analogues that are particular as to events rather than as to individuals. Thus if the distinction between generic and non-generic sentences is the result of a difference between generic and non-generic predicates, then the underlying mark of being a non-generic predicate is being particular as to events. Following this line of reasoning, we have to think of “is intelligent” as generalizing over particular events of being intelligent. In this way we can think of (2.57 a) as implying its non-generic analogue (2.57 b) following the same rule that leads from (2.56 a) to (2.56 b).

In the case of the non-generic predicate “is available”, we get the expected result that (2.58 a) being true at some time is not particularly good evidence that (2.58 b) is true at a different time.

(2.58) (a) Pat is available.

(b) Pat is available now.

In order to include sentences like “Pat eats fish” (2.56 a) and also like “Ravens are intelligent” (2.50 a) in the category of generics, it seems like we need to use two different criteria.

Conjecture 1. *A sentence is generic if it implies a non-generic version of itself being true an appropriate number of times. If “Pat eats fish” is true then “Pat is eating fish” is true for an appropriate number of times of evaluation.*

Conjecture 2. *A sentence is generic if its truth implies (in some sense to be defined) the truth of non-generic versions of itself. “Ravens are intelligent” implies “This raven is intelligent”.*

An initially promising approach for a unified account of the difference between (2.50 a) and (2.52 a), and the difference between (2.56 a) and (2.56 b), is to hypothesize that sentences that represent particular events have a time (or event or situation) variable that is supplied a value by the context of utterance, while sentences that do not represent particular events bind that same variable with a quantifier.⁴⁶ This hypothesis implies that the following criterion for being *non-generic*.

Conjecture 3. *A sentence that is true if, and only if, a particular event is occurring at the time of evaluation is not generic. “Pat is eating fish” is true if and only if there is an event occurring (Pat eating fish) at the time of evaluation, so it is non-generic.*

2.3.3.1 Going Below the Surface

Syntactic accounts of non-specific uses imply that non-specific interpretations will only occur when a definite description is in the scope of a variable-binding operator. Thus, the fact that non-specific interpretations occur in sentences without constituents that seem to be variable-binding operators looks like dispositive evidence against syntactic accounts. We are now considering the possibility that the theory of generics provides evidence that the apparent counter-examples to syntactic theories of non-specific uses of singular definite descriptions are *merely* apparent.

Simple present-tense sentences with verbal predicates are typically used to convey generalizations, and so it is reasonable to suppose that there is something like a generic variable-binding operator in their logical forms. However, as we have seen, definite descriptions can also be used non-specifically in other syntactic structures, like clauses with adjectival and prepositional predicates (e.g., “is fast” and “is under the table”). But it is not a characteristic fact about adjectival and prepositional predicates that they are used to convey generality. Thus evidence that there is an implicit variable-binding operator in simple present tense sentences with verbal predicates is only relevant to the analysis of non-specific definite de-

⁴⁶This is the approach developed by Kratzer (1995).

scriptions insofar as it supports the more general claim that the class of generic sentences, which cuts across traditional syntactic categories, is characterized by an implicit variable-binding operator. I will call this “the generic hypothesis”, and I turn now to consider it.

As we have seen, some copular clauses with singular definite subjects can be interpreted as generalizations, and some cannot. That there is no surface-level syntactic feature that marks the difference can be seen from examples like the following.

(2.59) He is orange.

Whether (2.59) is generic or not depends on who the referent of “he” is.⁴⁷ If he is my neighbor who just had the spray tan disaster, it is not generic. If he is my neighbor’s orange cat, then it is generic. What makes the difference is that being orange is a long-term, stable characteristic of the cat, and only a temporary state of my neighbor. Hence another test for being non-generic is whether a sentence makes sense when embedded in a “when” conjunction.⁴⁸

(2.60) When he is orange people laugh at him.

Whether (2.60) makes sense depends on whether it is about an orange cat or my temporarily orange neighbor. Examples like these are adduced as evidence that there are two predicates associated with the phrase “is orange”, one that is generic and one that is not.⁴⁹

⁴⁷The sense in which it “depends” on the referent is that the use of the sentence to talk about one thing rather than another is correlated with being interpreted in a way consistent with being classified as generic. If the generic hypothesis is correct then, in a different sense of “depends”, being generic depends on which of two possible semantic values “is orange” has.

⁴⁸See Kratzer (1995). Kratzer’s description of the test is that if a sentence makes an ungrammatical “when”-conjunction, then it is generic. However, her “bad” examples of this kind of sentence are plainly grammatical, there is just no reasonable interpretation of them. (For example, “When Harry has long arms he can reach the light bulb.”) The charge of ungrammaticality is in fact crucial to the claim that “Morris (the cat) is orange” is generic. The proposed reason for “When Morris is orange people laugh at him” being ungrammatical is that being generic is supposed to involve having a variable-binding operator that binds the hypothesized additional variable supplied by the generic version of “orange”. But, in fact, it is grammatical, so there is no evidence of an underlying ill-formed logical form.

⁴⁹Carlson (1977b) proposed an operator that transformed individual-level predicates into stage-level predicates. On Carlson’s original theory the two versions of (2.59) would be:

- (a) $\exists x(\text{Stage}(x, [[\text{he}]]) \ \& \ \text{Orange}_1(x))$
- (b) $\text{Orange}_2([[\text{he}]])$

The standard terminology is that the predicates of generic sentences are individual-level predicates in

The hypothesis that the same verb phrase sometimes represents a generic predicate and sometimes a non-generic predicate is also supposed to explain the contrast between sentences like the following:

(2.61) (a) The spare tire is in the trunk.

(b) The forensic examiner is in the trunk.

It is relatively easy to interpret (2.61 a) as either specific or non-specific, whereas (2.61 b) seems only to have a specific interpretation. If the generic hypothesis is correct, then the fact that the subject of (2.61 a) has a non-specific interpretation, while the subject of (2.61 b) has only a specific interpretation, is due to the fact that “in the trunk” is a generic predicate in (2.61 a) and a non-generic predicate in (2.61 b). That is to say, if the generic hypothesis is correct, the fact that being told (2.61 a), but not (2.61 b), on Monday is a reasonable basis for inferring that it will be true on Tuesday is a reflection of the fact that “in the trunk” means, strictly and literally, one thing in (2.61 a) and something else in (2.61 b).

There is no direct evidence of a (relevant) syntactic difference between (2.61 a) and (2.61 b), but according to standard theories of generics, we can infer that there is one because only such a syntactic difference can account for the logical difference between generic and non-generic sentences. Thus we are now looking for a pattern of logical facts that is best explained by the presence of a variable-binding operator in all generic sentences.

2.3.3.2 Simple Present and Present Progressive

When looking for a pattern of logical facts that would indicate the presence of an implicit, generic variable-binding operator, the simple present is a natural place to start. Sentences in the simple present tense can typically be understood as generalizations.

contrast to stage-level predicates. The terms go back to Carlson (1977a). Carlson proposed that some predicates apply to stages of individuals and some to the individuals themselves. Carlson categorized both time slices of four-dimensional concrete individuals and individual members of kinds as stages: what I call “individuals” are, according Carlson (1977a) individuals with respect to their time slices and stages with respect to the kinds of which they are members. In the literature, the terms “individual-level” and “stage-level” admit of a variety of interpretations, but they always mark a distinction between predicates that form generic sentences and those that do not.

(2.62) (a) Lola runs.

(b) Lola is running.

Sentence (2.62 a), on its most likely interpretation (on which it means at least nearly the same thing as “Lola is a runner”), is general in the sense that it says that Lola performs a type of action, but it does not specifically represent any particular occurrences of her performing actions of that type. If (2.62 a) is true, then, with some regularity, Lola must actually be running. This implication of (2.62 a) is paralleled by its syntactic relationship with (2.62 b).

Sentence (2.62 b) is the present progressive version of (2.62 a). And another way of describing the generality of (2.62 a) is to say that if it is true, then with some regularity, (2.62 b) must also be true. If we hypothesize that the syntax of (2.62 a) and (2.62 b) encodes the implication we have noted, then we might further predict that the logical form of the two sentences is something like the following.⁵⁰

(2.63) (a) $\text{Qt}(\text{Run}(\text{Lola}, t))$

(b) $\text{Run}(\text{Lola}, t_0)$

The idea here is that the verb denotes a relation that holds between individuals and times (the set of all individuals and times such that the individual is running at that time). The present progressive tense sets the value of the time variable (t) to the present, so (2.62 b) is true just in case Lola is running now ($\langle \text{Lola}, \text{the time of utterance} \rangle \in [[\text{run}]]$). On the

⁵⁰The current state of the art is that generic sentences are formed with a binary operator of the same kind that Lewis proposed for adverbs like “always” and “usually”. Nothing in my argument rests on the particulars of the variable-binding approach, so I will leave things as simple as possible. For the record, here is how Krifka et al. (1995) would formalize a few simple examples.

(a) Lola runs.

(a') $\text{GEN}[s,x](\text{appropriate-for-running}(s) \ \& \ x = \text{Lola} \ \& \ x \text{ is in } s; \text{running}(x,s))$

(b) Cats are furry.

(b') $\text{GEN}[s,x](\text{Cat}(x) \ \& \ x \text{ is in } s; \text{furry}(x, s))$

(c) Cats sleep all day.

(c') $\text{GEN}[s,x](\text{Daytime}(s) \ \& \ \text{Cat}(x) \ \text{and } x \text{ is in } s; \text{sleeping}(x,s))$

“GEN” is a binary operator that takes two open sentences as arguments. The first open sentence (before the “;”) is the restrictor and the second is the matrix clause. For a summary of the reasons for which the originally proposed monodic operator has been replaced by a binary operator, see Krifka et al. (1995).

other hand, the simple present is a quantification over times, so (2.62 a) is true just in case the open sentence “Running(Lola, t)” is true for an appropriate number of values of “ t ” ($Qt(\langle \text{Lola}, t \rangle \in [\text{run}])$). Thus we have the makings of a theory tying the semantics of the plain present tense to (i) its use for making general statements and (ii) its semantic relationship to the present progressive: the simple present is formed by binding the time variable in the verb with a quantifier. If this outline of an analysis is correct, then it is the first step towards verifying the hypothesis that generality is a result of quantification.

On the basis of the first example (2.62 a), the analysis of the simple present as a quantification over times looks quite plausible. However, consider the features that such a quantifier would have to have. How big must the set $[[\text{verb}(a,t)]]$ be in order for “ $Qt(\text{verb}(a,t))$ ” to be true? It is immediately clear that there is no fixed number of times or percentage of times or similar such thing that would work. (2.62 a) implies that “Lola is running” is true more often than “Lola runs ultra-marathons” implies that “Lola is running an ultra-marathon” is true, and thus no fixed threshold can be used to interpret “ Qt ”.

A more plausible guess is that in order for a simple-present sentence to be true, its present progressive counter-part must be true with a regularity appropriate to the activity denoted by the verb. If so, then the simple present represents a quantifier with roughly the same meaning as “regularly”. In that case, a simple present tense sentence would be true in case its present progressive form is regularly true, where what counts as *regular* would be a function of the meaning of the verb. A person who brushes their teeth regularly does so much more frequently than a person who runs marathons regularly. However, further examples show that this hypothesis is also untenable.

A first difficulty is that, for some activities, irregular occurrences are a sufficient basis for a simple present tense sentence to be true.

(2.64) Jill tortures baby animals.

(2.65) Roy skis the K-12.⁵¹

⁵¹Holland (1985).

For (2.64) and (2.65) to be true, Jill and Roy need only have done something once. (Perhaps, it also has to be possible that they will or could do the relevant activity again in the future.) Worse yet, for some simple present tense sentences, no activity needs to have taken place – it is enough for it to be possible.

(2.66) That car runs (the mechanic just got finished with it).

The truth of (2.66) depends on what will happen when the ignition key turns, not what has happened in the past (the car could be right off the assembly line). A similar but distinct use of the simple present tense is to express the fact that an individual has the right to do, or the role of doing, something.

(2.67) Vice president Joe Biden, Democrat, breaks ties in the senate.⁵²

Finally, there are also verbs that occur in the simple present that do not occur in the present progressive, but we can still apply the same logical form to them.

(2.68) (a) Leslie wants money.

(2.69) $Qt(\text{Want}(\text{Leslie}, \text{money}, t))$

To account for (2.62 a) and (2.64) - (2.68 a), we need a quantifier that can mean all of the following: that an activity is done regularly, that it was done at least once in the past and potentially again in the future, that the subject is currently capable of doing the activity, that the subject is currently sanctioned to do the activity, and that the subject (in some sense) does something at all times within the present horizon (or possibly just at the current time).

There is clearly some conventional relationship between the meaning of present progressive sentences and their simple present tense counterparts (when they have one). However, the evidence does not support the hypothesis that the syntactic relationship corresponds

⁵²http://www.nytimes.com/interactive/2014/upshot/elections-2014-make-your-own-senate-forecast.html?_r=0.

to a logical relationship. There is no formal logical derivation from the simple present to the present progressive (or vice versa) because the common denominator amongst the relationships between the pairs of examples we have looked at is so thin that the application of any rule that applies to all will be insufficient to derive much of the meaning of any one of the simple present sentences. If the simple present is a quantified version of the present progressive, then the nature of the quantifier is left completely open by the semantics of the tense system of English, and thus fails to explain the logical relationship between any given pair of sentences.

We may be convinced that the simple present is a quantified form of the present progressive. And further, we may be convinced that it is precisely *because* the simple present is quantificational that it typically expresses generalizations. However, we have not found the kind of robust logical pattern that could serve as evidence that sentences with different surface-level syntactic structures have shared underlying logical forms. It seems in fact that we have not progressed beyond our initial characterization of generic sentences as those that can intuitively be understood as generalizations.

A further difficulty for the hypothesis that generalizations have a common logical form is that present progressive sentences can, after all, be used to express generalizations. Sentence (2.70 a) passes all of the syntactic and grammaticality tests for being non-generic, but nonetheless has the very characteristics that the category of generic sentence, and specifically the generic operator, was hypothesized in order to explain.⁵³

(2.70) (a) The monarch is migrating.⁵⁴

(b) Monarchs are migrating.

⁵³Previously published examples of present-progressive generalizations did not have stage-level interpretations such that their predicates applied to the individuals generalized over (stage-level predicates) (Krifka et al. (1995)).

(a) The wolves are getting bigger as we go north.

This sentence does not describe an event in which individual wolves are participants, so the sentence is not considered generic in the relevant sense (it is not a stage-level predication).

⁵⁴http://www.yourhoustonnews.com/friendswood/news/students-at-bales-elementary-participate-in-symbolic-butterfly-migration/article_b4a6feae-c315-5373-bf1f-7f1a20504f98.html. This sentence is potentially a counter-example to the claim that generic sentences never represent an event (Krifka et al. (1995)).

- (c) This monarch is migrating.
- (d) When the monarch is migrating, it eats a lot.
- (e) There are monarchs migrating.⁵⁵

Sentence (2.70 a) can either express something about a single butterfly or about a kind of butterfly. So, we should expect “is migrating” to be a generic predicate. And in fact (2.70 c) is implied by (2.70 b), as it should be on this hypothesis. But (2.70 a) describes an event (on either interpretation). And it can be embedded into a when-conjunction. It also passes the original test proposed by Carlson (1977) for non-generic predicates (as exemplified by (2.70 e)). Thus having the characteristic that the generic operator is supposed to explain cannot be predicted by independent criteria. We are either left with no motivation for supposing there is a generic operator in sentences that pass the tests for being generic (the same phenomenon is possible without there being a generic operator beneath the surface), or else we are left with the theory that sentences contain the generic operator if and only if they have the characteristics that the theory seeks to explain with the generic operator. In other words, the theory is either unmotivated, or ad hoc.

Sentences (2.70 a) - (2.70 e) show that the independent criteria proposed as tests for genericity can conflict – there is no underlying logico-syntactic feature sentences must have in order to be generic. The temporal specificity of non-generic sentences cannot explain why

⁵⁵Another test for non-generic predicates is whether they fit grammatically in sentences like:

- (a) There are ravens intelligent.

The fact that intelligent doesn’t make sense in that frame is supposed to show that it is a generic predicate (Carlson (1977b); Kratzer (1995); Krifka et al. (1995)). However, many supposedly non-generic predicates make about as much sense.

- (b) There are ravens dusty.

The standard example that passes this test is “available”. It seems that this frame is somewhat idiomatic. Note the difference in naturalness between the following.

- (c) There are ravens clean.
- (d) There are dishes clean.

The meaning of (d) and the meaning of “There are dishes available” is very close, which might explain why it sounds much better than (c). It works much more easily with verbs than adjectives.

- (e) There are ravens squawking/shaving/migrating/growing.
- (f) There are ravens black/clean/tired/angry.

But adding "all over town" helps a lot.

This phenomenon might be illuminated by considering these sentences with respect to a question under discussion. As a response to “Do you have anything clean/loud/tired/angry?” (f) sounds fine.

they do not admit of interpretation as generalizations. Hence we can draw no conclusions concerning the underlying logical form of a sentence from the fact that it expresses a generalization. Still, by considering what it is about (2.70 a) that makes it defy the generic/non-generic predication hypothesis, we can get a better understanding of the phenomena the distinction was formulated to explain.

Examples like (2.70 a) are difficult to come by due to the relative rarity with which four characteristics co-occur. First, the sentence has to have an interpretation on which it makes a statement about a kind. Second, on the same interpretation of the sentence, it has to be construable as a generalization about the things that are members of that same kind. “The wild yak has black fur” is a generic sentence, when the subject is interpreted generically, because individual yaks have black fur. Sentences like “The smart phone is taking over the telecommunications market” are not generic because no individual phone is taking over the telecommunications market. Finally, in order for “ x is migrating” to count as a non-generic predicate, it has to be construed as denoting something temporary (like an event), and because of the present progressive tense, it has to be an event that is unfolding at the time of utterance.

So all together, a counter-example to the claim that present progressive sentences cannot be generalizations about the members of a kind needs to (i) attribute something to a kind, and (ii) the attribute must be constituted by its members simultaneously doing the same thing. The vast majority of present-progressive sentences do not meet these requirements because most kind attributes do not consist in simultaneous activity of kind members. Sentence (2.70 a) works because migration is a relatively extended seasonal activity, and, crucially, monarch butterflies only live in North America and thus they experience the same seasonal changes at roughly the same time.

Sentences that fit these requirements are not particularly common, but with the recipe in hand it is easy to come up with more.

(2.71) The arctic fox is growing its winter coat.

(2.72) The sugar maple is changing color.

(2.73) The American tax-payer is expecting economic turbulence.

Recall that the generic hypothesis says that there is something in the logical form of generic predicates that explains why sentences with definite description subjects can be interpreted as generalizations. But by considering what it takes to be a counter-example to the claim that present progressive sentences cannot be generic, we can see that the contrast between simple present and present progressive is actually explained simply by the *meaning* of the present progressive, in conjunction with facts about the world.

What actually allows for the interpretation of (2.70 a) and (2.71) - (2.73) as kind predications is the fact that we as language users can find plausible interpretations for them with their subjects referring to kinds. In this respect, these examples are no different than specific uses in that we consider potential referents on the basis of the plausibility of what the sentence would express if they were the referent.

2.3.3.3 Adjectival Predicates

Before concluding the discussion of generic sentences, let us consider how the insights gained from the analysis of present progressive and simple present verbal predicates apply to sentences with adjectival predicates. We will see that as with intransitive verb clauses, sentences with adjectival predicates can be interpreted as generalizations if and only if the attribute denoted by the predicate makes a kind a plausible referent for the subject phrase. Once again, there is no logico-syntactic feature that explains the patterns within the data.

Within the category of adjective phrases, there is no (surface) syntactical distinction that correlates with the generic/non-generic distinction.

(2.74) (a) The coral snake is green.

(b) Coral snakes are green.

(c) This coral snake is green.

(d) When this coral snake is green, it makes a buzzing noise.

(2.75) (a) The traffic light is green.

(b) Traffic lights are green.

(c) This traffic light is green.

(d) When this traffic light is green, it makes a buzzing noise.

Although they share the verb phrase “is green”, (2.74 a) is a generic predication while (2.75 a) is a non-generic predication. (2.74 a), unlike (2.75 a), is naturally interpreted as a generalization (about coral snakes). A second, correlated distinction is that unlike (2.75 d), (2.74 d) is difficult to interpret.

In the generics literature, examples like these are taken as evidence that a given surface-level predicate can have two underlying syntactic forms, one generic and one non-generic. But we are now in a position to see examples like these (which are numerous) as strong evidence *against* a syntactic distinction between generic and non-generic predicates. However, ruling out a syntactic distinction leaves the data unexplained. Therefore, to complete the argument let us now turn to consider what actually explains the distinction between (2.74 a) and (2.75 a).

It is clear that being a generic adjectival predicate has something to do with the relationship of the predicate to the subject. So, our first step must be to look for a relationship between predicates and subjects that determines whether the predicate is generic or non-generic. Here, the standard hypothesis is that generic predicates ascribe attributes that are in some sense essential.⁵⁶

(2.76) (a) Obama is tall.

(b) The president is tall.

(c) The (basketball) center is tall.

⁵⁶Kratzer (1995); Krifka et al. (1995).

(d) When Obama is tall, he hits his head going through doorways.

(e) When the president is tall, he hits his head going through doorways.

In (2.76 a) - (2.76 e), we can see that “is tall” is a generic predicate of “Obama” and “the center”, whereas it is a non-generic predicate of “the president”. One might conclude that the reason for this is that being tall is an essential property of being a center and being Obama, but an accidental property of being a president.⁵⁷ However, this cannot be quite right.

(2.77) Snow is white.

(2.78) This snow is white.

(2.79) When snow is white it is safe to eat.

If Obama is essentially tall, then snow is essentially white. And, as expected for a generic predicate, (2.78) seems to follow from (2.77). However, (2.79) makes perfect sense, which is what we would expect for a non-generic predicate.

From this example, we might merely infer that the “When NP is Adj, ...” test does not work. But the problem is deeper than that.

(2.80) The American tax-payer is broke.

Sentence (2.80) can (and is frequently used to) express a generalization about American taxpayers, but there is no sense in which it attributes an essential property. In fact, the recurrent popularity of the sentence stems from recurrent fears that the American taxpayer is being pushed from “his” natural and rightful position as a successful provider for his family – that is, essentially *not* broke.

⁵⁷I doubt being tall is an essential property of Obama in any sense that metaphysicians are apt to be interested in. It is, however, the term used in the generics literature (which generally admits that its use is vague). A more accurate term might be “characteristic trait (attribute or property)”. Obama is characteristically tall, and so is the center of a basketball team. Whereas the president (understood generically) is not characteristically tall. (Note that an alternative term for generic sentence is “characterizing sentence”.)

What was hypothesized to be a matter of essence vs. accident is instead two different things that generally, but do not always, coincide. First, sentences of the form “NP is AdjP” are interpreted as generalizations just in case it is plausible that the speaker is actually saying, of *a kind* (that can be conventionally referred to by NP), that it has a characteristic that can be conventionally represented by AdjP. Second, sentences of the form “When NP is AdjP, ...” are difficult to interpret when the referent of NP is presumed to always (within a relevant time frame or space of possibilities) have the attribute represented by AdjP. In most cases, the first condition is satisfied by attributions of characteristic traits, and such traits also satisfy the second condition. However, as we see with (2.80), traits that are not characteristic of a kind are sometimes attributed to them.

One thing that we can conclude from the fact that the various tests for generic/non-generic predicates do not always agree is that there is not actually a robust pattern in the phenomena, calling for an explanation at the level of logical form. If there were actually a distinction at the level of logical form, we should expect uniform behavior – no exceptions. The other thing that emerges from these adjectival-predicate examples is that whether a sentence’s subject gets a non-specific interpretation is, once again, a matter of a combination of (i) world knowledge and (ii) there being a plausible reading of the entire sentence with the subject referring to a kind.

2.3.4 Kind Attributes and Properties of Sets

Stepping back from the details of the theory of generic sentences, we can take a fresh, wide-angle look at why variable-binding analyses of generalizations expressed by sentences with singular definite subjects are bound to fail. The basic answer is twofold: (i) kinds have traits that are not inherited from their members, and (ii) whether a trait of some members of a kind is also a trait of the kind depends on a variety of factors.

(2.81) The tyrannosaurus rex is well-known to school children.

- (2.82) The whale developed into a sea creature over millions of years.
- (2.83) The domesticated pig arrived in America soon after Columbus.
- (2.84) The mosquito carries malaria.
- (2.85) The common loon lays two eggs a year.
- (2.86) The sequoia is a huge tree that only grows above 9000 feet.
- (2.87) (a) The sea turtle hatchling climbs out of its egg and crawls straight into the sea and begins to swim.
- (b) The sea turtle hatchling is not likely to make it all the way out of its egg before being eaten.

Absent prior theoretical commitments, what sentences (2.81)-(2.87 b) clearly illustrate is that there are a variety of ways that kind characteristics can be related to the characteristics of kind members. People can relate to kinds without directly relating to their members (2.81). Kinds can change over time in ways that none of their members do (2.82). A characteristic of a single member of a kind can be inherited by the kind (2.83). Kinds can have characteristics in virtue of the practical or psychological significance of the characteristics of a minority of their members (2.84); because a sub-kind (e.g., the female loon) has a characteristic (2.85); on the basis of what it would be natural for its members to be like (2.86); on the basis of the natural developmental trajectory of its members; or, finally, on the basis of raw statistics. (Note that (2.87 a) and (2.87 b) are not contradictory.) No uniform quantificational analysis can model this range of possibilities.

Aside from the kind-reference view for which I am arguing, another response to the failure of a quantifier-based explanation is to abandon quantifiers in favor of some other kind of variable-binding operator.⁵⁸ But the problem with this approach is that insofar as it avoids reference to kinds, it doesn't work (because of kind attributes that do not distribute

⁵⁸Leslie (2007).

to members of the kind), and insofar as it relies on reference to kinds, the variable-binding operator is pointless (because kind reference alone can account for the data). Aside from the problem created by the fact that not all kind characteristics are inherited from or shared with kind members, the relationship that does obtain between kind characteristics and member characteristics depends on the *kind* of kind involved. Living kinds have natural habitats and developmental trajectories that mediate their inheritance of member characteristics, while artefactual kinds have functions that mediate the inheritance of kind characteristics by members. Facts like these would have to be built into the semantics of a variable-binding operator, but there is no reason to think that these distinctions are specifically grammatical in nature.

Once it is recognized that the subjects of sentences such as (2.81) - (2.87 b) are referential, and it is recognized that definite descriptions can refer to kinds, then the need for a complicated and flexible type of variable-binding operator is obviated. In addition, a kind-reference analysis explains the reciprocal nature of kind and member characteristics. That the common loon mates for life has implications for each and every individual loon because they are all members of the kind, the species *Gavia Immer*. And so facts about the kind are facts about its members as well. And conversely, that a member of a kind has a trait is a fact about the kind. What it takes for particular traits of individuals to be sufficient for a kind to have a corresponding trait is simply not within the ken of semantics.

Recognizing the use of definite descriptions to refer to kinds sheds new light on the relationship between the representation of facts about individuals, and generalizations that pertain to all members of a class of individuals but are not statistical facts abstracted from individual facts. It is only once this new perspective is reached that a unified semantics for all of the different uses of definite descriptions becomes possible.⁵⁹

⁵⁹I received a great deal of help with the material in this Chapter from Joseph Almog, Andrea Bianchi, Sam Cumming, Kristina Gehrman, David Kaplan, and audiences at the Third Parma Workshop on Semantics and Pragmatics, at Kaplan Fest, and the SynSem Symposium at UCLA.

Chapter 3

Direct Reference in Indirect Discourse

3.1 Introduction

In this chapter I will extend the account of the semantics of singular definite noun phrases to include those within subordinate clauses. In keeping with recent literature, I describe the subject of this chapter as indirect discourse, but I include any sentence of the form “NP V (that) S”, whether or not “V” is a verb of communication. The basic semantic claim of the theory developed in Chapter 1 and Chapter 2 is that singular definite noun phrases refer to single entities, which include both individuals and kinds. What I will argue here is that indirect discourse semantics is not exceptional with regard to the semantics of noun phrases. Noun phrases work, semantically and communicationally, within embedded clauses as they work in un-embedded clauses.

On the pure-reference view, indirect discourse is not a special or problematic subject for semantics, but rather is one amongst the many sentential operator constructions, like negation and modal operators, that add complexity to linguistic representations but do not change the semantics of the constituents with which they combine. The pure-reference

analysis of indirect discourse utilizes the same referential semantics it applies to expressions when they are not embedded in a subordinate clause.

The basic claims of the first two chapters are: (i) singular definite noun phrases contribute only reference to the semantic content of complete utterances; (ii) narrow-scope, generic and predicational uses of singular definite noun phrases refer to kinds (as opposed to the members of those kinds that are referred to by specific uses of the same phrases); (iii) the choice of expression to refer to a given entity is governed by three factors, (a) communication of the identity of the referent, (b) communication of the point of the utterance, and (c) other communicative goals like being polite, funny or ironic. Thus evidence against the strictly semantic aspect of the theory of the last two chapters would be examples of utterances in which substituting co-referential, singular definite noun phrases would change the semantic value of the utterance. Evidence against the broader theory of communication in which the semantic theory is presented would be examples of utterances in which the choice of singular definite noun phrase is answerable to criteria other than (a) - (c).

Section 3.2 is an outline of an account of the semantics of indirect discourse which will serve as the basis for the analysis of embedded noun phrases. In Section 3.2, I argue that embedded noun phrases perform the same communicative functions as un-embedded noun phrases. In Section 3.4, I will consider a variety of examples which have been taken as evidence that substitution of co-referential noun phrases can change the truth conditions of an utterance. I will show that ambiguities, which occur also in un-embedded clauses, are at fault for the failures of substitution – in other words, substituted expressions that change the semantic value of an utterance turn out not to be actually co-referential. Finally, in Section 3.5, I will consider an argument that failures of symmetry in embedded copular clauses are counter-examples to purely referential accounts of names. I will argue that “is” does not mean “is identical with”, and so we cannot infer from the symmetry of identity that identificational clauses are semantically symmetric, and thus that failures of symmetry are consistent with a purely referential account of singular definite noun phrases.

3.2 Basic Semantics of Indirect Discourse

Pure-reference indirect discourse semantics is an application of the pure-reference theory of singular definite noun phrases to the subordinate clauses of indirect discourse. On the pure-reference theory the syntax of verbal subordination has no direct influence on the semantics of embedded noun phrases. The syntax of a sentence determines how the things signified by the constituents of the sentence are represented as relating to each other. A sentence like “Jane runs” represents Jane as doing the action of running. Embedding a simple sentence in indirect discourse does not change what the constituents signify. What changes is the relationship that they are represented as having to each other. “Mary thinks that Jane runs” represents Mary as being in the state of thinking that Jane does the action of running. Applying the pure-reference theory to indirect discourse merely involves taking into account the additional complexity that comes with having one clause embedded in another: embedding does not fundamentally change the semantic analysis of the embedded clause, and it does not involve introducing any new semantic levels or types, or revision of the basic theory. Consider the following example.

(3.1) Jill thinks that Mont Blanc is stunning.

We can describe the meaning of (3.1) from inside out, by looking first at the meaning of the embedded clause, which represents Mont Blanc as having the property of being stunning. Being part of the complement of a verb does not change the meaning of the embedded expressions. As a whole, (3.1) represents Jill as being related in a particular way to what is represented by the embedded clause, Mont Blanc’s having the property of being stunning. The complete sentence represents Jill as being in the thinking relationship with Mont Blanc being in the being relationship with the property stunning. It may be helpful to compare and contrast indirect discourse with a transitive verb construction.

(3.2) Jill kicked Mont Blanc.

Sentence (3.2) represents Jill as having done the action of kicking to Mont Blanc. From the point of view of semantics, the most interesting difference between the two sentences is that whereas (3.2) relates Jill to an object, (3.1) relates Jill to a complex of an object, a property and a relationship between them. In both cases, the subject is related via the denotation of the verb to the semantic content of the complement of the verb. In the case of indirect discourse, the meaning of the complement of the verb is semantically more complex than the direct object of a transitive verb.¹

3.3 Point of View in Embedded Clauses

A widely held view about indirect discourse is that embedded noun phrases systematically (though not necessarily always) represent what they stand for from the point of view of the subject of the main clause. Contrary to this standard view, we will see that the factors that influence the interpretation of noun phrases in simple sentences and subordinate clauses are the same.

The primary communicative role of embedded definite noun phrases is the same as that of un-embedded definite noun phrases – communicating the identity of the referent to the audience. The following four sentences can be used to describe a single state of affairs. Which one would be used depends on the speaker’s beliefs about how familiar the participants in the conversation are with Lady Gaga, and is independent of which of the four expressions the subject, Lena, would use to express the thought being reported.

(3.3) (a) Lena thinks a major celebrity called her this morning.

(b) Lena thinks the singer of “Poker Face” called her this morning.

(c) Lena thinks the most popular person on Facebook called her this morning.

¹This way of looking at the matter contrasts with standard semantic theories in which the meaning of an expression is a function of the meaning of its constituents – what goes in may not come out. The theory being described here holds that what goes in stays in. The semantic content of an utterance is built up from semantic parts in the way a sentence is built up from parts – syntactic complexity is mirrored in complexity of content (as opposed to complexity in the algorithm for calculating a possible unstructured content).

(d) Lena thinks Lady Gaga called her this morning.

Suppose that Lena is big fan of Lady Gaga, and that friends of hers tricked her this morning with a prank phone call. Lena's father asks her older sister Selma what all the commotion is about. She is going to respond with one of (3.3 a)-(3.3 d). Which of the four she chooses will depend on her expectations as to her father's familiarity with Lady Gaga. If he has no idea who she is, she will respond with (3.3 a). If he knows the song "Poker Face", but does not't know the name of the singer, then she will likely respond with (3.3 b). If he knows about Facebook, but perhaps not about the song "Poker Face", she might respond with (3.3 c). (If he knows about both, either sentence would be appropriate.) If he has not't been hiding under a rock, and he knows who Lady Gaga is, then most likely Selma will use her name, as in (3.3 d).

The difference in the appropriateness of (3.3 a)-(3.3 d) is a function of what it will take for the audience to catch on to what the speaker wants to convey, and has nothing to with the fact that they are indirect discourse reports. To see this, we can consider how Lena might report the cause of her excitement to her father. She could use the subordinate clauses of (3.3 a)-(3.3 d) (with "me" substituted for "her") depending on the very same factors that determines which of (3.3 a)-(3.3 d) her sister would use: the familiarity of the audience (her dad) with Lady Gaga.

Each of (3.3 a)-(3.3 d) can be used to describe the same state of affairs. Which one will be used is a matter of appropriateness to the task of informing a particular person with a particular level of knowledge about an individual involved in the situation being described. For our purposes, the main thing to note is that this phenomenon is independent of whether an expression is embedded in indirect discourse or not.

3.3.0.1 Whose Beliefs are Indicated by Embedded Expressions?

There are two different ways that simple assertoric assertions convey information about the point of view of the speaker. First, a sincere assertoric utterance is an explicit expression of

a belief of the speaker. If you say, “The hat is expensive”, that is clear prima facie evidence that you believe the referent of “the hat” is expensive. The words a speaker chooses to convey a message are a second source of information about how things seem from the speaker’s point of view. Insofar as a speaker is using language conventionally and sincerely, their use of a referring expression indicates that they believe that the referent satisfies the conventional constraints on the use of that expression. If you say, “The hat is very expensive”, people are apt to assume that you think the thing you described as expensive is a hat.

Standard theories of indirect discourse hold that when a referring expression is in a subordinate clause, it reflects the point of view of the subject of the matrix clause, rather than the point of view of the speaker.² Is indirect discourse a special context in which the use of a particular referring expression does not reflect the beliefs of the speaker, but perhaps those of the subject of the report instead? As we shall see, indirect discourse does not change the perspective we take descriptions to come from, and thus expressions used in indirect discourse reflect the beliefs of the speaker in the same way as in un-embedded uses. The use of an un-embedded descriptive noun phrase generally licenses the inference that the speaker believes the description to accurately describe the referent.

(3.4) The man over there is handsome.

When there is no indication to the contrary, a person who utters (3.4) is assumed to think that the person she is saying is handsome is a man.

(3.5) Lena thinks the man over there is handsome.

Just as the utterance of (3.4) gives evidence that the speaker believes the person described as handsome is a man, the speaker of (3.5) will be assumed to think that the person she is saying Lena thinks is handsome is a man. This observation is the key to understanding certain examples that are difficult cases for content theories of indirect discourse, but that can serve as poster children for the pure-reference theory.

²See for example Heim (1992).

In the scenario with Lena and Selma, the speaker and the subject of the matrix clause believed the same things about Lady Gaga. Thus any of the expressions Selma (the speaker) might have used to refer to Lady Gaga might also have been used by Lena (the subject of the matrix clause). We will now consider what happens in cases of conflict between the beliefs of the speaker and the subject.³

(3.6) (a) Jill thinks the lying bastard is honest.

(b) The lying bastard is honest.

A speaker who utters (3.6 b) says something that seems contradictory. Thus, if the embedded clause of (3.6 a) reflects the point of view of Jill in the way that (3.6 b) reflects the point of view of the speaker, then (3.6 a) should attribute a contradictory thought to Jill. What (3.6 a) plainly illustrates is that when there is a conflict between (i) the semantic content of an embedded clause and (ii) the conventional constraints on the use of a referring expression in that clause, the conflict is not interpreted as a contradiction within the beliefs of the subject of the matrix clause.⁴

This example poses no difficulty on the assumption that the role of the description “the lying bastard” is to help the audience identify its referent. In both (3.6 a) and (3.6 b), “the lying bastard” is used by the speaker of the sentence to identify a particular person. The speaker chooses a phrase that will let the audience know who she is talking about (she may of course have additional reasons for choosing a particular expression). The primary communicative role of “the lying bastard” is not to communicate that the referent is a lying bastard, but to let the audience know who is being referred to.

Once “the lying bastard” is recognized as a tool for securing reference identification, rather than an ingredient in the content of the sentence, the reason (3.6 b) is contradictory

³See Kaplan (1989).

⁴The standard strategy for dealing with such cases within a content theory, due originally to Kaplan (1968), is to existentially generalize over singular terms that designate the same thing as the term within the subordinate clause. Thus the analysis of (3.6 a) would be:

(i) $\exists x$ (the referent of x = the referent of “the lying bastard” & x = “ y ” & Jill thinks that $[[y$ is honest]])

(i') There is an x such that its referent is the referent of “the lying bastard” and it is “ y ” and Jill is in the thinking relation with the content of “ y is honest”.

while (3.6 a) is not comes into focus. (3.6 b) seems contradictory because the speaker uses the description “the lying bastard” to refer to someone, presumably because it is accepted by the participants in the conversation (the speaker included) that the person being referred to is a lying bastard, whom she then goes on to say is honest.⁵

In (3.6 a) the speaker uses “the lying bastard” to refer to someone whom she says that Jill thinks is honest. The would-be contradiction in (3.6 a) is not within Jill’s thought, but between the way the speaker represents a certain person and the way the speaker represents Jill as thinking of that person. According to the pure-reference theory, the embedded clause represents a certain person as being honest, and the whole sentence represents Jill as standing in the thinking relation to what the embedded clause signifies.

3.3.0.2 Don’t Change the Point

I have so far argued that it is not part of the conventional semantic meaning of indirect discourse that embedded definite noun phrases reflect the point of view or representational competence of the subject of the report. However, that does not imply that the choice of embedded noun phrases is never influenced by the speaker’s intent to convey something about the subject’s point of view. In addition to their primary role as referent identifiers, definite noun phrases, both embedded and un-embedded, can serve the secondary purpose of helping to communicate the point that the speaker is trying to convey by making the utterance. In some cases conveying the point of an utterance can involve indicating something about the point of view of someone referred to in that utterance.

(3.7) (a) Othello believes Iago is honest.

(b) Othello believes the man who *seems* honest *is* honest.

⁵Note that according to the pure-reference view, (3.6 b) is not semantically contradictory. Whether it comes across to the audience as a contradictory utterance depends on what they attribute the speaker’s choice of expression to. Generally, the use of an expression to refer to an individual indicates that the speaker believes the individual meets the conventional requirements for being a referent of that expression, but not always. An utterance of (3.6 b) will not come across as contradictory as long as the audience has an explanation for the choice of “the lying bastard” that is consistent with its referent being believed by the speaker to be honest. There is however a sharp contrast between (3.6 a) and (3.6 b) – no such special circumstances are required for the felicitous use of (3.6 a).

(c) Othello believes the flatterer is honest.

All three of (3.7 a) - (3.7 c) can be used to express the same fact, but they are not equivalent in terms of the point they are suited to make about Othello. A person offering evidence that Othello is naive might use (3.7 b) as an illustration of his character, whereas (3.7 c) would more likely be used to illustrate Othello's vanity. Whether (3.7 a) would (be suited to) convey anything particular about Othello's character depends on the wider conversational situation in which they are uttered.⁶

3.3.0.3 Irony

Another way that the use of a given expression can indicate something about the beliefs or point of view of someone other than the speaker is through irony and other mimetic forms of representation.

(3.8) Santorum thinks the communist in the oval office is destroying America.

It is certainly possible to imagine hearing a liberal Democrat saying (3.8), without taking her utterance as expressing or implying that she believes that the president is a communist (or in some relevant way like a communist). We can hear (3.8) as saying that Santorum thinks that the president, who he (Santorum) thinks of as being a communist, is destroying America. However, when we interpret (3.8) in such a way that it does not indicate that the speaker believes Obama is a communist, we are inferring that the speaker has an ulterior motive for her choice of expression. When we infer that the speaker is using a particular referring expression to make a comment on a person who would use that expression, we

⁶See Hitchcock (1944) for a very good example of a speaker "changing the point" by using a different phrase to refer to the same individual:

Kovac: Do you mean you want to turn the boat over to the man who sunk our ship and shelled our lifeboats?

Connie: I mean I want you to turn the boat over to the man obviously best qualified to run it.

There is no disagreement amongst the passengers of the lifeboat that the same man fits both the description used by Kovac and the one used by Connie.

This is a common literary technique not limited to indirect discourse. See Nichols (2014a); Booth (1974); Cohn (1984); Wood (2008).

hear the use of the expression as ironic. It is irony, not indirect discourse, that occasions the shift in whose point of view we interpret the description as reflecting. The same effect can be achieved with an un-embedded use of the same clause.

(3.9) I was watching Fox News and guess what I learned. The communist in the oval office is destroying America.

In (3.9), just as in (3.8), the use of “the communist in the oval office” indicates that someone, or some people, think, or say, that Obama is a communist. From the preceding sentence, who those people are is clear. Irony involves saying something as if from another point of view in order to make a comment on that point of view: “this is the sort of thing that these kind of people are apt to say”. It is always possible for a speaker to use irony to shift the point of view from which an utterance (or part of one) is interpreted, and so it does not tell us anything specifically about indirect discourse that expressions in an indirect discourse report can be used ironically. Even with expressions in subordinate clauses, the choices of point of view are not limited to the speaker and people referred to in the utterance. With a little imagination it is easy to think of scenarios where embedded expressions reflect the point of view of a third party.

3.4 Substitution Failures

We now turn to purported cases of co-referential substitution changing the truth conditions of an utterance.

3.4.1 Ambiguity

The pure-reference view implies that the only property of a singular definite noun phrase relevant to the truth or falsity of an assertoric utterance is the identity of its referent. Thus the theory predicts that substitution of co-referential definite noun phrases should not

change the truth value of an utterance. Some care is needed in assessing potential counter-examples to this hypothesis. Definite noun phrases can refer to different things on different occasions, and two potentially co-referential expressions can be interpreted as referring to different things.

(3.10) (a) Bill would like to be the president.

(b) Bill would like to be Obama.

(3.11) (a) Bill dreamt that he was the president.

(b) Bill dreamt that he was Obama.

The pairs (3.10 a) and (3.10 b), and (3.11 a) and (3.11 b) can differ in truth value despite the fact that Obama is the president. But this has nothing in particular to do with indirect discourse, because the embedded clauses of (3.10 a) and (3.11 a), “He is/was the president” are ambiguous.⁷

According to the pure-reference theory, “the president” can refer to the kind *president of the USA*, or to the sole member of that kind (Obama). But, however the ambiguity in the subordinate clauses is explained, it is already there before embedding.⁸ Ambiguity can also obscure the point of view that is reflected in the choice of a given expression.

(3.12) (a) Bill thinks the math teacher is a/the spy. #In fact he doesn’t believe there is a spy.

(b) Bill thinks the spy is a/the math teacher. In fact he doesn’t believe there is a spy.

In (3.12 a) and (3.12 b), we see that there is a contrast between what the embedded subject and predicate nominal indicate about the subject of the main clause’s (Bill’s) point of view. One might conclude that the embedded predicate nominal refers from the point of view of the main-clause subject – that Bill would use that same expression to refer to the

⁷I am assuming that the complement of “would like” is a full clause with an elided subject. Compare, “Bill would like Linda to be the president.”

⁸See Nichols (2014b).

same individual. However that hypothesis would only be supported by these examples if embedded predicate nominals always referred to an individual that fit the descriptive information conveyed by the expression. But, according to the kind-reference theory, these predicate nominals (which are predicational on the relevant interpretation) refer to kinds, not individuals. Thus, the contradictory nature of (3.12 a) is explained without reference to how things would be designated from Bill's point of view.

(3.13) Bill thinks the math teacher is a member of the kind *spy*. In fact he doesn't know that anything is a member of the kind *spy*.

According to the kind-reference theory, (3.13) is an accurate gloss of (3.12 a). It makes transparent why (3.12 a) seems contradictory, and that it does not depend on how Bill would refer to anything. This explanation of the difference between (3.12 a) and (3.12 b) predicts that we should be able to replace predicate nominals with others that refer to the same kind without manifesting the contrast between (3.12 a) and (3.12 b).

(3.14) (a) Bill thinks the robber is a/the thespian. #In fact he doesn't believe there is a dramatist.

(b) Bill thinks the robber is a/the dramatist. #In fact he doesn't believe there is a thespian.

(c) Bill thinks the robber is a/the thespian. But he doesn't know the word "thespian".

Both (3.14 a) and (3.14 b) manifest the same kind of contradiction as (3.12 a). However, the first sentences of (3.14 a) and (3.14 b) are consistent with Bill not even knowing the meaning of the predicated nominal. Thus, the predicate nominal does not reflect Bill's point of view or how he would refer to things represented in the embedded clauses of indirect discourse reports about him.

3.4.2 Identification and Definition

In the last section, we looked at failures of substitution due to kind/individual-reference ambiguity in predicate nominals. In this section we will look at substitution failures due to use/mention ambiguity.

Consider the following situation. Anne never really liked her name. When she went away for college she decided to go by her middle name, “Theresa”. Anne’s cousin Darlene has a hard time remembering that Anne wants to be called “Theresa”. Darlene is talking to some of Anne’s friends, and one of them begins looking very confused – he has no idea who the conversation is about, though he can tell that Darlene expects him to. Someone decides to clear things up by saying:

(3.15) David doesn’t know that Anne is Theresa.

The issue this sentence raises for the pure-reference theory is that it seems like semantic analysis of this sentence would assert that David fails to know something trivial – though intuitively it does not. If the embedded sentence means that Anne is identical with Theresa, then because Anne is actually identical with Theresa, it means that Anne is identical with Anne. In that case, the whole sentence denies that David knows that Anne is Anne (which he does know despite his confusion).

A pre-theoretical explanation of what (3.15) means is that David doesn’t know that “Anne” refers to (or is a name of) Theresa. He doesn’t have two people in mind, nor does he think that he does. He simply doesn’t know who Darlene is talking about, and the reason why is that he doesn’t know that Theresa is also called “Anne”. In fact the use of identificational copular constructions to convey information about an expression is common.

(3.16) Maybe she didn’t hear you correctly and thought you said lesbian? Or maybe she doesn’t know what a thespian is and assumed you meant lesbian.⁹

⁹Found on YahooAnswers.com as an answer to, “I asked my friend if she and her mother were thespians. She slapped me! Is she being a drama queen?”

In this example it is clear that “doesn’t know what a thespian is” is used to express the state of affairs of not knowing what the word “thespian” means. A thespian is an actor, and the author of (3.16) is not suggesting that the person in question does not know what an actor is. Consider also the following sentences from the directions to a set of exercises on a standardized test of memory.¹⁰

(3.17) A baloo is a bear.

(3.18) Sculch is junk.

(3.19) Wuzzle means to mix.

(3.20) Alate means to have wings.

Here we have “is” used to form nominal definitions in simple sentences. Interestingly, the writers of the test use “is” for definitions of nominals and “means” for definitions of verbs. The systematic use of “is” instead of “means” for defining nouns is evidence that identificational sentences are conventional ways to define nouns, and thus that, when the context makes the use clear, noun phrases can conventionally be used to refer to themselves. Replacing “a baloo” in (3.17) with “a bruin” would not change the truth value of the sentence but would certainly change its semantic content. But that is no indication that the two expressions contribute different semantic content when they are used to refer to bears.

The sentence, “Anne is Theresa” can be used conventionally with the literal meaning that “Anne” refers to Theresa. Thus the fact that (3.15) can be used, conventionally, to literally mean that David does not know that “Anne” refers to Theresa is exactly what we should expect if embedding does not change the semantics of embedded expressions. Also, that names can sometimes refer to themselves is consistent with the theory that definite noun phrases contribute only reference to complete utterances. In some cases, noun phrases refer to themselves. Sentences like (3.17) and (3.18), from widely used educational material,

¹⁰The California Achievement Test.

are evidence that using a noun phrase to refer to itself is conventional, and even considered to be conventional by the guardians of American English.

3.4.3 Alter Egos

The holy trinity has been making trouble for referential semantic theories since the early middle ages. It seems to follow from the doctrine of the holy trinity that different descriptions of one entity (the triune God) are true depending on whether that entity is referred to as “Jesus Christ”, “The Holy Ghost”, or “God the Father”. The holy trinity is a special case, and perhaps it can be dismissed as an issue for semantics on the grounds that the joint supposition of three persons and only one God are logically inconsistent.¹¹ Personally, I doubt that the impossibility of the holy trinity really makes it irrelevant to semantic theory, but in any case, the very same issues are now raised with staple examples in semantics and the philosophy of language involving much less improbable characters.

Although the metaphysics is different, super heroes with secret identities and authors with pen names pose the same challenge for direct reference today as the holy trinity did for Augustine and Ockham.

(3.21) Lex Luthor believes that Superman can fly but that Clark Kent cannot.

Sentences like (3.21) have been used as evidence of the fact that co-referential names cannot always be substituted without changing the semantic value of an utterance. The relevance of such sentences to the semantics of names and indirect discourse depends of course on whether “Superman” and “Clark Kent” are, in fact, co-referential. The evidence in support of co-referentiality is that “Superman is Clark Kent” is true. If Superman and Clark Kent were identical, then “Superman” and “Clark Kent” would be co-referential. But just as in the case of the holy trinity, there is some difficult-to-swallow metaphysics underlying the supposed semantic issue.

¹¹See Geach (1967) for a logical theory of identity consistent with the doctrine of the holy trinity.

Sentences like (3.21) are only semantically interesting on the supposition that Superman and Clark Kent are identical. However, though there is some of the metaphysical slipperiness of the holy trinity involved here, there is relatively clear evidence that “Superman” and “Clark Kent” name distinct entities. Superman may be the same person as Clark Kent, but when alter-egos are involved, that does not mean the two are identical. Brief research into superhero comics reveals that a superhero’s secret identity is also an alter-ego. In other words, Clark Kent and Superman are not identical. One of them is the alter-ego of the other. (Which one is which has apparently changed over the years. In early comics, Superman used the Clark Kent identity to hide his powers. Later it was Clark Kent who used the Superman identity to protect those close to him from the danger of being associated with a super hero).¹²

It is true that Superman is Clark Kent, but it does not follow that they are identical. The inference from “Superman is Clark Kent” to “‘Superman’ is co-referential with ‘Clark Kent’” depends on the interpretation of “is” as “is identical with”. As a general matter, that is rarely a plausible gloss for “is”, and the fact that it does not work in these cases should not be so surprising. (See also Section 3.5 below for further discussion of “is” and the relation of identity.) And because they are not identical, “Superman” and “Clark Kent” are not co-referential. Thus superheroes with secret identities do not reveal hidden failings in direct-reference semantics theories.¹³

Sometimes a person with two names has two identities, and sometimes they just use more than one name. In some cases it is not easy to tell which is the case.

(3.22) Students who do not know that Johannes de Silentio is Kierkegaard will fail the quiz.

(3.23) Ernie was surprised to learn that Cassius Clay is Muhammad Ali.

¹²See Wikipedia (2014a).

¹³See Saul (1997). In her examples, e.g., “Clark went in and Superman came out” there is essential dependence on the fact that Superman and Clark are two personae, even if they are one man, or one super hero.

Pen names and names changed in association with a personal transformation are somewhat like names associated with alter egos and somewhat like co-referential names. How to tell one from the other is not a matter of semantics. It is enough to point out the slippery possibilities that will go along with marginal cases.

3.5 Identificational Predications and the Relation of Identity

In the last section, we looked at apparent failures of co-referential substitution. The potential problem for the pure-reference view was that if co-referential terms could not be substituted inside indirect discourse, then the meaning of the embedded expressions would seem to go beyond reference. What we saw was that names and descriptions that could not be substituted turned out not to actually be co-referential. The substitution failures were evidence not of an intensional context but of an ambiguity.

Embedded identity sentences raise a different issue having to do with the positions in which the expressions occur in the embedded clause. Identity is a symmetric relation, but reversing the order of the terms of an embedded identity sentence can change the meaning of an indirect discourse report. As such, we need to see whether the embedding is the cause of the asymmetry, and, if it is, whether it is evidence that non-referential semantic content is part of the meaning of indirect discourse reports.¹⁴

Samuel Cumming (2008) developed an example with a proper name as a nominal predicate specifically to show that names in indirect discourse can have both *de re* and *de dicto* interpretations. The example is supposed to show that a name embedded in indirect discourse can represent its referent from the point of view of the subject of the report.

In *Love's Labour's Lost*, Katherine and Rosaline trick their lovers by tipping them off

¹⁴Note that the asymmetry exists with definite descriptions too, and un-embedded at that. The explanation of the difference in informativity of “My neighbor is the one on the left” compared to “The one on the left is my neighbor” should mesh with indirect discourse semantics. The semantics of copular clauses is a rich topic unto itself which cannot be done justice here.

as to who will be wearing what at a masked ball, and then switching. The trick works. Biron dances with a woman wearing what Rosaline promised that she would be wearing, and he thinks he is dancing with Rosaline. In fact the woman with whom he is dancing is Katherine. Observing the situation, Maria remarks:

(3.24) Biron thinks Katherine is Rosaline.

Maria's remark correctly describes the situation. On the other hand, if she had reversed the names so that "Rosaline" was the subject and "Katherine" was in the predicate, as in (3.25), her remark would have been false.

(3.25) Biron thinks Rosaline is Katherine.

We have in (3.24) and (3.25) two embedded identity sentences in which the same names appear, but in one case in one order, and in the other case in the opposite order. The two names are not co-referential so this not a putative failure of co-referential substitution. Rather the idea is that the order of the names should not change the referential-level meaning of the embedded clause as a whole, and so, if only the referential-level of meaning mattered to the meaning of the indirect discourse reports, they should have the same meaning. Identity is a symmetric relation, and the embedded clauses are identificational, so, the story goes, if it were only the entities referred to by the names that played a role in the semantic content of the sentences, they ought to have the same truth conditions. Since they don't have the same truth conditions, then at least one of the names is not purely referential.

Let us go through the case now a little more slowly and carefully to see why the example is supposed to make trouble. It is generally believed that sentences of the form "NP is NP" can either be predicative or identificational. For instance, "Obama is the president" can either predicate being the president of Obama, or, in an appropriate conversational circumstance, it can express the fact that Obama is the same person as the president. Names do not generally work predicationally, like "the president" can, and they certainly don't seem to have predicational meanings in the sentences we are interested in here. The

embedded sentences in (3.24) and (3.25) do not seem to be of the predicational sort, so “Katherine is Rosaline” seems to mean, at least roughly, that Katherine is identical to Rosaline.

Identity is a symmetric relation: for any x and any y , if x is identical to y , then y is also identical to x . Thus if Katherine is identical to Rosaline, then Rosaline is also identical to Katherine. To say that a relation is symmetric is just to say that if it holds in one direction it holds in the other. But in the case of identity, holding in one direction seems to be just the same thing as holding in the other. Beyond the relationship of mutual implication, it is hard to see what the difference is between Katherine being identical to Rosaline, and Rosaline being identical to Katherine. After all, if Katherine and Rosaline are identical, their being identical is only one thing standing in a relationship to itself.

Coming back to identificational sentences, suppose rather than attributing anything to Biron, Maria were to assert one of the following:

(3.26) (a) Katherine is Rosaline.

(b) Rosaline is Katherine.

If the two sentences assert the identity of Katherine and Rosaline, then, it seems, they describe the same situation. It makes no difference to an equation which term is to the left of the identity sign and which is to the right. If Katherine is identical to Rosaline, then “Katherine” and “Rosaline” refer to the same person anyway, so how could it make a difference in which order the two names appear? So, (3.26 a) and (3.26 b) appear to describe the same situation.

The difference between (3.24) and (3.25) is just that one embeds (3.26 a) and the other embeds (3.26 b). And (3.26 a) and (3.26 b) seem to describe the very same situation. Thus, it seems that the influence of (3.26 a) and (3.26 b) on (3.24) and (3.25) is not limited to the situation they describe: they describe the same situation, but have different semantic contributions. If so, then their semantic contributions go beyond the referential-level of meaning.

On the assumption that there is a de re/de dicto distinction, the difference in meaning between (3.24) and (3.25) can be explained as follows. If, in both (3.24) and (3.25), the subject of the embedded clause is de re, while the predicate is de dicto, then the two statements involve different ways of the subject thinking about someone: one statement involves one guise and the other a different guise. The two then have different semantic ingredients, and thus will clearly describe two different situations (only one of which matches the plot of *Love's Labour's Lost*).

According to the de re/de dicto explanation, (3.24) attributes an identity thought to Biron. The referent of the first term of Biron's identity thought is Katherine (who has been referred to by Maria's use of "Katherine"). The second term of the identity is the name "Rosaline", which refers to Rosaline, but in contrast to "Katherine", Biron is being described as using the name (it is supposed to be his "dicta"), and its use is supposed to capture Biron's perspective on the situation (as opposed to the speaker, Maria, using the name to refer to the object of Biron's thinking).¹⁵ On the other hand, in (3.25), "Rosaline" is used by Maria to refer to Rosaline, and "Katherine" is attributed to Biron, and so it expresses part of the content of Biron's thought. The fact that one of the names is being used referentially, while the other expresses part of the content of Biron's thought, is what breaks the symmetry between the two embedded identity sentences, and therefore explains the difference in meaning between the two indirect discourse reports.

The challenge raised by Cumming's example for the pure-reference theory is that it cannot rely on a de re/de dicto ambiguity in the names to break the symmetry. The pure-reference theory cannot say that the predicate nominal signifies Biron's thought. Rather, it appears stuck accepting that both names are referential. And indeed, I think the difference between (3.24) and (3.25) can be explained while accepting that both embedded names are used referentially by the speaker of the sentences.

¹⁵What I describe here as "the de re/de dicto explanation" diverges from the account developed in Cumming (2008). Cumming's view is that the relevant semantic characteristic of "Rosaline" is the discourse referent it denotes. It is beyond the scope of the present discussion to do justice to Cumming's novel account of the de re/de dicto distinction.

Instead of looking for some ambiguity in the names, we should look for the source of the asymmetry in the structure of the embedded sentences. A critical move in the claim that (3.24) and (3.25) are evidence of a *de re/de dicto* ambiguity is the claim that the embedded clauses ought to be symmetric because identity is a symmetric relation. However, even if so-called identity sentences such as (3.26 a) and (3.26 b) do genuinely describe identities, and being identical is symmetric, it still does not follow that the two identity statements have the same meaning or even describe the same situation. Identity statements in English need not be symmetric just because identity relations are symmetric (or because identity sentences in the language of arithmetic are symmetric).

Before considering the symmetry of identity sentences in general, it may be useful to note that when we spell out what is meant by (3.25), it does not appear to be symmetrical, but yet does not seem to go beyond referential meanings. Put in other words (3.25) says that Biron thinks one particular person, Katherine, is a different particular person Rosaline. They are not the same person, and there seems to be a difference between thinking someone is Katherine and thinking they are Rosaline.

Aside from the fact that Katherine and Rosaline are not the same person, neither the speaker nor Biron think they are the same person, so we have no reason to think that for either of them, or from their perspectives, thinking someone is Katherine and thinking she is Rosaline is thinking the same thing.¹⁶ Theory aside, it is clear that switching the names should make a difference to the meaning, precisely because of who the names name.

It would be quite a different claim from (3.25) to say that Biron thinks that Katherine and Rosaline are just one person. Biron thinks there are two people, it's just that he thinks the one he is dancing with, who is Katherine, is someone else, Rosaline. Sentence (3.25) does not even suggest that Biron thinks that Katherine and Rosaline are collectively just one person. We can say consistently, Biron thinks Rosaline, the person he is dancing with, is

¹⁶As Cumming points out, both the embedded subject and predicate nominal are open to co-referential substitution (including using descriptions that Biron would not recognize as accurate). As such it seems that thinking of *x as being Rosaline* is being described in an objectual way in (3.25) – it does not depend on a particular perspective on Rosaline.

Katherine, and he also thinks Rosaline is outside smoking a cigarette.¹⁷ To put the matter in a semi-formalized, artificial manner: $\exists xy((\text{Katherine} = x \ \& \ \text{Rosaline} = y) \ \& \ \text{Believes}(\text{Biron}, x = y))$ is not an accurate representation of (3.25), not because it is too de re, but because Biron does not think Katherine and Rosaline are identical.

In order to give maximal plausibility to the argument that Cumming’s example does call for a de dicto analysis, we implicitly relied on considering what an identity sentence describes when it is true. If Katherine and Rosaline are identical, then saying that Katherine is Rosaline seems to be the same as saying that Rosaline is Katherine. But, if they are not in fact identical, the statements no longer appear to describe the same situation.

True identity sentences, like “Hesperus is Phosphorus” seem to be symmetric, and seem to mean roughly the same thing as equations ($H = P$). On the other hand, false identities do not seem symmetric, and do not seem equivalent to equations. Assuming the meaning of a sentence is prior to its truth, how should we reconcile these differences? One possibility to consider is to look at identity sentences as predications. Rather than a function of two variables ($=(x,y)$), the structure would be more like ($\text{is-}y(x)$). What it is to be y depends on what y is. Looking at the matter thus, it is easier to see “is a cat” and “is Hesperus” as two possible versions of the same form. It’s just that to be Hesperus is different than to be a cat, partly because cats are different than planets, but partly because Hesperus is a particular planet.

When x is y , to be x is the same as to be y . So the symmetry of true identity sentences follows from intersubstitutability of both the subject and the predicate nominal. Whereas the asymmetry of false identities follows from the difference between being one thing, and being a different thing. On this view, we can understand why “Biron thinks that Katherine is Rosaline” means something different than “Biron thinks that Rosaline is Katherine”, even though the names are being use exclusively to refer what they are names of.

¹⁷Compare “You think the killer and the desecrator are the same” from *The Exorcist*. Here Father Karras is in fact describing Lt. Kinderman as thinking that only one person is referred to by the two phrases “the killer” and “the desecrator”. Notably, he does not use an identificational copular construction of the sort so often discussed by semantic theorists.

3.6 Conclusion

Indirect discourse semantics is the semantics of clauses in which the complement of the verb is a sentence. Any sentence can be the complement of a subordinating verb; thus getting indirect discourse right depends to a large extent on getting a lot of other things right first. In terms of its place in the larger argument of my dissertation, my aim in this chapter has been to show that facts about indirect discourse are not incompatible with the pure-reference account of definite noun phrases. To show that indirect discourse does not present problems for pure reference, I showed that once the semantics of definite noun phrases is properly theorized, most of the difficulty of indirect discourse semantics is already taken care of. The remainder required a reconsideration of the semantics of copular constructions. By looking at how copular constructions are actually used in English, we saw that the idea that they have the same semantics as identity sentences in formal languages – an idea that is at the heart of a variety of semantic puzzles – does not fit the data.

Appendix A

Distributively Determined Noun Phrases¹

The topic of this dissertation is the semantics of singular definite noun phrases. As such, the semantics of distributively determined noun phrases like “every student” is not directly at issue.² However, the interpretation of singular definite noun phrases can be influenced by distributively determined noun phrases.

(A.1) Every boy thinks he will win the prize.

Sentence (A.1) could be used to describe a situation in which a group of boys are watching a game show, and each of them thinks that a particular male contestant is going to win a particular prize. If the sentence is used in that way, then the interpretation of “he” is as described in Chapter 1. Alternatively, the same sentence can be used to describe a situation in which a group of boys is competing for a prize, and each one of them thinks, “I will win the prize”. On this second use of the sentence, “he” is not understood as referring to a single individual. Rather, in some sense, it represents all of the boys.

¹This appendix is a development of earlier work done in collaboration with Jessica Pepp and Joseph Almog. See Pepp et al. (2015).

²I follow the terminology of Huddleston and Pullum (2002).

“He” can only be interpreted in the second way when it is in a precisely definable syntactic relationship with a distributively determined noun phrase.³ For this reason, we should expect this type of interpretation of “he” to depend on the semantic content of “every boy”. And so, in order to apply a theory of singular definite noun phrases to examples like (A.1), some understanding of distributively determined noun phrases is required. (The account developed in this appendix is preliminary and presented without argument.)

Distributively determined noun phrases are like singular (definite and indefinite) descriptions in that (i) they comprise a determiner plus a common noun phrase, (ii) they take singular verb agreement, and (iii) they can be either the subject or the object of a verb. On the other hand, distributively determined noun phrases are unlike singular descriptions in that (i) they cannot be the complement of a verb, (ii) they are never interpreted as signifying a single individual, and (iii) they have a syntactically definable sphere of semantic influence over the interpretation of other noun phrases.

A few examples will illustrate some of the relevant characteristics of distributively determined noun phrases.

(A.2) (a) Every student has five dollars.

(b) The students have five dollars.

Both (A.2 a) and (A.2 b) describe a group of students, and both attribute having five dollars. Sentence (A.2 b) can be interpreted either collectively or distributively: it can either describe a situation in which the group of students collectively has a total of five dollars, or a situation in which each student has their own five dollars. In contrast, (A.2 a) only has the distributed interpretation.

(A.3) (a) Every boy who is competing thinks he will win.

(b) Every boy is competing. He thinks he will win.

³Singular noun phrases only have this kind of interpretation when they are c-commanded by a distributively determined noun phrase. For a standard definite of c-command, see Carnie (2013).

As in sentence (A.1), sentence (A.3 a) has an interpretation such that “he” signifies one individual, and also an interpretation such that “he”, in some sense, represents all of the boys being described. In contrast, “he” in example (A.3 b) can only be interpreted as signifying a single individual.

What follows is a sketch of a semantic theory of distributively determined noun phrases that is compatible with the theory of singular definite noun phrases developed in this dissertation. I treat distributively determined noun phrases as indefinite descriptions plus a substitutionally interpreted quantifier.

(A.4) A student has five dollars.

Sentence (A.4) attributes having five dollars to a single student. On a natural interpretation, (A.4) means that one out of some particular group of students has five dollars, not that at least one of all the students in the world does. Likewise, on a natural interpretation, (A.2 a) means that each of some particular group of student has five dollars, not that every student in the world does. If they were to be spoken in the exact same circumstances, the group of students relevant to the interpretation of both (A.2 a) and (A.4) would be the same. Thus, sentence (A.2 a) is true just in case sentence (A.4) is true no matter which of the students in the relevant group “a student” refers to.

(A.5) Every[a student has five dollars].

We can analyze a sentence with a distributively determined subject as a sentence with a singular indefinite subject, which is then modified by a quantificational expression such that the whole sentence is true just in case the sentence in the scope of the quantifier is true no matter which of the students it is used to refer to. The phrase “a student” is still treated as referring to single individuals, but, one at a time, to more than one single individual.

If a distributively determined noun phrase is the object of a verb, then the quantifier only has scope over the object position.

(A.6) (a) The students like every teacher.

(b) The students like every[a teacher].

Sentence (A.6 a) is true just in case “The students like a teacher” is true no matter which of the teachers “a teacher” refers to.

Noun phrases that are understood as co-referential with a previous indefinite noun phrase are definite, even though the audience may only be able to identify their referents as the things referred to by previous expressions.

(A.7) A boy thinks he will win the prize.

In sentence (A.7), “he” is used to refer to some individual that the audience is expected to be able to identify. One possibility is that the audience is expected to identify the referent of “he” as the referent of “a boy”. If not, there will be some other connection from the use of “he” to its referent that the speaker expects the audience to be able to use for identification.

We are analyzing “every CNP” as comprising a quantifier “every” and a singular indefinite noun phrase “a CNP”. Crucially, if there is a second noun phrase in the scope of “every” that is understood as co-referential with “a CNP”, then it is a definite noun phrase.

(A.8) Every[a boy thinks he will win the prize].

By treating “every boy” as “every” plus “a boy”, we can see both how to interpret “he” on the model of (A.7), and also how to explain why it is a singular definite expression. Given an interpretation of (A.1) such that “he” represents each of the boys, “he” is a singular definite because it is understood as co-referring with the singular indefinite subject, “(a) boy”.

So-called donkey anaphora can be analyzed in the same way.

(A.9) (a) Every girl who received a prize thinks she deserves it.

(b) Every[a girl who received a prize thinks she deserves it].

In (A.9 a), “(a) girl” and “a prize” are indefinite because the audience is not expected to be able to identify their referents. On the other hand, “she” and “it” are definite because the

audience is expected to identify their referents as the referents of the preceding indefinite noun phrases “a girl” and “a prize”.

Finally, definite descriptions can also be interpreted in a similar way.

(A.10) (a) Every girl who received a gift removed the wrapping paper very neatly.

(b) Every[a girl who received a gift removed the wrapping paper very neatly].

In (A.10 a), “the wrapping paper” is understood to refer to the paper in which the referent of “a gift” was wrapped.

There is quite a bit more to be said on the subject of distributively determined noun phrases, and on their influence on the interpretation of other noun phrases. My intention here has been to give an idea of how the account developed in this dissertation can be extended to cover bound interpretations and donkey anaphora.

Bibliography

- Aristotle (1963). Categories. In Ackrill, J., editor, *Aristotle, Categories and De Interpretatione*. Clarendon Press.
- Booth, W. C. (1974). *A Rhetoric of Irony*. University of Chicago Press.
- Brennan, S. E. and Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6):1482–1493.
- Carlson, G. N. (1977a). *Reference to kinds in English*. Garland Publishing.
- Carlson, G. N. (1977b). A unified analysis of the English bare plural. *Linguistics and Philosophy*, 1(3):413–456.
- Carnie, A. (2013). *Syntax: A Generative Introduction, 3rd edition*. Blackwell Publishing.
- Clark, H. and Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22:1–39.
- Cohn, D. (1984). *Transparent Minds: Narrative Modes for Presenting Consciousness in Fiction*. Princeton University Press.
- Corazza, E., Fish, W., and Gorvett, J. (2002). Who is I? (indexical reference, answering machine paradox). *Philosophical Studies*, 107(1):1–21.
- Cumming, S. (2008). Variabilism. *Philosophical Review*, 117:525–554.
- Cumming, S. (2013). From coordination to content. *Philosophers Imprint*, 13(4).

- de Saussure, F. (1983). *Course in General Linguistics*. Open Court.
- Donnellan, K. S. (1966). Reference and definite descriptions. *Philosophical Review*, 75(3):281–304.
- Doron, E. (1988). The semantics of predicate nominals. *Linguistics*, 26:281–301.
- Fara, D. G. (2001). Descriptions as predicates. *Philosophical Studies*, 102(1):1–42.
- Fara, D. G. (frth). Names are predicates. *Philosophical Review*.
- Fine, K. (2007). *Semantic Relationism*. Blackwell.
- Geach, P. (1950). Russell’s theory of descriptions. *Analysis*, 10(4):84–88.
- Geach, P. (1967). Identity. *Review of Metaphysics*, 21:3–12.
- Groombridge, J. J., Jones, C. G., Bayes, M. K., van Zyl, A. J., Carrillo, J., Nichols, R. A., and Bruford, M. W. (2002). A molecular phylogeny of African kestrels with reference to divergence across the Indian Ocean. *Molecular Phylogenetics and Evolution*, 25(2):267 – 277.
- Heim, I. (1990). Presupposition projection. In van der Sandt, R., editor, *Reader for the Nijmegen Workshop on Presupposition, Lexical Meaning, and Discourse Processes*. University of Nijmegen.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9:183 – 221.
- Higgins, F. (1979). *The pseudo-cleft construction in English*. Garland Publishing.
- Hitchcock, A. (1944). *Lifeboat*. Twentieth Century Fox.
- Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, 3(1):67–90.

- Hobbs, J. R. (2010). Clause-internal coherence. In Kühnlein, P., Benz, A., and Sidner, C., editors, *Constraints in Discourse 2*, pages 15–34. John Benjamins Publishing Company.
- Hobbs, J. R., Stickel, M. E., Appelt, D. E., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, 63(1):69 – 142.
- Holland, S. (1985). *Better off Dead*. A & M Films.
- Huddleston, R. and Pullum, G. K. (2002). *The Cambridge Grammar of the English language*. Cambridge University Press.
- Kaplan, D. (1968). Quantifying in. *Synthese*, 19(1-2):178–214.
- Kaplan, D. (1989). Demonstratives. In Perry, J., Almog, J., and Wettstein, H. K., editors, *Themes From Kaplan*. Oxford University Press.
- Kaplan, D. (1990). Words. *Aristotelian Society Supplementary Volume*, 64:93–119.
- Kratzer, A. (1995). Stage-level and individual-level predicates. In Carlson, G. N. and Pelletier, F. J., editors, *The Generic Book*. The University of Chicago Press.
- Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Chierchia, G., and Link, G. (1995). Genericity: An introduction. In Carlson, G. N. and Pelletier, F. J., editors, *The Generic Book*. The University of Chicago Press.
- Kripke, S. A. (1977). Speaker’s reference and semantic reference. In French, P. A., Uehling Jr, T. E., and Wettstein, H. K., editors, *Studies in the Philosophy of Language*, volume 2, pages 255–296. University of Minnesota Press.
- Kripke, S. A. (1979). A puzzle about belief. In Margalit, A., editor, *Meaning and Use*, pages 239–83. Reidel.
- Leslie, S.-J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21(1):375–403.

- Lewis, D. (1975). Adverbs of quantification. In Keenan, E. L., editor, *Formal Semantics of Natural Language*, pages 178–188. Cambridge University Press.
- Liebman, D. (2011). Simple generics. *NouÛs*, 45(3):409–442.
- Love, C. (2014). Busting 10 snake-related email hoaxes. [Online; accessed 3-November-2014].
- Moore, G. E. (1944). Russell’s theory of descriptions. In Schilpp, P. A., editor, *The Philosophy of Bertrand Russell*, pages 175–225. Tudor.
- Nichols, P. (2013). Direct reference in indirect discourse. <https://ucla.academia.edu/PaulNichols>.
- Nichols, P. (2014a). Reference and communication. <https://ucla.academia.edu/PaulNichols>.
- Nichols, P. (2014b). Singular reference and generality. <https://ucla.academia.edu/PaulNichols>.
- Pepp, J., Almog, J., and Nichols, P. (2015). A unified treatment of pro-nominals in ordinary English. In Bianchi, A., editor, *On Reference*. Oxford University Press.
- Raffraya, C. N., Pickeringb, M. J., and Braniganb, H. P. (2007). Priming the interpretation of noun-noun combinations. *Journal of Memory and Language*, 57(3):380–395.
- Reynolds, T. A. (2014). Dispatches from Dar. *Emergency Medicine Journal*, 31(1):7–8.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69.
- Rothschild, D. (2007). Presuppositions and scope. *Journal of Philosophy*, 104(2):71–106.
- Saul, J. M. (1997). Substitution and simple sentences. *Analysis*, 57(2):102–108.
- Schlenker, P. (2008). Be articulate: A pragmatic theory of presupposition projection. *Theoretical Linguistics*, 34(3):157 – 212.

- Spade, P. (1994). *Five Texts on the Mediaeval Problem of Universals: Porphyry, Boethius, Abelard, Duns Scotus, Ockham*. Hackett.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457.
- Stanley, J. and Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15(2&3):219–61.
- Strawson, P. F. (1950). On referring. *Mind*, 59(235):320–344.
- Strawson, P. F. (1964). Identifying reference and truth-values. *Theoria*, 30(2):96–118.
- Tabossi, P. and Johnson-Laird, P. (1980). Linguistic context and the priming of semantic information. *Quarterly Journal of Experimental Psychology*, 32(4):595–603.
- Van der Sandt, R. A. (1989). Presupposition and discourse structure. In Bartsch, R., van Benthem, J. F. A. K., and van Emde Boas, P., editors, *Semantics and Contextual Expression*. Foris Publications.
- Wikipedia (2014a). Clark Kent — Wikipedia, the free encyclopedia. [Online; accessed 17-November-2014].
- Wikipedia (2014b). Imperial presidency. Online; accessed 3-November-2014.
- Wikipedia (2014c). United states secretary of homeland security — Wikipedia, the free encyclopedia. [Online; accessed 31-October-2014].
- Wood, J. (2008). *How Fiction Works*. Farrar, Straus & Giroux.