

2013-05-28

The Variety of Names

Fredrik B. Haraldsen

University of Miami, fredrikhar@yahoo.no

Follow this and additional works at: https://scholarlyrepository.miami.edu/oa_dissertations

Recommended Citation

Haraldsen, Fredrik B., "The Variety of Names" (2013). *Open Access Dissertations*. 1025.
https://scholarlyrepository.miami.edu/oa_dissertations/1025

This Open access is brought to you for free and open access by the Electronic Theses and Dissertations at Scholarly Repository. It has been accepted for inclusion in Open Access Dissertations by an authorized administrator of Scholarly Repository. For more information, please contact repository.library@miami.edu.

UNIVERSITY OF MIAMI

THE VARIETY OF NAMES

By

Fredrik B. Haraldsen

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

June 2013

©2013
Fredrik Haraldsen
All Rights Reserved

UNIVERSITY OF MIAMI

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

THE VARIETY OF NAMES

Fredrik B. Haraldsen

Approved:

Otávio Bueno, Ph.D.
Professor of Philosophy

M. Brian Blake, Ph.D.
Dean of the Graduate School

Risto Hilpinen, Ph.D.
Professor of Philosophy

Simon Evnine, Ph.D.
Associate Professor of
Philosophy

Stephen Neale, Ph.D.
Professor of Philosophy
City University of New York

HARALDSEN, FREDRIK B.
The Variety of Names

(Ph.D., Philosophy)
(June 2013)

Abstract of a dissertation at the University of Miami.

Dissertation supervised by Professor Otávio Bueno.
No. of pages in text. (357)

I defend a form of semantic pluralism about proper names. Proper names are, like pronouns, singular terms, and as such rigid and semantically simple. Though occurrences of proper names are often directly referential, I argue that proper names also – like pronouns – sustain anaphoric occurrences, in which case they work roughly in the manner of first-order variables. Instead of a unified theory of names in terms of their contributions to propositions expressed by sentences in which they occur, I argue that singular terms are unified in terms of a requirement that they take unstructured values, a requirement issued by compositional constraints associated with grammatical theory, and that these constraints are flexible with regard to the particular semantic values of occurrences of such terms. I develop a two-tier semantic framework to model the discourse contributions of such terms, and motivate a form of semantic two-dimensionalism that is more austere than most existing two-dimensional systems but that is nevertheless powerful enough to provide novel solutions to a range of traditional problems associated with singular terms, such as the problem of the contingent *a priori*, the problem of empty names, and – more tentatively – Frege’s puzzle.

ACKNOWLEDGMENTS

I will keep this short. It is impossible to thank everyone who has, in some way or other, made the writing of this dissertation possible, but it is impossible to get around the members of my committee: Otávio Bueno, Risto Hilpinen, Simon Evnine, and Stephen Neale, for their comments, help, patience and willingness to read this whole thing with a critical eye. I would furthermore extend my gratitude to the Department of Philosophy at the University of Miami in general, but must in particular mention Peter Lewis and Amie Thomasson, who have provided helpful feedback on parts of this work without having had any particular obligation to do so.

I would furthermore thank current and recent graduate students who have been willing to listen to me ranting about descriptive names, empty names, anaphora and indexicals at unreasonable lengths – though I cannot name them all, I need to mention at least Ryan Lake and Micah Dugas. In particular, however, I must thank my wife and fellow graduate student Robin Neiman for her patience and willingness to hear me out, even on occasions when I was probably talking *to* her rather than *with* her – her feedback and comments have been invaluable.

Parts of the work in this dissertation have also been presented at the Logic and Language conference in Aberdeen, September 2010, the 2011 ICLII conference in San Sebastian, Spain, and – not the least – the 2012 Central Division APA meeting in Chicago, and I wish to thank the participants there for helpful comments and feedback. In a different way it is hard to see how this work could have been completed without

generous grants from The Royal Caribbean Cruiselines (administered by the Norway America Association) and the 2012 International Student Summer Scholarship from the University of Miami.

Finally, I would have been completely at loss without the continued assistance and help from Lisa Tafares and Alex Puente, and they deserve the most heartfelt gratitude for their efforts.

TABLE OF CONTENTS

INTRODUCTION	1
1. RIGIDITY & REFERENCE	17
1.1. The Standard Arguments	17
1.1.1. The Modal Argument	18
1.1.2. <i>De Jure</i> Rigidity and the Semantic Argument	20
1.2. <i>De Jure</i> Rigidity and Reference	22
1.2.1. Kaplan on Direct Reference	23
1.2.2. Direct Reference and Object Dependence	26
1.2.3. Non-Referential Views	27
1.2.4. The Role and Explanatory Power of Rigidity	30
1.3. Rigidity Without Designation	33
1.3.1. Reference and Semantic Value	34
1.3.2. Two Shortcomings of (RD)	36
1.3.3. Recanati's Lexical Feature REF	40
1.3.4. Conclusion	42
1.4. A Concise History of Rigidity	44
1.4.1. The Need for Objectual Quantification	45
1.4.2. Quine's Paradox: Structured Vs. Unstructured Expressions	48
1.4.3. Lessons from the History of Modal Logic	51

1.5. Conclusion	54
2. SEMANTIC SIMPLICITY	56
2.1. Reference and Denotation	59
2.1.1. Two Kinds of Designation	60
2.1.2. Reference and Structure?	61
2.2. Meta-Linguistic Descriptivism and the Syntax of Names	66
2.2.1. The Syntax of Proper Names	67
2.2.2. Meta-Linguistic Descriptivism: Nominal Description Theory	68
2.2.3. Meta-Linguistic Descriptivism, Indefinites and Names With Free Variables	72
2.3. Deriving New Expressions	75
2.3.1. Irregular Occurrences	75
2.3.2. Partially Descriptive Names	77
2.3.3. Names in Language	79
2.4. Grammatical Kinds and Semantic Values	81
2.4.1. Preliminary Remarks: Compositionality and Semantic Value	82
2.4.2. The Semantics and Grammar of Singular Terms and Predicates	84
2.4.3. Unified Pluralism	86
2.4.4. Montague's Interpretation of Names As Quantifiers	88
2.5. Conclusion	92

3. THE CONTENT AND CHARACTER OF NAMES	94
3.1. Constants and Variables	95
3.1.1. Variables and Constants in Natural Language	95
3.1.2. Free Variables	98
3.1.3. Some Clarifications	102
3.1.4. Conclusion	103
3.2. Names As (Directly Referential) Constants	105
3.2.1. Indexicals and Constants	105
3.2.2. Reference and Character: Individuating Names I	108
3.2.3. Reference and Character: Individuating Names II	111
3.2.4. Names, Characters and Senses	112
3.3. Names As Pronouns	117
3.3.1. Names and Pronouns As Definite Expressions	119
3.3.2. Names As Anaphora	122
3.4. Conclusion	125
4. NAMES AS VARIABLES, NAMES AS CONSTANTS	128
4.1. Introducing Names	129
4.2. The Problem of Descriptive Names	132
4.2.1. Julius and the Semantic Argument	133
4.2.2. Semantics Vs. Metasemantics	136
4.2.3. Julius and the Contingent <i>A Priori</i>	140

4.2.4. Donnellan's Solution	144
4.3. Names As Variables	150
4.3.1. Names with Indefinite Antecedents	152
4.3.2. Alternative Accounts: Evans and Berger	158
4.4. Conclusion	162
4.5. Appendix: Arbitrary Names	164
5. PRONOUNS	170
5.1. E-Type Pronouns	170
5.1.1. Evans's Referential Treatment of E-Type Pronouns	172
5.1.2. D-Type Pronouns	173
5.1.3. Situations	177
5.1.4. Definiteness and Semantic Simplicity	179
5.2. Dynamic Frameworks	182
5.2.1. Discourse Referents and DRT	186
5.2.2. DPL: A Brief Outline	190
5.2.3. Whither Dynamic Semantics?	194
5.3. Notes on the Semantics of Variables, and Static Interpretations	199
5.3.1. Do We Need Values for Variables?	200
5.3.2. The Semantics of Variables: A Suggestion	202
5.3.3. Static Approaches to Cross-Sentential Variable-Binding	205
5.4. A Dynamic Pragmatic Account	210

5.4.1. Updateable Identity Functions (UIF)	211
5.4.2. A Pragmatic Interpretation	219
5.5. Conclusion	234
6. PLURALISM IN CONTEXT	237
6.1. Unified Pluralism	239
6.1.1. The Unified Theory	240
6.1.2. Names and Descriptions	242
6.1.3. Names Vs. Pronouns	245
6.2. The Ubiquity of Descriptive Names	246
6.2.1. Descriptive Names as Speech Act Failures	247
6.2.2. The Utility of Descriptive Names	249
6.3. Names As Pronouns	251
6.3.1. Names and Pronouns: Tracking Individuals	253
6.3.2. Names and Pronouns: Binding, Accommodation and Presupposition Projection	256
6.4. Anaphoric Names Across Contexts	259
6.4.1. Tests for (Semantic) Binding?	260
6.4.2. Escaping the Context of Introduction	263
6.4.3. The Notion of Context, and Theories of Reference	271
6.4.4. Transitions from Anaphoric to Deictic Names	274
6.4.5. Transitions in the Other Direction: Indeterminacy?	275

6.5. Conclusion	276
7. SOLUTIONS	278
7.1. Two-Dimensional Semantics	280
7.1.1. The Two-Dimensional Framework	282
7.1.2. Worries Concerning the Two-Dimensional Framework	286
7.2. Variables	289
7.2.1. Pronouns in Modal Contexts	289
7.2.2. “Julius” As a Variable	292
7.3. Empty Names	295
7.3.1. Descriptive and Empty Names	296
7.3.2. Empty Names and Fiction	298
7.3.3. Empty Names and Name-Using Practices	303
7.4. Application I: A Variable-Name Solution to Empty Names	308
7.4.1. Negative Existentials	309
7.4.2. Vulcan the Hypothetical Planet	314
7.4.3. Empty Names, Gappy Propositions and the Role of Descriptions	318
7.5. Application II: Frege’s Puzzle	325
7.5.1. A Solution to Frege’s Puzzle: The Basic Moves	327
7.5.2. The Role of Contexts	331
7.5.3. Kripke’s Puzzle	335

7.5.4. Geach Sentences	337
7.6. Conclusion	339
REFERENCES	341

INTRODUCTION

In this dissertation I defend a form of *semantic pluralism* about proper names. Names make different semantic contributions to sentences in which they occur, and their contributions correspond (in a sense to be made precise) to the different contributions of individual variables and constants of first-order logic. Drawing a distinction between two semantic types of names has been suggested before, but no previous account has to my knowledge developed the distinction in a fully satisfactory, systematic manner. Evans (1982) takes all names to be referential, though he employs a Fregean, truth-theoretic framework that allows different types of referential relations. To do so, he needs to assume a robust theory of reference on which the relation is determined in part by speakers' cognitive situation and epistemic relations. For Berger (2002), the distinction between different types of names is ultimately meta-linguistic and not reflected in the semantics of sentences in which names occur. A meta-linguistic pluralism, however, is too weak to solve the classic problems associated with referential views, in particular the problem of empty names or the problem of the contingent *a priori*.

The version of pluralism defend in this dissertation is semantic, and although I wish to avoid committing myself to a particular view of propositions, I assume a roughly neo-Russellian perspective. I also defend a view of reference that can be termed *metaphysically deflationary* – it sharply distinguishes questions concerning semantic properties from epistemic or metaphysical concerns about the referents or our epistemic or cognitive access to them. There is ample semantic and even linguistic evidence for a pluralist interpretation, which we can recognize regardless of how we divide labor

between semantics, epistemology or metaphysics. Nevertheless, philosophers of language have generally been reluctant to seriously consider semantic pluralism. Part of the reason, I suspect, is that the question of the semantic contribution of singular terms is often approached from the wrong end. If a systematic theory is assumed to require that all names make the same contribution to propositions – for instance that sentences containing singular terms in subject position uniformly express *singular propositions* – then a pluralistic account will not be viable. But a unified contribution to propositions cannot be the correct criterion for a unified theory of names. The first part of this dissertation, Chapters 1–4 is devoted to showing this, and to explaining what the correct level of unification may be.

First, there is the issue of explanatory role. Direct reference theories are motivated by certain semantic data, such as *de jure* rigidity. As I argue in Chapter 1, although *direct reference* seems to provide plausible explanations for a range of phenomena, a theory that requires that names make a single type of contribution to propositions will not have sufficient explanatory scope. Not only are there compelling counterexamples; such views also fail to explain features of natural language they should have explained, such as the rigidity of variables. Second (and accordingly), this approach does not extend to singular terms in general. *Pronouns* are singular terms, yet they clearly do not make a single kind of contribution to propositions. Demonstrative pronouns are directly referential; anaphoric pronouns work in a manner closely resembling individual variables of first-order logic. So if unification in terms of the propositional contributions of names were required one would have to give up the hope of a unified theory of *singular terms*.

A Systematic Pluralist Account

An explanation of the data can instead be given at the level of *requirements* associated with the grammatical category to which names and singular terms belong. All singular terms are semantically simple expressions (Chapter 2); their contents are semantically unstructured and their contributions sensitive only to individuals (and not descriptive material, for instance), and they can be defined in terms of functional role (captured by Montague's category *e*). That singular terms are associated with this requirement captures the distinguishing feature that makes singular terms useful – their role is to keep track of individuals independently of the various roles they may play (whereas predicates pick out those roles and quantifier phrases pick out individuals *in virtue* of their roles).

The requirement does not guarantee that all singular terms make the same type of contribution to propositions. In particular, it applies to pronouns as well as names, since it does not distinguish between constants and variables. Defining the semantic category in terms of grammatical category and associated requirements is sufficient to provide a systematic theory that predicts most data predicted by direct reference theories, and which also has a wider scope than accounts that start from contributions to propositions. The next, novel step is to argue that names exhibit the same semantic flexibility as pronouns. As Kripke's semantic argument shows, some name occurrences are indeed directly referential, and to explain such occurrences something more than satisfaction of the requirement is needed. But, as I show in Chapter 4, not all name occurrences are directly referential in this sense.

It is important to distinguish properly between *expressions* and *occurrences* of expressions in contexts. The requirement that the value of a singular term is an individual

is associated with the *expression*. Expressions do not make semantic contributions; *occurrences* – that is, an utterance or token inscription of a name – make semantic contributions. The contributions of names vary since the *expression* “Jack” may be used in different ways in different contexts. Some *occurrences* (uses or inscriptions) refer; some should be interpreted as what amount to being bound by antecedent occurrences of descriptions. Semantic pluralism does not operate with two types of names (expressions), but two types of *occurrences* or ways of interpreting utterances of names. The contribution of an occurrence is determined by context.

A crucial feature of singular terms is that they require a distinction between their linguistic meanings, their *characters*, and the semantic contributions of their occurrences, their *contents* (Kaplan, 1977). The lexical meaning of a singular term is a description – “the most salient male” or “the thing called ‘Jack.’” The semantic role of a singular term is to contribute individuals. The content is thus not the lexical meaning; the character instead helps *determine* content in a context. Now, it is common to say that the content is the individual that character and context pick out. In our framework it is more accurate to say the content of a singular term – regardless of whether it occurs as a constant or a variable – is *nothing but* an individual, and whether context seeks to supply a particular referent depends on the context.

Bound Names

There is ample evidence that names can occur in contexts where they do not refer but should rather be interpreted as variables (semantically) bound by some antecedent description. Among the most familiar cases is Evans’s “Julius,” introduced to designate

the inventor of the zip (whoever it may be). Given the stipulation, *Julius is the inventor of the zip* is knowable *a priori*. Since Julius, whoever he was, only contingently invented the zip, the proposition is also contingent. On Millian views, the proposition *Julius invented the zip* is equivalent to *Whitcomb Judson invented the zip*, so the latter is contingent but knowable *a priori* as well.

Millian views may, not implausibly, take *Hesperus is F* and *Phosphorus is F* to be equivalent, even in belief contexts, and even if the equivalence may appear informative to speakers (Salmon, 1986; Soames, 2002) – at least if belief contexts are closed under first-order validity. Furthermore, not all examples of the contingent *a priori* are obviously objectionable (“I am here now”). Nonetheless, *Whitcomb Judson invented the zip* should not be knowable *a priori*, so the purported equivalence of *Julius invented the zip* and *Whitcomb Judson invented the zip* is disastrous. Since *Julius invented the zip* is licensed by “someone invented the zip” (*a priori*), not only are belief contexts closed under first-order operations; since *Julius invented the zip* = *Whitcomb Judson invented the zip*, belief contexts will accordingly be closed under operations that are clearly *invalid* in first-order logic (a worry that does not arise in Hesperus/Phosphorus cases). Even appeals to special mechanisms for interpreting belief ascriptions (e.g. Crimmins, 1994) will not help. Since *Julius invented the zip* is knowable *a priori*, so is *Judson invented the zip*, and that follows directly from the identity of the propositions, regardless of such mechanisms. This result is a *reductio* of a Millian treatment of “Julius” (I discuss Donnellan’s purported solution in Chapter 4).

Occurrences of names in *explicitly* bound position are rare, but examples are provided in Geurts (1997), Elbourne (2005) and Cumming (2008). The fact that it is

possible to bind names is enough to show that not all names are directly referential. The important question thus becomes not a matter of whether they exist but of how common they are and the explanatory scope of the semantic pluralism they entail.

Names and Pronouns

I argue that the relationship between names and pronouns justifies viewing names as a special kind of pronoun. Names are, like pronouns:

- Associated with a character–content distinction
- Assigned content only in context
- Singular terms, and hence rigid and semantically simple.
- Definite expressions
- Able to occur in both bound and free positions

There are important differences between proper names and ordinary pronouns such as “she” or “it.” Names are for instance far more specialized – we possess a limited stock of ordinary pronouns, but a very large potential stock of names. The primary difference between names and ordinary pronouns, however, is a difference in (type of) character. The linguistic meaning of a pronoun calls for the most *salient* individual that fits a particular criterion (“male,” “female,” “thing”). The character of a name *n* is, roughly, “the thing called *n*.” There is no requirement that the individual picked out be the most *salient*, and we can hence use names to differentiate between individuals that e.g. belong to the same predicate extension and retrieve individuals from far earlier in a discourse, or

even across discourses. An ordinary pronoun cannot, since an ordinary pronoun, e.g. “she,” by default picks up the most salient female.

This difference makes names – in particular deictically used names – very useful, but it predicts some other differences in discourse roles between names and pronouns. Whereas the salience inflection predicts that ordinary pronouns pick up individuals (or *discourse referents*) automatically, a name must be explicitly attached to an individual. Furthermore, the characters of names call for *deference* to one’s linguistic community. What others call *n* will, in part, determine what an agent is referring to with *n*, so there is no infallible guarantee that the speaker knows the exact contribution of a name she uses even if the character provides the name with a determinate content. In particular, names can be associated with *name-using traditions* that may supply referents even when none is demonstratively available. Though I do not provide a theory of reference, it is at least plausible that the contribution of a name occurrence is often determined through a (perhaps causal) link to an individual originally baptized with that name (the expression).

If a referent *is* available in some way or other the nature of the character predicts that a name occurrence refers. Names prefer accommodation of referents over binding, and it is generally *expected* that a name occurrence pick up something other than the most salient discourse referent (otherwise “he” or “she” would be the preferred expression). Ordinary pronouns, on the other hand, exhibit a preference for binding; without additional effort (e.g. pointing), the salient male will be a discourse referent provided by an antecedent for “he” that serves as the binder.

If a name is introduced by description (what I call an *anaphoric* (use of a) *name*), a tradition associated with the name will not supply referent. Should we instead say that

an occurrence of “Julius” picks up the *antecedent* used to introduce it in a previous discourse? A name-using tradition can supply *contents*, but the content of an anaphoric name is not the antecedent description (its content is rather that of a variable semantically *bound* by that antecedent in the previous context). In Chapter 6 I argue that anaphoric names must be bound by locally available antecedents (though as opposed to ordinary pronouns, names support presupposition projection). In fact, anaphorically used names do, for a number of reasons, rarely escape their contexts of introduction. And when they appear to do so, there is no guarantee that the description associated with the name in a new context is the same as in the previous context. Descriptive names such as “Julius” (the expression) thus have potentially unstable meanings, and an occurrence of the expression in two different contexts may make different contributions. For anaphoric names, sameness of contribution across contexts may be a matter of degree.

Context, reference, and binding

A brief outline of the semantics of singular terms might go as follows. The proffered content of a singular term is that of a free variable. Fundamentally, “Jack is tall” and “she is smart” are “ x is tall” and “ x is smart,” respectively. The *expression* is associated with a requirement that “ x ” be taken by some individual of the domain, but the two sentences are truth-evaluable only relative to a context. There are two options.

First, we can draw on *extra-linguistic* context to provide a *referent*. In that case, the occurrences of “ x ” are directly referential, and are, at the level of logical form represented as a constant. Parallel to how the values of constants in first-order logic are assigned by the interpretation function, as part of determining a model, when “ x ” is

directly referential its value is assigned as a step in determining the *truth conditions* of sentences in which it occurs. Identifying Jack and *she* is necessary to identify the *meanings* of the sentences, and the values are fixed prior to evaluating the sentence at a world. Demonstrative occurrences are directly referential occurrences (Kaplan, 1977), and there is evidence – for instance from Kripke’s semantic argument against descriptivism (Kripke, 1971) – that names also have directly referential occurrences.

Second, one can draw on the *linguistic* context to find an antecedent binder.¹ Just as character can help locate an extra-linguistic value in a context, it can help locate an antecedent that can be used to bind it. For instance, “she is smart” may be bound by “some professor recently developed a new theory of *X*.” In that case, “she is smart” is true if the professor in question is smart. The *content* of “she” is nevertheless not the content of “the professor who recently developed a new theory of *X*” (this can be shown by embedding sentences containing anaphoric pronouns in modal contexts; see Chapter 5). The antecedent description constrains who may be *she* at a given world, but only individuals can be the content of “she.” It is, however, not specified *which* individual “she” refers to. It is for instance eminently possible to grasp the two-sentence fragment without knowing who she is, and in counterfactual scenarios, someone else might make the fragment true. The individual is no part of the *truth-conditions*. Rather some individual makes the sentence true (or not) at a given circumstance of evaluation.

Formally, I represent “*x*” as an identity function on a domain, or – more precisely – an inclusion function from the domain of the context into the domain of the predicate of

¹ It should also be noted that the *binding* in question is not *syntactic binding*, but some form of *semantic binding* – the antecedent issues a constraint on possible assignments, either by quantifying over assignment functions that are picked up by the pronoun, or by restricting the domain of potential satisfiers. We discuss these issues in Chapter 5.

an open sentence “ x is F .” For purposes of illustration, let the intension of a name be a function from worlds to individuals (I don’t think much hinges on this representation). Singular terms are, as mentioned, associated with a character–content distinction. The character of a singular term is associated with an identity function from context to the intension. (Of course, the character also expresses constraints on *where* to find a value, but these constraints are no part of the semantics.) If the context is *extra-linguistic*, then the input is supplied from the “outside.” The assignment is determined prior to circumstances of evaluation, and the identity function that is the proffered content of the name is “locked” to this input – it becomes, in effect, a constant (and total) function, which resets the values of the intension.

If, on the other hand, the relevant context is *linguistic* – determined by the contributions of previous sentences in a discourse – then the job of the character is to locate an appropriate antecedent binder. The individuals that satisfy the antecedent are the input to the identity function. This class of individuals varies across worlds, so the input (and output) of the identity function interpreting the pronoun varies with world of evaluation of the *antecedent*. The input nevertheless resets the intension to individuals of the antecedent class. It follows that the pronoun is rigid. Even if the pronoun occurs within the scope of a counterfactual operator, the relevant assignments are supplied by the world of evaluation of the antecedent, not the counterfactual world. The term is not referential, since the input varies with world of evaluation of the antecedent. In Chapter 5 I sketch a system on which free variables are represented as functions interpreted relative to a restriction on the domain of individuals. The restriction is determined by the

antecedent and recorded in a pragmatic layer of discourse, the *salience structure* of the discourse, the role of which is to keep track of anaphoric relations.

Applications

Semantic pluralism provides a degree of flexibility that we can use to solve several problems associated with direct reference views. An interesting upshot of treating certain names as bound is that we can mimic a form of two-dimensional semantics. As opposed, for instance, to Chalmers (1998) or Jackson (1998) the resulting system is consistent with a basically neo-Russellian framework, yet is nevertheless *semantically* two-dimensional (as opposed to the meta-linguistic two-dimensionalism of Stalnaker, 1998). It also has the significant advantage of eschewing the special operators defined in Davis & Humberstone (1980) and any commitment to several intensions being associated with single sentence occurrences. Variable binding and operator scope suffice to capture much of the power of two-dimensionalism, and these are features that any adequate semantic theory must recognize in any case.

As we show in Chapter 7 this “pseudo-two-dimensional” interpretation easily gets around apparently contingently *a priori* sentences. It also accounts for *empty names* such as “Vulcan,” introduced to designate the planet purportedly between Mercury and the Sun. I argue that empty names are anaphoric names. Their semantic contributions are hence determined relative to descriptive antecedents. By treating empty names as variables, atomic sentences containing them will come out false but meaningful, and the pseudo-two-dimensional framework can be used to account for a range of apparently conflicting intuitions associated with empty names. For instance, insofar as empty names

are semantically bound we obtain an elegant interpretation of negative existentials such as “Vulcan does not exist” (see Chapter 7); the contribution of “Vulcan” is determined by a presupposed or explicitly given antecedent. The sentence denies this antecedent, yet is nevertheless true (and non-contradictory).

Although it is not entirely clear that Frege’s puzzle *ought to* receive a semantic solution, I nevertheless sketch, somewhat tentatively, how pluralism and the pseudo-two-dimensional system could provide one. If “Jack believes that Hesperus is bright” and “Jack believes that Phosphorus is bright” can differ in truth-value it must be because the semantic content of the *that*-clause is somehow sensitive to Jack’s cognitive situation (otherwise Frege’s puzzle would not be a puzzle). We may capture such sensitivity by treating names in belief contexts as bound by presuppositions associated with the names. Now, it is presumably the context of the speaker *ascribing* the beliefs that determines these binders; the ascriptions will nevertheless be correct insofar as any of the provided binders are *extensionally* equivalent to anything Jack associates with Venus (not necessarily with the names – we can ascribe beliefs even if Jack does not know the names, or speaks a different language). The content of the ascriptions will hence not capture Jack’s cognitive states with any precision, of course, but that is not the point. Instead, semantic binding can provide a way of ascribing beliefs to Jack that makes the ascriptions sensitive to the fact that Jack is in *some* cognitive state that may not line up with the real world, and in such a way that the ascriptions can be invoked to explain his actions or reasoning in a manner consistent with viewing him as a rational agent.

Brief Chapter Outline

Chapter 1 starts from the observation that proper names are rigid. Indeed, names are *de jure* rigid; that is, rigid in virtue of semantic properties of the name rather than extra-linguistic (e.g. essential) properties of the referent. *De jure* rigidity calls for an explanation that avoids robust metaphysical or epistemological commitments. *Direct reference* might qualify as such, but I argue that an explanation in terms of the contribution a name makes to a proposition is inadequate, since it leaves certain natural language phenomena unexplained, such as the rigidity of variables. I suggest an explanation in terms of *grammatical function* rather than an explanation that relies, even minimally, on the referent. I also present a brief history of rigidity to illustrate how rigidity emerged from the development of the semantics of modal logic – in particular the move to *objectual quantification* in modal contexts requires rigid variables – and briefly what the history of modal logic shows us about natural language semantics.

Chapter 2 presents and develops the idea that rigidity is a result of a grammatical requirement associated with singular terms. Names are *semantically simple* and take individuals (and only individuals) as satisfiers. The requirement is associated with all singular terms, regardless of whether they function as constants or variables. I argue that alternative accounts, in particular forms of meta-linguistic descriptivism that take names to be semantically complex, are unsatisfactory. I also briefly discuss so-called “partially descriptive” names and occurrences of names in unusual grammatical constructions, and argue that such examples provide no compelling reason to modify the conclusion that names are semantically simple, complete noun phrases.

After a brief discussion of what the difference between *functioning as an individual variable* and *functioning as an individual constant* amounts to in natural language interpretation, Chapter 3 shows that some singular term occurrences must be interpreted as constants and are directly referential. Singular terms are associated with a character–content distinction, and I argue that all occurrences of a name “Jack” share a single character; a specific interpretation is provided for particular occurrences relative to contexts of utterance. The final part of the chapter lays the foundation for treating names as special-duty pronouns.

Chapter 4 discusses the evidence for the central claim in this dissertation: that names also occur in (semantically) bound positions. The evidence for this claim does not require commitment to any particular theory of reference. There is ample semantic evidence against treating all name occurrences as directly referential, independently of any particular (meta-linguistic) theory of *how* a referential occurrence comes to have the interpretation it has. Although it is correct that we often use names that function as variables – *anaphoric names* – because we find ourselves in situations where we have no cognitive or epistemic access to a referent, this does not mean that the semantic distinction between anaphoric and referential (or *deictic*) occurrences relies on epistemic or cognitive distinctions. Instead, this observation explains *why* natural language needs the semantic distinction between deictic and anaphoric occurrences.

Chapter 5 seeks to establish that anaphoric occurrences of pronouns must be interpreted as bound variables, and to explain how to achieve binding across sentence boundaries. I argue that *D*-type accounts, even if embedded in a sophisticated situation semantics, do not quite provide the solutions we need – in particular, such accounts do

not support a unified theory of singular terms – and find other static frameworks wanting (relying e.g. on parametrical covariance or contextually retrieved “off-stage” quantifiers). To achieve cross-sentential binding we need a dynamically updated discourse parameter. Representational dynamic theories, such as Discourse Representation Theory (DRT), may provide adequate accounts, but at least classical versions of DRT (such as Kamp & Reyle, 1993) encounter difficulties separating representational and model-theoretic levels if applied to all singular terms. “Pure” dynamic views, such as DPL (Groenendijk & Stokhof, 1990) or update semantics, avoid the worries. Although I have no conclusive argument against such accounts, they depart rather significantly from classical semantic treatments (despite reassurances from some defenders), and it would at least be worth trying to develop a more conservative extension of classical, truth-conditional semantics. I suggest (and outline) a two-tier account. Singular terms are represented as free variables interpreted as identity functions, but anaphoric relations are encoded in a (pragmatic layer of) *salience structure* of a discourse. The role of the salience structure is to encode anaphoric relations by providing appropriately restricted domains to serve as domains for the identity functions representing anaphoric singular terms. Identity functions thus serve the role of assignment functions in versions of DRT.²

In Chapter 6 I return to names. I explain in some detail why we need anaphoric names, discuss criteria for determining whether a given occurrence is anaphoric or deictic, and the extent to which it is up to the speaker – or even whether the speaker

² There may admittedly be other accounts that would have served our purposes as well, for instance by treating anaphoric pronouns as choice functions (Winter, 1997) or ϵ -terms (Peregrin & Heusinger, 2004), or by building the values of pronouns “locally,” as in the variable-free systems of Jacobson (1999) or Barker & Shan (2008).

needs to be aware of – whether a given occurrence is deictic or anaphoric. Although anaphoric names rarely escape their contexts of introduction there are exceptions. I describe the properties of context-escaping anaphoric names as well as how names originally introduced anaphorically can come to sustain deictic uses. A significant portion of the chapter is devoted to the relationship between names and ordinary pronouns, and how the differences are at least primarily upshots of differences in lexical meaning.

Chapter 7 applies our considerations to the problem of the contingent *a priori*, the problem of empty names, and various versions of Frege's puzzle, and compares the pluralist approach to alternative attempts to solve these problems.

CHAPTER 1: RIGIDITY & REFERENCE

I take it to be well established that ordinary proper names are, as argued by Kripke (1972), Kaplan (1977, 1989) and others, rigid. As a tentative definition I use Kripke's definition of a "rigid designator" as reconstructed by Kaplan (Kaplan, 1977, p.569):

(RD) A designator d of an object x is rigid, if it designates x with respect to all possible worlds where x exists, and never designates an object other than x with respect to any possible world.

(RD) is deliberately non-committal in one important respect. What, exactly, does d designate with respect to a world w the domain of which does not contain x ? There are two possible answers. Call a designator *persistently rigid* if it designates nothing with respect to w , and *obstinately rigid* if it still designates x in w (Salmon, 1982).³

1.1. The Standard Arguments

The thesis presented in *Naming and Necessity* is a two-step one. In addition to the claim that proper names are rigid designators, Kripke also argued against views according to which the meaning or semantic contribution of a name is given by a definite description, thereby showing that the descriptive picture of intentionality and content is wrong.

³ In *Naming and Necessity* Kripke seems to vacillate between treating names as persistently rigid (see for instance p.146) and obstinately rigid (see for instance p.78).

1.1.1. The Modal Argument

The gist of the *modal argument* for the rigidity of names is roughly as follows (my version is based on Stanley, 1997): Let n be a proper, ordinary language name. Suppose, for reductio, that (1) is *not* the case:

(1) If n designates o , then n designates o rigidly.

Given (RD) there are two possibilities compatible with the falsity of (1) (assuming that n designates at most one thing in each world):

- i. There is a world w in which o exists, but is not designated by n .
- ii. There is a w in which o does not exist, and n designates something other than o .

Note that the semantic content of n must be individuated prior to evaluating sentences in which n occurs with respect to possible worlds. Worlds in which n has a different semantic content are irrelevant. Given this assumption, and assuming that i) is true, then if n designates o with respect to the actual world, there is a world in which o exists but is not designated by n . Hence, (2) is true (where n is used, not merely mentioned):

(2) $\exists x(x = n \wedge \Diamond(x \text{ exists} \wedge x \neq n))$

But (2) is intuitively false, for it implies that Mark Twain, say, could have existed without being identical to Mark Twain. Similarly, if ii) is true, (3) is true:

$$(3) \exists x(x = n \wedge \Diamond(n \text{ exists} \wedge x \neq n))$$

And (3) seems false as well, for it implies that Mark Twain is identical to someone who could have failed to exist even if Mark Twain exists.

According to Kripke, definite descriptions are not rigid designators, so names cannot be equivalent to definite descriptions. To see this, replace n in (2) and (3) with a description (e.g. “the last great philosopher of Antiquity”), and notice that the results are *not* intuitively false. Thus, we have at least some inductive evidence for the non-rigidity of descriptions. More importantly, we have a reductio of the idea that “Aristotle” is equivalent to “the last great philosopher of Antiquity.” A definite description “the F ” denotes whoever satisfies “ F ” in a given world, whereas n – even though the referent of n happens to be the F in the actual world – denotes its referent in any world regardless of whether the referent is the F in that world. So, the semantic content of a proper name is not equivalent to the content of a description. Furthermore, assume that the content of “Mark Twain” were “the greatest American satirist of the 19th century” and consider the following sentence pair:

(4) Mark Twain lacked financial acumen

(5) The greatest American satirist of the 19th century lacked financial acumen

Consider a counterfactual situation w where i) Mark Twain lacked financial acumen but ii) was not the greatest American satirist of the 19th century, and iii) the greatest American satirist of the 19th century did not lack financial acumen. (4) is clearly true with respect to w , while (5) is equally clearly false. To put the argument slightly differently, consider w again. The following two claims, (6) and (7), seem true:

(6) The proposition that Mark Twain lacked financial acumen is true in w .

(7) The proposition that the greatest American satirist of the 19th century lacked financial acumen is false in w .

Hence the content of a proper name is not the content of a definite description.

1.1.2 De Jure Rigidity and the Semantic Argument

In fact, Kripke made the stronger claim that proper names are *de jure* rigid designators, where an expression d is *de jure rigid* just in case semantic rules link d to the object x it designates without mediation and without requiring that x satisfies some associated description in order for it to be the designation of d . That is, an expression is *de jure* rigid if “the reference of a designator is *stipulated* to be a single object” (Kripke, 1971, [1980] p.21, ft. 21). This characterization leaves the notion a little vague. Let us first characterize it negatively.

Some descriptions satisfy (RD). For instance, “the actual author of *Huckleberry Finn*” designates Mark Twain with respect to all possible worlds insofar as Mark Twain is the author of *Huckleberry Finn* with respect to the actual world. Likewise, “the square

of 3” (presumably) designates 9 in all possible worlds because 9 is necessarily the square of 3. Rigid definite descriptions satisfy (RD) in virtue of describing a condition that the denotation satisfies necessarily. In the case of “the actual author of *Huckleberry Finn*” the *non-semantic* fact that Mark Twain satisfies the condition described by “the actual author of *Huckleberry Finn*” with respect to all worlds ensures that the description is rigid. It is, however, rigid as a result in part of how things are with the individual Mark Twain. To find the designation of the description we must interpret its parts and then determine, by non-semantic means, which object satisfies the (semantically expressed) condition the definite description describes.

The designatum of “Mark Twain,” Mark Twain, is a different matter. Neither the name nor its parts expresses conditions that an individual satisfies in virtue of having any particular non-semantic property. The name is *unmediated* by descriptive conditions. A more apposite picture is that of a tag fastened onto the designation by stipulation and maintained by convention. Of course, it is a conventional matter that “square of” picks out the function it does, and that “nine” picks out the integer it picks out. It is, however, not a conventional matter (at least not in the same sense) that the property expressed by “the square of” yields the number denoted by “nine” for the number denoted by “three” as argument. The rigidity of a *de jure* rigid expression is, on the other hand, a matter of semantic properties of names, or of conventional, direct links between the name and its referent, and not a consequence of what properties the referent happens to have. In other words, *de jure* rigid designators satisfy (RD) in virtue of semantic conventions; *de facto* rigid designators happen to be rigid in virtue of non-semantic facts about the designatum.

The claim that names are *de jure* rigid designators is more controversial than the claim that they are rigid. The modal argument establishes only that names are non-equivalent to descriptions satisfied by different objects in different worlds. Evidence for *de jure* rigidity is provided by Kripke's *semantic argument*. A name "Aristotle" cannot be equivalent to a description cast in general terms along the lines of "the last great philosopher of antiquity" since "Aristotle was the last great philosopher of antiquity" is not *analytic*. Even if a linguistic community associated "the greatest pupil of Socrates" with Aristotle, "Aristotle was the greatest pupil of Socrates" should make a false claim about Aristotle, not a trivially true claim about Plato. The semantic argument does not rely on modality, though it has consequences for representations of modal contexts. If correct, it appears for instance to suggest that if the non-modal sentence "Aristotle was *F*" were evaluated at a different world, the referent of "Aristotle" should remain the same since it implies that the meaning depends on nothing but the referent, and a change in referent would entail a change in meaning. It also means that the value of "Aristotle" must be settled prior to the evaluation of "Aristotle is *F*" at a circumstance of evaluation, and hence that *locating the* referent of "Aristotle" is necessary to determine the intended meaning, and not only truth-value, for "Aristotle is *F*."

1.2 *De Jure* Rigidity and Direct Reference

(RD) is a criterion for the behavior of names in modal contexts, but the semantic argument suggests that there is a non-modal explanation for rigidity. Moreover, speakers do intuitively not need to have any understanding of modality to fully understand

sentences containing proper names. Consider speakers in a linguistic community who use proper names but have no modal notions (and are otherwise like us). Upon learning modal notions, for instance how to use “□,” would the speakers have to make any new stipulations concerning proper names? I am inclined to say “no.”⁴ This community, whose members already use and understand proper names, would only need to grasp “□” in order to use sentences involving proper names and “□” the way we do. If this is right, rigidity must be a consequence of a more fundamental property of names, though one that comes to the surface in particular in modal contexts.

1.2.1 Kaplan on Direct Reference

One common explanation is that names inherit rigidity from the fact that they are *directly referential*. Although rigidity does not entail referentiality, direct reference entails rigidity.⁵ Kaplan characterizes *de jure* and *de facto* rigidity as follows:

⁴ Smith claims that it is “self evident” that they do not (Smith, 1984, p.187), but this seems a bit strong.

⁵ *De jure* rigidity has sometimes even been taken to be equivalent to *direct referentiality*. While the identification makes the notion relatively clear I think it leaves us with too few resources to explain certain semantic phenomena (see below). One type of explanation of *de jure* rigidity that is particularly misleading is one along the following lines: “an expression *d* is a *de jure* designator of *x* if the fact that *d* designates *x* is independent of any (contingent) property *x* realizes.” First of all, the occurrence of “contingent” ensures that this definition will not be sufficient to distinguish *de jure* from *de facto* rigid designators (yet it could hardly be removed, since whether *d* designates *x* is hardly fully independent of what properties *x* realizes – it depends, for instance, on *x* having the property of *being designated by d in the actual world*). Second, and more importantly, it makes rigidity too closely associated with objects having (or not having) properties and hence with issues of metaphysics. I have argued elsewhere that rigidity has no (non-trivial) metaphysical implications (Haraldsen, 2012), and *de jure* rigidity should be a semantic property of certain expressions, not a metaphysical property of their designata. Another mistaken way of characterizing the distinction is:

[A] *de facto* rigid designator has constant reference with respect to all possible worlds in virtue of expressing a condition which attributes a certain individual essence to an object, so that its rigidity traces to extra-linguistic modal facts; a *de jure* rigid designator [...] has constant reference in virtue of the semantical rules of the language, so that it does not merely turn out that the term is rigid (McGinn, 1982, p.99).

A *de facto* rigid designator of x designates x in virtue of meeting some condition, i.e. via mediation of some description. *Ergo*: (2) A *de jure* rigid designator of x designates x *not* in virtue of x meeting some condition, i.e. *not* mediated by a description. [Kaplan, 1989, p.568]

According to Kaplan the feature captured by the notion of direct reference

[...] is not the *fact* that [an expression] designates the same object in every circumstance, but the *way* in which it designates the same object in every circumstance [Kaplan, 1977, p.495].

If a name is directly referential, then the semantic contribution of a name to the semantic content of a sentence in which the name occurs is just the referent, the object designated by the name. Since the name has no descriptive content that could be used to individuate different referents in different worlds, they must be rigid. If we think of propositions as Russellian structured entities, the semantic contribution of a directly referential term is

This characterization is insufficient, however (in fairness, though, McGinn calls the definition a “preliminary characterization,” so I am reluctant to saddle him with it). As Almog points out (Almog, 1986), on this definition actually-rigidified descriptions count as *de jure* rigid, insofar as it is the semantic rules for the actually-operator that determine that “the actual F ” denotes the same object with respect to every world. Actually-rigidified descriptions also satisfy the conditions for being *de facto* rigid, of course, since they are rigid in virtue of expressing a descriptive condition. To be theoretically useful, however, the categories should be mutually exclusive. Furthermore, the definition makes it hard to characterize descriptions such as “the person identical to Aristotle” or “the square of 3;” that is, descriptions that *contain* referential expressions. The fact that “the person identical to Aristotle” rigidly denotes Aristotle is determined at least in part by semantic rules of the language, namely the rule that “Aristotle” is a rigid designator. But even though “the person identical to Aristotle” contains a *de jure* rigid term, the description as a whole should be *de facto* rather than *de jure* rigid. McGinn argues that the definition also entails that ordinary speakers must be aware of certain complex semantic rules that they are certainly not aware of. However, it is strictly speaking sufficient that ordinary speakers *follow* these rules, not that they are aware of them; and if rigidity is, as I will claim, the result of basic *syntactic* features it is hardly surprising that speakers are able to follow them without being explicitly aware of them (after all, ordinary speakers are remarkably good at following ordinary grammatical rules without being aware of them).

the unstructured object itself, whereas the semantic contribution of a description is some complex entity built up by various attributes (see e.g. the preface to Kaplan, 1977).⁶

Whereas rigid designation is a modal notion, *direct reference* is a semantic notion, and since direct reference implies rigidity it gives us a non-modal explanation of the modal behavior of names. Relative to the intensional semantic framework Kripke probably had in mind it may be difficult to keep the semantic and modal aspects completely apart; if we analyze propositions as (functions from) possible worlds, the semantic content of a proper name is a constant function from worlds or world-time pairs to individuals, functions that would be partial or total depending on whether names are persistently or obstinately rigid. While this framework distinguishes “Aristotle is *F*” from “the last great philosopher of Antiquity is *F*” – the second may yield different values for different sequences of indices as argument; the first will not – it does not explain *why* names define constant functions or provide a way of distinguishing direct reference from *de facto* rigidity. If names are directly referential, however, the interpretation of a name is determined in a context in advance of evaluating a sentence at various worlds. One could think of names as constants whose values are supplied by the model itself rather than determined by satisfaction of some condition – identifying the referent of the name is a step in identifying the intended model rather than determined *in* the model. Names are thus insensitive to the world of evaluation. In that case they are obstinately – if *o* is the value of *n* independently of the evaluation of ... *n* ... with respect to a world *w*, it is

⁶ In fact, Kaplan argues that the idea that a term is directly referential is equivalent to the doctrine that sentences in which such terms figure express Russellian propositions. This claim, however, is based upon certain assumptions about the nature of propositions that are far from uncontroversial (e.g. Schiffer, 2002).

irrelevant to the semantic value of n at w whether o exists in w .⁷ The account suggests a version of *Millianism*, according to which proper names are directly referential in virtue of being, in Barcan Marcus's words, mere "tags" (Barcan Marcus, 1961); Mill's own formulation is similarly that names are "attached to the objects themselves, and are not dependent on the continuance of any attribute of the objects" (Mill, 1843, p.20).

1.2.2 Direct Reference and Object Dependence

A more flexible characterization of direct reference could be given in terms of *object dependence*. Roughly, the semantic content of a sentence "... d ..." is object-dependent if the truth *conditions* of the sentence depends on how things are with d . As Martí (2003) puts it the position is characterized in terms of the *presence* of an element – the referent – that makes the truth conditions object-dependent. That is, object-dependence requires that if there are mediating features between the name and the referent, these are truth-conditionally irrelevant. The Millian definition, on the other hand, is characterized by the *absence* of any kind of semantically mediating feature between term and reference. Hence, the Millian definition entails object-dependence but not vice versa.

Whereas defenders of a Millian view of names usually adopt a Russellian view of propositions, even Fregean frameworks could accept direct reference if defined in terms of object-dependence, insofar as grasping the contribution of a name to the truth

⁷ We may need a free logic to *evaluate* the claims involving o , but that is a different matter. Now, one worry that putatively arises by associating names with constants of first-order logic is syntactic. First-order constants are (by stipulation) complete noun phrases that have one and only one referent each but this is not obviously true for proper names in natural language. On one way of individuating names, several referents can share a name, and names can arguably also be quantified over (see in particular Katz, 2001; Bach, 2002; Baumann, 2010). Bach takes the observation to be evidence for meta-linguistic descriptivism, which we discuss in Chapter 2. I think a form of demonstrative or indexical theory is preferable to meta-linguistic descriptivism, and I argue that rather than being quantified *over* names sometimes occur *bound* by quantifier phrases.

conditions of sentences depends *only* on identifying the referent. For instance, whereas the object itself is the sole contribution to the truth conditions of an utterance and it is sufficient to understand an utterance to have *de re* knowledge of the referent, grasping the *proposition* or *thought* expressed by utterances such sentences may involve something more (e.g. an object-dependent *sense*). Non-Millian direct reference views include the “modest”⁸ truth-theoretic Fregean theories defended by Wiggins (1975), McDowell (1977, 1984), Evans (1982), and Sainsbury (2002, 2005). (See also Recanati, 1993, for a view on which singular terms are associated with descriptive information but the presence of an additional semantic feature REF guarantees that the truth-conditions of the occurrence of a sentence containing a name depends *only* on the extension of the name.) In the following I concentrate on Russellian versions of direct reference theories, but hopefully much of what I say can be incorporated in a truth-theoretic view.

1.2.3. Non-Referential Views

The idea that names are *de jure* rigid or directly referential is controversial. It has been argued that names are equivalent to descriptions that always take wide scope in modal contexts (Dummett, 1973, 1981, 1991; Stanley, 1997; Sosa, 2001). If this were the case then names would satisfy the modal argument, but not obviously the semantic argument, and names would not be *de jure* rigid or referential. Alternatively, it has been argued that

⁸ A theory of meaning is modest (the term was introduced by Dummett, 1975) if the meaning of an expression cannot in general be analyzed in terms of elements or mechanisms in a manner that would be sufficient to explain the meaning of the expression to someone who does not already understand it. On McDowell’s view, for instance (McDowell, 1977), the meaning of a name “Hesperus” is given by the semantic axiom “‘Hesperus’ refers to Hesperus,” grasp of which is sufficient to understand “Hesperus,” even though no further analysis of the meaning of the expression can be provided – the axiom is not supposed to provide a definition of the sense, but to *display* it. In general, grasping homophonic referential (or *interpretive*) axioms will be necessary and sufficient for understanding a referential expression, but does not provide an *analysis* of its meaning.

names are equivalent to *actually-embedded descriptions* (Nelson, 2002). The actually-operator assigns to a predicate its extension at the actual world with respect to all worlds, so the individual that satisfies the description will not vary across worlds, making the description *de facto* rigid. I will not discuss these views at any length here. Such views have been forcefully criticized in e.g. Soames (2002, 2005), Everett (2005), and Forbes (2011). I will, however, mention a general worry. Notice that the sentence “Alexius is a philosopher” *appears* to be entirely non-modal. If “Alexius” were shorthand for (or even has its reference determined by) an actually-rigidified description some modal operator or at least some apparently irreducible modal feature would be involved. Understanding “Alexius,” however, does not seem to require understanding the @-operator – a linguistic society could surely have names without having the operator (Davies, 2006). More importantly, analyzing names as @-rigidified description incurs an explanatory commitment that is not easily cashed out. Descriptivism is motivated by the cognitive significance of the use of names, for instance by the fact that co-referential names appear not to be substitutable in belief ascriptions; but cognitive significance does not explain why a name behaves as a *rigid* rather than non-rigid description.

If names are directly referential, then their rigidity is explained – without invoking modality – by the absence of any semantic element in its interpretation that could be used to individuate a different referent at a different world. Hence, direct reference explains, as mentioned earlier, the rigidity of a name in virtue of its semantic properties, and the name’s modal behavior is predicted by those non-modal properties. If names are @-rigidified descriptions, on the other hand, modal properties are built into the *semantic* properties of names. The behavior of the name in modal contexts is correctly predicted,

of course, but if we ask *why* names behave the way they do in modal contexts we encounter what seems to be an uncomfortably tight circle. The obvious answer, that a name behaves as a rigid rather than non-rigid description as a result of the @-operator, is circular since the justification for the idea that names are rigid is precisely that it gives them the correct modal behavior. But that leaves the fact that names behave the way they do in modal contexts unexplained; instead, the correct modal behavior is achieved by stipulating that the name has a certain modal property, namely that its extension is fixed to the actual world, and it is hard to see any *independent* motivation for the mechanism that ensures the behavior, the insertion of the @-operator. Hence, the motivation for *rigidified* descriptivism seems *ad hoc*, introduced to accommodate Kripke's data by stipulating that sentences containing names have a particular modal profile. There are hence *prima facie* methodological worries with rigidified descriptivism.

Wide-scope descriptivism has related problems. Dummett recognizes that taking wide scope in modal (though not hyperintensional) constructions should be the result of a syntactic requirement, but in addition to the problem of how to deal with mixed modal–hyperintensional contexts (Soames, 2002), the move seems to lack independent justification. It is also unclear how to motivate a syntactic or meta-linguistic requirement different from the one the defender of direct reference would appeal to, namely a requirement along the lines of “takes an individual and only an individual as value.” There may be ways around these objections, or reasons not to worry too much, but at least direct reference theories have an advantage in providing a non-modal, non-circular explanation of the behavior of names in modal contexts. It is for instance plausible that

the explanation for the modal behavior of a name is a result of some semantic rule (perhaps derived from a cognitive need to distinguish objects from the roles they play).

1.2.4. The Role and Explanatory Power of Rigidity

Historically it was often thought that rigidity was an obscure essentialist doctrine. Quine, for instance, claimed that “[a] rigid term differs from others in that it picks out an object by its essential traits” (Quine 1977, p.118). Others have argued that the modal truth conditions of a sentence involving a rigid designator are ascertained by the *haecceity* or “primitive thisness” of the designatum (see Adams, 1979). The idea that rigidity should do some metaphysical work is also explicit in Lewis’s rejection of the rigidity of names (Lewis, 1986). Much work by Kripke and subsequent authors has been directed toward showing that rigidity does not have essentialist consequences (for a lucid survey, see Sullivan 2005). For instance, even if metaphysical considerations should favor the existence of *haecceities* the issue is orthogonal to the question of rigidity; singular term rigidity is, as Kaplan points out, “prior to the acceptance of essentialism, not tantamount to it” (Kaplan 1986, p.265). Instead of expressing metaphysical commitments, the idea that names are rigid designators dispels certain mysteries concerning how we ensure that sentences containing names behave the right way in modal contexts – instead of having to search for the proper referent in a world we simply stipulate that we are talking about a world containing this referent.

By the same token the behavior of names with respect to modal operators and the truth conditions of modal sentences containing names, are not *explanations* of their rigidity. Rather, intuitions about the truth-values of counterfactual sentences containing

names are *evidence* for the claim that names are rigid. As just argued this observation is worrisome to views that analyze names as actually-rigidified descriptions, since the behavior of sentences containing names would then also figure in the purported explanation for why names are rigid. The worry, however, is in fact an upshot of a general worry about the explanatory role of rigidity.

It is commonly assumed that *rigidity* plays an important explanatory role, for instance in accounting for *a posteriori* necessary truths. To take an example at random:

For example, given that “Hesperus” and “Phosphorus” are rigid designators, and given that they are coreferential in the actual world, it follows that they designate the same object in all possible worlds. As a consequence, “Hesperus is Phosphorus” is necessary if true, even though its truth or falsity can only be known *a posteriori*. The rigidity of proper names does, then, seem to play a key role in the explanation of the necessity of *a posteriori* identity statements between names.
[Haukioja, 2012, p.400]

But insofar as rigidity is a modal notion to begin with one may worry about the extent to which modal truths *could* be explained by the rigidity of names. It is for instance not immediately clear why rigidity should even be explanatorily prior to the necessary truth of identity statements involving names.

Direct reference, a semantic property, is explanatorily prior to the necessary truth of “Hesperus = Phosphorus.” Since names are referential, neither “Hesperus” nor “Phosphorus” could refer to anything but Venus without a change in meaning. Given that the names mean the same in modal contexts it follows that “Hesperus = Phosphorus” is necessarily true without reference to extra-linguistic properties of the referents. It might of course be that Venus is essentially a planet, but since “Hesperus” and “Phosphorus” are directly referential, “Hesperus = Phosphorus” is necessarily true (true in all worlds

where Venus exists) even pending this metaphysical assumption. Of course, we need to assume that “ $\alpha = \alpha$ ” is a necessary truth (at least in worlds where α exists); this is not an essentialist assumption about α but a consequence of the reflexivity of identity.

The explanation for the necessity of identity just given has no substantial essentialist consequences, and entails that standard metaphysical arguments for contingent identity, such as “Ship of Theseus” cases, questions about the relationship between statues and what they are made of, and questions about personal identity over time, are turned into questions of whether the same referent is assigned, in which case identity statements would be true, or whether different referents are assigned, in which case identity claims could be false. In other words, these issues become metaphysical concerns regarding the persistence conditions of an individual referent over time or across worlds, or concerning the identity conditions for the referents. Direct referentiality does not suggest any particular answer to such questions.

But the explanation of the necessity of identity just given does not invoke *rigidity* either. Nor is rigidity needed to explain why “Mark Twain” is not equivalent to “the author of *Huckleberry Finn*.” Since the semantic content of “author of *Huckleberry Finn*” is a property (or at least a function from individuals to truth values), not an individual, and predicates have different extensions at different worlds, the denotation, Mark Twain, is no part of the semantic content of the predicate. Since Mark Twain *is* constitutive of the meaning of the name, the name and the description (at least assuming that definite descriptions are quantifier phrases whose meaning are functions from the meaning of their parts) cannot have the same meaning. So the difference in meaning is predicted without invoking rigidity, and even if definite descriptions were rigid, a description could

not be semantically equivalent (in the sense of making the same contribution to a proposition) to a name of its denotation. Hence, just as metaphysical intuitions about truth-conditions of modal and counterfactual sentences are evidence for the rigidity of names, rigidity is a datum for an adequate semantic account of names rather than a fundamental feature that does explanatory semantic work on its own. Though it has not always been understood in this manner (see Soames, 2002, for instance) direct reference facilitates a very metaphysically deflationary view of the semantics and modal behavior of proper names (at least as long as we keep the question of semantics apart from the meta-linguistic question of *how* a name, or occurrence of name in a sentence, gets the meaning it has). In the next section we will, however, argue that *de jure* rigidity is the result of even more fundamental constraints associated with grammatical role.

1.3. Rigidity Without Designation

In the last section we gave some reasons for thinking that rigidity should be explained by non-modal properties, and that direct reference appeared to offer a plausible explanation. The feature we are after is the feature of names that explains their truth-conditional contributions and why names for instance function as rigid designators. In this section I argue that insofar as direct reference is understood in terms a relation between a name and a particular referent as expressed by instances of a disquotation schema of the form “‘*n*’ refers to *n*,” then direct reference cannot be an adequate explanation of *de jure* rigidity. A point that is easy to overlook – especially if one presupposes a particular view of *propositions* – is that rigidity is (in a sense to be explained) more fundamental than

designation. There are expressions that behave, in modal contexts, in a manner that seems to require the same explanation that is used to explain the rigidity of names, but which are not designators at all. I will argue that the fact that certain expressions are rigid designators of particular objects is explained by more fundamental constraints on possible values associated with the grammatical category to which they belong and a principle of compositionality (I develop a more precise account in following chapters).

1.3.1. Reference and Semantic Value

The category of singular terms, and the properties that define singular terms, should be identifiable at the level of grammatical category. This is the functional level at which such terms are distinguished from other types of terms with other types of semantic values, and hence the level that determines which type of semantic value an expression receives. We know that something is a singular term by knowing its grammatical function. Assuming that such terms are directly referential we can derive rigidity from strict compositionality. Insofar as the meaning of a sentence is a (strict) function of the meanings of its parts, the occurrence of any non-ambiguous, non-vague expression (after indexical elements and context-sensitivity are resolved) should

- I) Receive an interpretation independently of the interpretation of other elements in the sentence, such as modal operators.
- II) Any occurrence of elements that are represented as the same at the level of logical form should receive the same value, where semantic value is understood as the entity that grammar and context assigns to the expression.

Let us call this the same-semantic-value requirement, or SSV. Even descriptivist views accept SSV (at least with respect to modal contexts insofar as these are not what Frege would call “*ungerade*”); descriptivists would rather deny that the semantic value of a name is just a referent but, perhaps, a Fregean sense. If names are directly referential, however, the value of a name is a referent, and SSV requires that it refers to the same individual on all occurrences, regardless of the scope of modal operators. But then it follows that the name satisfies (RD).

Importantly, this derivation of rigidity does not quite require that a name is directly referential (in the sense that it actually refers to an object), since the derivation does not assume that the semantic value of a sentence depends on a *particular assignment* of a value to a singular term. It requires only a certain restriction on the *kinds* of values the singular term can take. The *requirement* that the name can only take an individual as value is sufficient, assuming SSV, to obtain rigidity, so this inference to rigidity from grammatical constraints entails no commitment for instance to whether the name’s satisfier or referent is determined prior to circumstances of evaluation or not. In particular, distinguishing between individual variables and constants is not obviously possible at this level. If this is correct it is not a characteristic property of names or singular terms *as a grammatical category* that such terms correspond to constants of first-order logic, refer to particular individuals or even satisfies the disquotational schema of reference above. I will later argue that singular terms – even names – can occur as both constants and variables. All singular terms are nonetheless (*de jure*) rigid, non-

descriptive, sensitive only to extensions (and belong to Montague's category *e*), since these properties *do* follow from the characteristic properties of singular terms.

The idea that names function as individual *constants* or are directly referential is not inferred from requirements associated with grammatical category, and is not required for giving a non-modal explanation for rigidity, but is motivated by at least two factors: first, characterizing the semantics of names in terms of their contributions to a proposition rather than grammatical role (and the type of semantic value grammatical role requires); second, Kripke's Semantic Argument. The semantic argument goes further than what is required for *de jure* rigidity. For an expression " α " to satisfy the semantic argument, which implies a commitment to permutation invariance under assignments, " α " must be directly referential (and the referent of " α " must be determined prior to evaluation, i.e. be determined as part of determining the intended *truth conditions* of a sentence). The semantic argument is hence an argument for direct reference, not for *de jure* rigidity. Not all singular terms (indeed, not all names, as we argue in Chapter 4) satisfy the semantic argument, and satisfying the semantic argument requires constraints that go beyond the constraints that define the semantic contribution of singular terms and distinguish them from other types of expressions, and beyond what is needed to account for the fact that all singular terms are rigid.

1.3.2. Two Shortcomings of (RD)

(RD) captures one way in which expressions can be said to be rigid, but not the most fundamental one. This is easy to overlook partially, presumably, because of a tendency to

focus on *reference* (in fact, *de jure* rigidity is often equated with direct reference⁹). *De jure* rigidity should not be viewed as a matter of referring but as a feature explained by semantic requirements associated with grammatical category. To argue for this we will show that if direct reference and *de facto* rigidity were the only two notions of rigidity in play, there would be expressions whose behavior in modal contexts would be inexplicable. These are best characterized as rigid, even *de jure* rigid, since they are associated with the grammatical requirement that they take only individuals as values, but they do not *designate* these values. Consider:

$$(8) \exists x(Fx \wedge \Diamond \neg Fx)$$

This sentence is true if something is *F* but could have failed to be so. A proper evaluation of (8) requires that the same individual is assigned to both occurrences of “*x*,” even though the conjuncts are evaluated with respect to different worlds (assignments that provide different values to the two occurrences are irrelevant), so “*x*” is rigid.

It is common to treat free variables under assignments of values to variables as directly referring to the assigned value. If assignments are associated with worlds, and circumstances of evaluation are time-world pairs, then the value of the variable that satisfies the sentence will vary with the circumstances of evaluation. It is also worth pointing out that given standard Tarskian semantics the truth of a sentence containing

⁹ E.g. in Recanati (1993) and Fitch (2004). The idea that expressions can be *de jure* rigid but not directly referential is suggested in Kaplan (1989, p. 577 ff 25; Salmon 2003), but the examples they have in mind are names disguised as (rigid) descriptions, since these are – strictly speaking – rigid by *stipulation*. However, the notion of stipulation that grounds *de jure* rigidity is not just any act of stipulation – the central idea is that the rigidity in question is determined not by individuation through composition of semantic values at a circumstance of evaluation, but by assignment to a name.

variables does not even depend on the (sequence of) element(s) actually assigned by the assignment function, but on the existence of an appropriate *variant* (sequence of) element(s). If “ F ” is ‘is a current president of the US’ then “ $\exists xFx$ ” is true since Obama is, in fact, such a president but could have failed to be so, and in Tarskian semantics “ $\exists xFx$ ” is satisfied relative to an assignment that assigns Duns Scotus to “ x ” as long as Obama occurs in the right position in a variant assignment. Surely this does not mean that “ x ,” as it occurs in “ $\exists xFx$ ” refers to Scotus (see Salmon, 2005). On the other hand, the values of “ x ” are *assigned*, not compositionally determined, so it would be a mistake to think that “ x ” is *de facto* rather than *de jure* rigid. Nor does “ x ” have any descriptive content – it is *bound* by a descriptive expression, but there are no mediating descriptions between the variable (under an assignment) and the value assigned.

Even if free variables refer under an assignment, the variables in (8) are bound. An assignment is a step in a recursive definition of truth for a quantified sentence, not part of the semantic content of a sentence in which the variable occurs bound. Regardless of what individual is assigned, the assignment does not turn up in *the proposition expressed* by (8) or any sentence that is properly regimented as (8). The proposition is thus not object-dependent, since it is no barrier to fully grasping the truth conditions of (8) that one fails to know which element satisfies 8 or even whether such an element exists. Bound variables do not *designate* individuals at all; rather, variables *range* over a domain of individuals. But the bound variable is nevertheless best characterized as rigid, given that *any* proper assignment to “ x ” must assign the same individual to both occurrences regardless of the modal operator – just *reading* the sentence “there is something which is F , and it could have failed to be F ” shows that the *bound* variable –

here “it” – must be understood as a rigid expression. It is a singular term, non-descriptive, and associated with the same grammatical category as a name. Intuitively the rigidity of the variable should receive the same explanation as the rigidity of a proper name. The rigidity in that case is not a matter of satisfaction, but of a constraint on semantic values issued by syntax *prior to* evaluation for satisfaction of a sentence containing it. But if *de jure* rigidity is a matter of *reference*, understood as a relation between an expression and a selected object, we have no explanation for the behavior of “x.” Of course, the worry is not be very pressing unless there are expressions in natural language that correspond to or are most appropriately represented as bound variables. I argue (Chapter 5) that there are such natural language expressions, the most obvious being anaphoric pronouns.

Now, it is of course possible to argue that the bound variable is, technically speaking, not rigid since it is not a designator and does not satisfy (RD). This, however, is mostly a terminological point. What we are looking for is the feature that explains the behavior of names and singular terms in modal contexts, and we have argued that the explanation can be cast in terms of grammatical constraints on what value an expression can take. The similarity between names and variables in modal contexts is sufficiently striking to justify the claim that the feature that explains the behavior should be the same, and it cannot be a Millian notion of direct reference. Questioning whether the phenomenon to be explained, the characteristic modal behavior, is strictly speaking captured by the notion of *rigidity* amounts to questioning whether rigidity is any interesting semantic property at all compared to some slightly more general property. I will persist in calling the feature “rigidity.” Since the value is not a matter of being

determined through satisfaction of a compositional condition expressed determined by the expression, the variable is not properly characterized as *de facto* rigid either.¹⁰

A second possible worry is the apparent rigidity of *empty names*. Empty names do not refer but appear to be meaningfully used in counterfactual claims. Defenders of direct reference may argue that since “Vulcan” fails to designate it must be a confusion to think that it *rigidly* designates anything. The question, however, is why empty names behave as if they were rigid in ordinary use. Insofar as Kripke’s modal argument is compelling at all, examples using “Vulcan” should be equally compelling to an audience who believed Vulcan existed. Such names are deficient and seemingly parasitic on ordinary names; but what feature of ordinary names are ordinary speakers ‘carrying over’ to uses of such parasitic names, e.g. when they make counterfactual claims? The fact that they do make such claims suggests that there is some rule or meta-linguistic requirement attached to names as a category. The best explanation for why rigidity systematically extends to empty names is that the requirement is associated with a *syntactic* or grammatical kind – the type of proposition expressed by a sentence containing a name is secondary.

1.3.3. Recanati’s Lexical Feature REF

As an alternative to Millian views, where direct reference is understood in terms of object-dependence rather than singular Russellian proposition, Recanati (1993) offers an account along the lines we are looking for. Recanati argues that the rigidity of proper

¹⁰ Admittedly an argument is required here. In (8) one could argue that the value of the variable *is* determined in this manner, since even the second occurrence of “*x*” occurs within the syntactic scope of the quantifier (and may be considered, in principle, as *part* of a single, restricted quantifier). In a natural language translation of (8), however, the expression corresponding to the variable (e.g. “it”) will occur outside the c-command of the quantifier. As we show in Chapter 5 the occurrence of “it” needs a separate semantic value.

names is the result of structural features. Ordinary proper names are equipped with a lexical feature REF that guarantees that they are directly referential, and that any description users associate with them is truth-conditionally irrelevant. Recanati does not provide a complete story of *why* names have this property, deriving it instead from their modal behavior. Unfortunately, since REF is also what explains the modal behavior of names, we get a relatively tight explanatory circle (though Recanati does offer some considerations pertaining to mental content that may be seen as a partial justification). Semantically the REF-feature is a property that names automatically take on by being the type of linguistic entities they are. Insofar as names are (semantically) equipped with REF we have an explanation of why empty names are rigid, since the use of an empty name invokes a rule associated with the kind of expression it is. One remaining worry is that if REF is something along the lines of a semantic rule speakers may need to be at least tacitly aware of it, or at least intend to use proper names in a manner that can be modeled as involving the feature.

I suggest that the feature REF is superfluous. Singular terms, including names, demonstratives and variables, are *de jure* rigid as a result of being grammatical units that have no scope potential and no parts among which their semantic contribution can be distributed. Rather, their values are assigned and therefore they have – as opposed to quantifier phrases – no semantic structure that can be used to individuate different individuals at different worlds (I develop the positive view further in Chapter 2). That the role of a name is to contribute an individual does not entail that all name occurrences will contribute *specific* individuals. In combination with SSV this functional characterization implies rigidity. However, although the requirement ensures that names are rigid it does

not entail that a given name satisfies Kripke's semantic argument, since it does not guarantee that the truth *conditions* of a sentence will depend on which individual is assigned to a name, even if the expression is *de jure* rigid. If "Jack" functions as a bound variable in "Jack is *F*," then "Jack" should *not* be immune to reassignment if the world of evaluation for the sentence is changed, though it is rigid – immune to reassignment – relative to modal operators occurring within a sentence, something that, say, the extensions of predicates or definite descriptions are not (the significance of this becomes particularly vivid when we consider, in Chapter 5, how variable expressions are semantically bound by presuppositions or antecedents occurring far earlier in a discourse). If names satisfy the semantic argument – and I will argue that many names do – the features required to do so come in addition to the features required for names to be *de jure* rigid.

1.3.4. Conclusion

Singular terms are noun phrases (or occur as noun phrases in natural language). Noun phrases can also be quantifier phrases. Quantifier phrases do not rigidly designate extensions; rather, they have complex (second-order) functions as semantic values and *determine* different extensions (or operations on extensions) relative to different possible worlds. As a working hypothesis we may perhaps follow Neale (1993, 2009), Dever (2001) and Sullivan (2007, 2009) in taking all noun phrases to be either (rigid) singular terms whose role is to pick out individuals, or structurally complex quantifiers whose contributions to semantic evaluations of sentences in which they occur are complex intensions. Although some quantifier phrases, such as "the smallest prime number,"

happen to (*de facto*) rigidly designate their extensions, they are not rigid designators in virtue of the semantic function of such phrases, but because of the properties of the individuals in their extensions.

Since singular terms have no semantic structure that could determine a referent relative to a circumstance of evaluation, they must have their values assigned (or stipulated e.g. by demonstration). Their semantic values are hence unstructured as well; that is, although the individuals in question may of course be metaphysically complex, the interpretation of the term is not sensitive to such complexity. Since the values are not determined by semantic composition of constituent elements, such terms are rigid. Even if some singular terms are directly referential the fundamental property is the *requirement* that such expressions take only individuals as values. In (8), for instance, there is a requirement that the same individual be the satisfier of both occurrences of “x.” Variables contain no compositional structure that would allow them to individuate different individuals at different possible worlds either, so they behave rigidly, though they do not designate, or refer, since the truth conditions of sentences containing them are insensitive to *which* element of a domain they pick out.

The account we defend is hence as follows. We give up the idea that a unified account of the semantic contribution of names can be given in terms of their contributions to propositions, in the sense that all sentences containing names (in subject position) express singular propositions. Instead, I claim that the unifying feature of all singular terms, including names, is to be found at the level of the *requirement* that only individuals – extensions of names – matter to the truth-conditions of sentences containing names. This requirement governs the interpretation of all singular expressions. The

requirement, however, is non-committal with respect to whether a particular *occurrence* of a singular term in a sentence should be interpreted as directly referential – as corresponding to an individual constant – or whether the occurrence should be interpreted as insensitive to *which* individual would make the sentence true – as corresponding to an individual variable.

A proper interpretation of (8) would of course require that both occurrences of “*F*” received the same interpretation as well. In that sense, we could, perhaps somewhat misleadingly, say that predicates are (trivially) *de jure* rigid. The contents of predicates are strictly speaking assigned as well. However, the content of a predicate is a property, or a function from individuals to truth-values, and although which property (or function) a predicate expresses is strictly speaking determined by convention or stipulation it is not a stipulated matter which individuals have the property (or which individuals the function takes to truth). Although predicates determine extensions at every circumstance of evaluation, extensions are not the semantic contents of predicates. What makes variables and constants special is that their contents are insensitive to anything but their extensions.

1.4. A Concise History of Rigidity

The historical source of the idea that singular terms are rigid designators is the development of the semantics of modal logic. Insofar as quantifiers are interpreted objectually – which is needed to avoid modal paradoxes – constants and variables of first-order logic are prime examples of rigid designators in virtue of a requirement on what types of values these terms can have. What *Naming and Necessity* establishes is the

further, empirical fact that proper names of ordinary language are rigid designators as well. Since there are different ways of ensuring that names are rigid designators, it goes beyond what is established in *Naming and Necessity* to claim that proper names should always be treated as logical constants in formal representations of natural language sentences or utterances. Yet the semantics of modal logic reflects the truth conditions of modal sentences in natural language, and the fact that the *de jure* rigidity of individual terms in modal logic emerges as a consequence of measures needed to avoid paradoxical results is not irrelevant for how we understand the natural language expressions the regimentation of which seemed to give rise to the paradoxes.

1.4.1. The Need for Objectual Quantification

The need for rigidly designating expressions had been recognized long before *Naming and Necessity*, and rigidity was introduced (if not under that name) in connection with the development of axiom systems for quantified modal logic. Conceptual or substitutional interpretations of quantified modal logic championed by Carnap (1947) ran afoul of several paradoxical or at least highly counterintuitive results¹¹, such as Quine's challenges (Quine 1943, 1947, 1956), the most important of which being the argument that the principle of substitution does not hold: If a formula Φ contains a term α , and if " $\alpha = \beta$ " is true, substituting β for α in Φ should not affect the truth-value of Φ . Yet if Φ is $\Box(8 = 8)$, which is true, then since "the number of planets" is substitutable for "8" –

¹¹ See Stanley (1997) for a compact but excellent overview over the history; see also Soames (1995) for a similarly insightful overview over the history of the notion of rigidity.

they are extensionally equivalent – $\Box(8 = \text{the number of planets})$ should be true as well.

But the latter is false on its most natural reading.

Furthermore, not only is *substitution* into modal contexts incoherent, according to Quine; *quantification* into modal contexts is incoherent as well. If substitutional quantification is assumed, a formula $\Box\Phi$ attributes a property (necessary truth) to a sentence. Now, Carnap famously attempted to reduce necessity to analyticity. On Carnap's view $\Box\Phi$ says that Φ is analytic. This would make existential introduction, the inference from $\Phi(\alpha)$ to $\exists x\Box\Phi(x/\alpha)$ tantamount to quantifying into quotational contexts. Not only is quantification into quotational contexts rather dubious in itself¹²; the example with the number of planets suggests that *how* an object is named (“8” or “the number of planets”), and not only the object named, matters to whether it satisfies an open formula of modal logic. The upshot is that there was no notion of satisfaction available in terms of which the truth of $\Box\Phi(x)$ could be defined.

Church (1943) took variables in modal languages to range over individual concepts, functions from possible worlds to extensions, and substitution to be allowed for terms that denote the same individual concept (rather than concepts with the same extension). Roughly, a formula containing free variables does not determine a proposition under an assignment of values to the variables alone; to associate a proposition with the formula one also has to assign senses or concept functions to the variables. This approach may block Quine's paradoxes, but it still leads to highly undesirable results. The sentence “ $\exists x\Box(x \text{ numbers the planets})$ ” should intuitively be false insofar as it seems to say that there is an object which is such that it necessarily numbers the planets (the example is

¹² Though see e.g. Kaplan (1986).

from Stanley, 1997). According to Church's approach it is true, however, since in every world the number of planets numbers the planets – the individual concept expressed by “the number of planets” satisfies the open formula “ $\Box(x \text{ numbers the planets})$,” and it is unclear how Church could allow *any* sentence to have the (false) reading just discussed.

The first step toward solving the problems is to recognize that quantification into modal contexts should be understood as objectual rather than substitutional quantification; first-order quantifiers range over objects, not individual concepts (see especially Barcan Marcus, 1961; but also Smullyan, 1947). Objectual quantification is not susceptible to Quine's worries about quantifying into modal contexts, since these worries relied on satisfaction being construed in terms of the truth of closed sentences containing names of the satisfiers. The purported paradox concerning the number of planets, for instance, requires that two closed sentences under the scope of modal operators differ only in the names they contain for the objects (“8” and “the number of planets”). As mentioned, Quine took the alleged failure of substitution to show that whether an object satisfies an open formula (under the scope of a modal operator) depends on how the object is named. But if quantifiers are objectual, an open formula $\Phi(x)$ is satisfied by an assignment A if the object or sequence of objects that A assigns to “ x ” (or a variant) does, in fact, satisfy Φ – irrespective of whether Φ is in the scope of a modal operator, and independent of how and even if we name this object.

But this understanding of quantified modal logic also implies that variables are rigid, and in fact, rigidity (even if it was not called “rigidity”) was very much present in Kripke's semantics for modal logic (Kripke, 1959, 1963). Consider the sentence $\exists x \Box \Phi(x)$. On an objectual interpretation, it is true just in case there is some assignment A

that assigns an object o to “ x ,” where o satisfies the open formula $\Phi(x)$ with respect to every world. Since for every world, what we want to evaluate is whether or not o satisfies $\Phi(x)$ we must ensure that “ x ” designates o with respect to all possible worlds. So on an objectual interpretation of quantified modal logic, variables are rigid (with respect to assignments), and *de jure* since there is nothing further to the semantics of variables over and above the stipulation that (with respect to an assignment that assigns o to “ x ”) “ x ” designates o in every world relative to A .

1.4.2. Quine’s Paradox: Structured Vs. Unstructured Expressions

The move to objectual quantification does not immediately allay Quine’s worries. We are still saddled with the idea that if $\Box(8 = 8)$, then $\Box(8 = \text{the number of planets})$ as well. To solve the problem, we must deny that $8 = \text{the number of planets}$ (or interpret ‘the number of planets’ itself as a term). A plausible answer is that “8” and “the number of planets” (on the false reading of “ $\Box[8 = \text{the number of planets}]$ ”) belong to types of expression that cannot both flank a (first-order) identity sign at the same time, and that “8” and “the number of planets” are intensionally non-equivalent and cannot be substituted *salva veritate* in intensional contexts. This can be accomplished if we treat “the number of planets” as a structured expression, for instance as a quantified expression in the manner of Russell – that is, as something along the lines of $\exists x(Nx \wedge \forall y(Ny \rightarrow y = x))$ – and “8” as a singular term (in which case the *true* reading is obtained by letting the ι -operator take wide scope over the modal operator). “Necessarily, 8 is the number of planets” would thus read $\Box \exists x(Nx \wedge \forall y(Ny \rightarrow y = x) \wedge x = 8)$, which is false.

To see, at least in principle, where Quine's argument goes wrong we must distinguish the principle of substitutivity for (singular) terms (PS) and the principle of substitutivity of materially equivalent sentences (PM) (I follow Neale, 1993):

PS: $\Phi(\alpha), \alpha = \beta \therefore \Phi(\beta)$

PM: $\Phi(p), p \equiv q \therefore \Phi(q)$

These principles are not equivalent. While PM implies PS the converse does not hold. A *non-extensional context* can be defined as a context in which PM does not generally hold. Let us define "intensional context" as any non-extensional context in which PS holds. There may be non-extensional contexts that are not intensional, such as contexts created by propositional attitude operators ("hyperintensional" or "oblique" contexts; they will not matter to us at the moment). Now, given the necessity of identity, PS entails that first-order singular terms satisfy RD (relative to assignments). Quine's paradox can thus be viewed as an argument against the claim that modal contexts are intensional contexts in the sense just defined, or equivalently, as attempting to establish that any context that satisfies PS also satisfies PM.

As pointed out by Smullyan (1947) and especially Neale (1990, ch. 3; 1993) Quine's argument is based on the following inference: "the number of planets is 8; necessarily $8 > 7$; hence, the number of planets is necessarily > 7 ." It goes through if interpreted as " $8 = \text{the number of planets}, \Box(8 > 7), \models \Box(\text{the number of planets} > 7)$." But if "the number of planets" is treated as a Russellian quantified expression and "8" as a constant then we will at most be able to derive $\exists x(Nx \wedge \forall y(Ny \rightarrow y = x) \wedge \Box(x > 7))$,

which says – plausibly – of the number which in fact numbers the planets (in the world of evaluation) that it is in fact necessarily greater than 7. And that is surely true. The false sentence, according to which the number that in fact numbers the planets and is such that it is necessarily greater than 7, necessarily numbers the planets – $\Box \exists x(Nx \wedge \forall y(Ny \rightarrow y = x) \wedge x = 8)$ – cannot be derived from Quine’s premises.

Insofar as the quantifiers are objectual, PS holds in intensional contexts. But to substitute “8” and “the number of planets” in an intensional context one needs PM, not PS. Whereas “8” and “the number of planets” are coextensional with respect to the actual world, they are not necessarily co-extensional. To avoid Quine’s paradox we assume that the predicate components of quantifier expressions do not receive the same extensions at all worlds. Thus the extension of “ N ” varies, and PS will not license the substitution of extensionally equivalent sentences in modal contexts even if it licenses substitution of extensionally equivalent terms. Quine’s claim that contexts that satisfy PM also satisfy PS thus begs the question, and as long as the semantic contribution of the definite description is not just its extension, whereas the contribution of a singular term is just the extension, the purported paradox for quantified modal logic is avoided.

In fact, the response just given is more complex than we strictly speaking need. In “the number of planets is 8,” the description occurs predicationally, and it has been relatively common to treat the sentence as a subject-predicate sentences (following Higginbotham, 1987); that is, $N(8)$. In that case, given that “8” is a (rigidly designating) constant the paradox does not even get off the ground unless the predicate “being a number of planets” necessarily entailed (analytically) “being greater than 7,” which it

does not.¹³ Note that neither solution requires that numerals are actually disguised definite descriptions rather than terms (and should strictly speaking be treated as quantifiers rather than constants) – the crucial point is rejecting the claim that “8 is the number of planets” is a singular identity statement (though it may contain one) and that “8” and “the number of planets” can be substituted *salva veritate* in intensional contexts.

1.4.3. Lessons from the History of Modal Logic

The drive towards recognizing rigid designators originated in technical, logical and semantic developments in modal logic and its model theory.¹⁴ From there it is a non-

¹³ Indeed, we could even lift the assumption that definite descriptions are quantifier phrases. For instance, if definiteness is defined in terms of *familiarity* (Heim, 1983), we could interpret the occurrence of “the *F*” as anaphoric on a previous occurrence (or presupposition) of an indefinite expression introducing an *F* to a discourse, and perhaps even as a variable bound by the antecedent quantifier (as in Groenendijk & Stokhof, 1991). In the latter case “the *F*” would be a term. On familiarity theories definite descriptions take wide scope by default, which makes the rigid reading the primary one. It is, however, eminently possible to let modal operators affect the assignments. Though I leave the details aside for now, since the interpretation of anaphoric expressions is in general achieved by updating assignment functions it would be natural to take modal operators to be expressions that signal assignment updates as well, ensuring the correct, false reading of Quine’s conclusion. The important point is not that definite descriptions are quantifiers rather than terms (see for instance Heim, 1990, for solutions based on directly manipulating world or situation variable); the important point is that the occurrence of a definite description contributes something other than its extension to evaluations of sentences; in the case of the familiarity view just described the crucial feature is that the modal operator signals a reassignment (presumably from the class of *F*s at a given world) to the variable that represents the definite expressions. “8” does not. It is important to notice that despite treating the definite description as a variable, it is, semantically speaking, not *merely* a variable. An anaphoric pronoun “it” does not license reassignment in modal contexts (I argue for this in chapter 5), so the presence of the predicative element (at least) is semantically efficacious in the case at hand.

¹⁴ Of course, the development of rigidity is more complicated than the brief historical outline provided here. Although the need for rigid designators in modal logic was relatively well known prior to Kripke, there was (apparently) little agreement concerning which terms to treat as rigid, especially with respect to modeling natural language claims. Treating names simply as constants would apparently have the effect of making them what Russell would have termed “logically proper names” (Russell, 1905; 1910–11), and the possibility of empty names remained a major worry, at least to those who to a lesser or greater extent shared Russell’s epistemological concerns. There was, however, a recognized need for complete symbols for individual constants (if they were Russellian definite descriptions they would be incomplete), where empty names were given, in a Fregean spirit, arbitrary but specific referents (see for instance Kalish & Montague, 1964, especially p.242–244). A further twist was introduced by Donnellan’s famous distinction between attributive and referential uses of descriptions (Donnellan, 1966) and (what amounts to) the distinction between *de re* and *de dicto* necessities. Hughes & Cresswell (1968), for instance, recognized the necessity of certain identity statements but also accepted contingent (*de dicto*) identities. Thomason &

trivial step to *Naming and Necessity*'s empirical thesis that proper names in natural language are rigid designators. In fact, Kripke's thesis was even stronger, namely that ordinary proper names at least roughly correspond to individual constants, as shown in his response to Dummett's suggestion that the rigidity of names reduces to questions about the scope of descriptions (Dummett, 1973) in the preface to the 1980 edition: "The intuition is about the truth conditions, in counterfactual situations, of (the proposition expressed by) a *simple* sentence [i.e. a sentence not containing modal operators]" (Kripke, 1971[1980], p.12). Ensuring that a name has the same value when changing the circumstances of evaluation requires that it functions like a constant rather than a variable, and suggests direct reference.

Nevertheless, the hypothesis that ordinary proper names in natural language are rigid is not based on a speculative argument by analogy from the semantics of modal logic. The explanatory direction is the other way around; natural language sentences such as "the number of planets is necessarily 8" are intuitively false, and the rigidity of individual terms allows us to facilitate the correct truth conditions in a formal system, thus providing (inductive) evidence for the hypothesis that names are rigid. Although the step from observations about developments in formal modal logic to natural language is non-trivial, the semantics of modal logic provides not just a theoretical framework for

Stalnaker (1968) incorporated Donnellan's distinction in their modal logic by treating certain constants as "substance terms" (equivalent to "rigid designators") for names, the referents of which were determined by a referential description, but they rejected the idea that names in general could be treated in this manner. Kaplan may have come closest to Kripke's ideas in his search for designators the reference of which was "freed from empirical vicissitudes" (Kaplan, 1968, p.190). The rigidity of ordinary proper was suggested by Barcan Marcus (1961), but she offered no argument in support of this claim (see Soames, 1995). Montague (1973) also treated proper names as (intensionally) rigid designators, but did so by stipulation for reasons relating to his treatment of semantic types; proper names are higher-order functions – that is, quantifier denotations – associated with particular meaning postulates. An obvious obstacle to rigidity was the confused idea that rigidity was a hopelessly obscure essentialist notion that implied a type of Platonism anathematic to empiricism, a worry not fully laid to rest until the advent of direct reference theories.

representing the logical form of natural language sentences, but data for semantic hypotheses concerning natural language. The putative paradoxes arise not because of inconsistencies in the formal systems, but because the systems fail to model inferences we naturally make using natural language. Insofar as formal solutions to the paradoxes suggest a firm distinction between quantifier phrases and singular terms, it is likely that it does so because it reflects a natural language distinction.

As we have seen, however, even though rigidity per se is inextricably tied to modality, the *source* of rigidity ought to be found in more basic, non-modal facts about proper names in natural language. Singular term rigidity is already motivated by the adoption of objectual quantifiers, not through separate assumptions concerning formal representations of modal claims. If names work like constants in formal regimentations this representation reflects important semantic properties of proper names in natural language. Nonetheless, we should be careful about what we take the analogies between singular terms and constants to show. It may be a category mistake to say that names *are* variables or constants. What we claim to show is rather that to capture at least the first-order truth-conditions of sentences containing singular terms, we must regiment such terms using constants and variables. The language of first-order logic is artificial, and it is certainly not clear that a first-order regimentation captures all there is to the nature of singular terms or the semantics of sentences containing them (the dynamics of *discourse* being a feature obviously lost in a first-order representation). That such languages are tailored toward capturing distinctions we make in natural language nevertheless suggest that distinctions in first-order languages track significant distinctions in natural language.

Even if there may be more to singular terms, the comparison with first-order languages help us illuminate the extent and nature of central distinctions in natural language.

1.5 Conclusion

Names are *de jure* rigid designators, and this fact needs a (non-modal) explanation. Direct reference is an obvious candidate. However, insofar as direct reference is understood simply as something that is achieved by the satisfaction of a disquotation schema of the form “‘*n*’ refers to *n*,” we lose important insights about the relationship between names, variables, and objectual quantification. The purported paradoxes of modal logic suggest that quantifier expressions in natural language are objectual, and that there is a semantically crucial distinction between semantically complex expressions, such as definite descriptions, and individual terms, such as names. The former individuate denotations or extensions in terms of what roles they play (or predicates they fall under); the latter are used to identify the individuals themselves, independently of their roles or properties. That an expression is *de jure* rigid means that it is an expression used to pick out an individual independently of its role or properties. Names perform this function; so do variables. Variables do not *refer*, however, so *de jure* rigidity is not a matter of referring; variables are used to pick out individuals, but are insensitive to *which* particular individual is assigned – understanding a sentence containing a bound variable (for instance an anaphoric pronoun) does not require identifying the individual that satisfies it. A unified account of names in terms of a purported unified contribution to propositions can hence not be extended to a unified account of singular terms, since

singular terms do *not*, as a category, make unified contributions. I propose instead that the fundamental, unifying property of singular terms, names and variables, is the *requirement* that they take only individuals as values. In Chapter 2 I outline a framework for interpreting singular terms in light of this requirement, and which hence encompasses both expressions that work as constants and expressions that work as variables.

This referential requirement is not sufficient to ensure that terms satisfy Kripke's semantic argument. Insofar as names generally *do* satisfy the semantic argument something more is needed. I will later (chapter 3) argue that in addition to the referential requirement, name occurrences often inherit their values from demonstratives and have their referents determined pre-semantically. This is not a necessary condition on names, however; there are names occurrences that function differently, and since names make different contributions to propositions in different contexts, it is a mistake to assume that uniformity of proposition type is the constraint by which to develop a systematic semantic theory even of proper names only. By shifting the unifying feature to the level of requirements it is possible to provide a framework for the interpretation of *all* singular terms, including pronouns. The cost of the view is the rejection of the idea that such terms uniformly contribute a referent to a proposition. But the latter idea that is clearly unsustainable as a general theory of singular terms in any case if anaphoric pronouns are singular terms, which they are.

CHAPTER 2: SEMANTIC SIMPLICITY

In this chapter I argue that names are “semantically simple” expressions that are rigid in virtue of having (unstructured) semantic contents, defend this claim against competing views that treat names as determiner phrases, and provide the foundations for a unified theory of singular terms. By “semantically simple” I mean that an interpretation requires assigning an individual to the term rather than a complex compositional structure or complex intension – finding the semantic value of the expression is not a matter of composing the semantic elements of its parts (if any). Now, there are arguably other types of expressions that are simple as well, such as simple predicates, but singular terms are distinguished by

- i) Taking individuals as semantic values (or, more generally, their semantic values are dependent on nothing but individuals) and no descriptive content.
- ii) Being complete noun phrases.

The role of singular terms is to attach an utterance to items in the world, independently of the properties these items have or what roles or relations they enter into.

It does not follow from these conditions alone that names are directly referential; variables are not. Not all expressions belonging to the category of singular terms refer, though it *is* required that the satisfiers are individuals (and only the individuals). I suggest that the characteristic feature is the *requirement* that the referent or satisfier is (only) an

individual.¹⁵ Now, Millian theories tend to take all names to be (equivalent to first-order) constants. This, I claim, is a mistake. The universal theory of singular terms is the theory that associates them with semantically simple individual terms; whether occurrences of such expressions should be interpreted in a manner corresponding to variables or constants must be determined by context. Giving up the Millian assumption that names are always constants does not mean that there is no unifying feature all singular terms share. Rather, we locate the universal traits, *de jure* rigidity and non-descriptiveness, at the level that actually distinguishes singular terms from other types of expressions such as predicates and quantifiers. By encompassing both occurrences that function as constants and occurrences that function as variables our criteria will, as opposed to Millian views, allow us to provide a unified theory of singular terms.

To establish that singular terms function either as individual constants or as individual variables depending on context (I discuss exactly what I take this to mean in the next chapter) I need to establish the following:

I) Variables and constants of first-order logic have no descriptive content, and are semantically simple. For the analogy between singular terms and individual terms of first-order logic to hold, singular terms should be semantically simple, complete noun phrases as well. This is the goal of the present chapter.

¹⁵ Of course, semantic considerations place no restrictions on what an individual might be, or what quantifiers can range over. That an individual is “simple” or “unstructured” is meant to be a matter of how it is picked out by a linguistic expression (by assignment), not a metaphysical claim about such individuals. If we can refer to or (first-order) quantify over complex properties, for instance, then complex properties are unstructured individuals for this purpose since the semantic values of the expressions used to pick them out are insensitive to the (metaphysical) structure.

II) I take it to be the case that, although names can function as both constants and variables, they tend to function as the former. Thus, I need to establish how and why names function as constants (in most cases); that is, that the value (referent) of a proper name is determined in advance of evaluating sentences in which the name occurs at a circumstance of evaluation. This is the goal of Chapter 3.

III) In Chapter 4, show that names cannot always be interpreted as constants. I show that the distinction between various occurrences of names corresponds (with some constraints) to the distinction between deictic and anaphoric occurrences of pronouns. Indeed, we may say that names *are* a special type of pronouns.

IV) Chapter 5 argues that anaphoric pronouns are semantically simple. Although the semantic value of a constant is the individual it refers to (or a constant function from worlds to individuals), the value of a *bound* variable is more mysterious – yet if noun phrases in natural language can function as variables they should (arguably) have semantic values as well, in a manner consistent with our criteria.

V) By investigating names in various contexts we will also formulate certain criteria for when they function in what respect. This is the goal of Chapter 6.

V) Finally, in Chapter 7, I describe applications of the view, show that it is discuss main alternatives, and solve outstanding questions surrounding the semantics of names in a more elegant manner than most alternatives.

In section 2.1 I draw some preliminary distinctions. In 2.2 I respond to views that treat names as complex noun phrases, and in 2.3 I tie up some loose ends from 2.2. In 2.4 I describe the positive view and its implications.

2.1. Reference and Denotation

An expression is semantically simple if its meaning cannot be distributed among its parts. “Some fish” is not simple; its meaning is composed, in a particular manner, of the meaning of “some” and the meaning of “fish.” Although “Alexius Meinong” contains the elements “Alexius” and “Meinong” it is semantically simple since the meaning of “Alexius Meinong” is not a matter of composing the meaning of “Alexius” and the meaning of “Meinong.” The role of a *noun phrase* is to introduce or activate the individual or individuals that are the topic of the sentence. Noun phrases can do this either by specifying a role or property and use a determiner to express the portion of individuals filling the role or having the property, or by using a singular term to pick out an individual independently of the role the individual plays. This distinction is crucial. A noun phrase is simple iff it picks out individuals directly, and a semantically complex quantifier if it picks it out in virtue of a role.

2.1.1 Two Kinds of Designation

Names designate individuals. Other expressions can designate individuals as well. For instance, the predicate “first president of the US” designates George Washington.

Predicates designate individuals by expressing or describing a property or role an individual may play. The difference is captured by the distinction between *referring* and *denoting*. Whereas *referring* is a relation between terms and designata that obtains in virtue of convention or stipulation, *denoting* is a relation between an expression *E* and the object that satisfies the compositionally determined conditions expressed by *E*.

To illustrate the distinction Russell famously pointed out that

The author of Waverley is not merely conventionally a name for Scott; the element of mere convention belongs here to the separate words *the* and *author* and *of* and *Waverley* [...] A man’s name is what he is called, but however much Scott had been called [“the author of *Waverley*”], that would not have made him be the author; it was necessary for him actually to write *Waverley*, which was an act having nothing to do with names. [Russell, 1911, p.123]

To determine what “the author of *Waverley*” designates one must compare the (in this case) empirical condition compositionally determined by the semantic values of the various parts to the world. This is not necessary for “Scott.” A related example is:

To wonder what number is named by the German *die Zahl der Planeten* may betray astronomical ignorance, but to wonder what number is named by the German *Neun* can only indicate linguistic incompetence [Kaplan, 1968, p.194]

It is entirely possible to understand what “die Zahl der Planeten” means without knowing what it designates, but not what “neun” means without knowing what it designates.

Now, by varying the scope of the modal operator in “the president of the United States could have failed to be the president of the United States,” we can obtain a reading

on which it is true (the *de re* reading, on which the modal operator takes narrow scope) and one on which it is false (the *de dicto* reading, on which it takes wide scope). It is important to distinguish between *de re* and *de dicto* readings, on the one hand, and *attributive* and *referential* readings of descriptions, on the other. A *de re* reading is not a referential reading. It is eminently possible to fully grasp the proposition expressed by the *de re* reading of “the president of the US could have failed to be president of the US” (the true reading) without having the faintest clue who the president is. The contribution a *de re* description makes to a proposition is *not* the denotation, but a complex functional entity. A *referentially used* description is a description used as a singular term. A referentially used description contains no semantic requirement that the referent actually satisfies the description,¹⁶ since it contains no semantic parts that semantically express such a condition. A description used *de re*, on the other hand, *denotes* the individual that satisfies its predicative material. Hence a *de re* rigid description can only be *de facto* rigid, whereas a referentially used description is *de jure* rigid.

2.1.2 Reference and Structure?

A referential use of a description is accomplished by disregarding its semantic structure and use it as if it had a simple, non-compositional semantic value. One removes the information contained in description and use it as an informationless tag. The opposite move, to use a *name* as if it had semantic structure, seems impossible (at least without shifting grammatical category). This observation suggests that there is a close connection

¹⁶ There may be a pragmatic or presemantic requirement (e.g. a requirement at the level of character), but this is a different matter.

between reference and semantic structure (see Neale, 1993, 2009; Dever, 2001; Sullivan, 2007, 2009).

The connection may be something along the lines of the following. It is possible to treat *all* simple expressions of a language as *de jure* rigid designators of their semantic values, or perhaps the properties (or properties of properties) that determine those values (by “semantic value” I mean what it contributes to a proposition). A predicate is in that case a *de jure* rigid designator of a characteristic function from individuals to truth values (namely the one determined by its lexical meaning – the function is defined for all worlds, although what individuals it takes to *truth* will of course vary); a determiner is a *de jure* rigid designator of a particular higher-order function (a function from individuals to truth values, to a function from a function from individuals to truth values, to truth values, which will yield different quantifier phrases for different ordered pairs of predicates – or their corresponding values – as inputs). Since the particular semantic value of an expression is insensitive to other expressions in a sentence (though given the particular semantic values of other expressions the semantic values can of course be *input* to the values of more complex expressions), and insensitive to world of evaluation, they are *de jure* rigid.

Now, the role of a *noun phrase* is to pick out one or more individuals. It can perform this role in one of two ways. A *quantifier phrase* – which is, by nature,¹⁷ not simple – expresses a partition on a set of individuals, the extension of a predicate or predicates. The *lexical meaning* of a predicate is a *role* (a property or relation). So a

¹⁷ I leave aside type ⟨1⟩ quantifiers here, primarily (and I hope I am excused) because they raise too many questions; it is for instance not entirely implausible to think of singular terms themselves as constituting a special kind of type ⟨1⟩ quantifiers here (cf. Montague, 1973).

quantifier phrase picks out individuals in virtue of the roles they play; that is, a predicate describes a role, and the determiner expresses an operation on that role by relating a given portion of the role-players to a different predicate regardless of who those role-players may be at any given world of evaluation. A *singular term*, on the other hand, does not and cannot pick out individuals in this manner. Since singular terms are semantically simple, complete noun phrases – and therefore *de jure* rigid designators of their values – they do not combine with predicates to individuate objects in virtue of the roles those individuals play. Insofar as the role of noun phrases is to pick out individuals, singular terms are referential – their values must be individuals, not roles. By equating *de jure* rigidity and semantic simplicity, and assuming that the role of noun phrases is to pick out individuals, singular terms are *de jure* rigid in virtue of being semantically simple.

Of course, an alternative is to claim that singular terms really *do* pick out roles (that is, the name “Jack” picks out the role of *being Jack*), and perhaps combine with (null) determiners to form determiner phrases. I discuss this option in 2.2. Certain Fregean views may also claim that the value that the name rigidly designates, according to the understanding of *de jure* rigidity just given, is something other than an individual. On the face of it a view that allowed a name “Jack” to pick out something not reflected in the *syntax* of “Jack” (i.e. either the individual Jack, on referential views, or the role *being Jack*, according to metalinguistic descriptivism) may in this context look like a violation of an immensely plausible disquotational link between a word and a semantic value. Things are admittedly more complex, and I briefly take up Fregean views in Chapter 3.

Even if the role of semantically simple noun phrases is to pick out individuals, it does, not follow that any semantically complex expression is non-referential, though it is

a plausible hypothesis.¹⁸ Nor does the suggestion imply any particular way of distinguishing *semantic reference* from *speaker reference*. A definite description, for instance, is used referentially only if it is used as a semantically simple expression; that is, used in a manner that avoids treating the elements of the description as expressing a complex, descriptive condition that the individual one aims to individuate must satisfy. Whether an expression is semantically simple or not is a matter of whether it has a semantically simple value, regardless of whether receiving that value is a matter of speaker reference or semantic reference.

It is important to emphasize (following Neale, 2009) that the claim that referential expressions are semantically simple is an empirical claim about natural languages. It is no problem to develop formal languages that allow structured referential terms, but natural languages do not work that way. The reason may be that the operations of natural language reflect certain cognitive needs (Føllesdal, 1986; Neale, 1993), such as the need to recognize (and talk about) individuals across changes in roles, and also to recognize roles independently of the individuals (if any) who fill them, or to talk about individuals *in virtue* of the roles they fill. Modal and temporal contexts provide evidence for such distinctions. We need a linguistic means of tracking Obama through the various roles he plays, but also a way of tracking the role of *being president of the United States*, regardless of who actually fills it. (Note that explaining the distinction between referential terms and quantifier phrases by cognitive needs does not mean that semantic distinctions reduce to cognitive distinctions; rather, cognitive distinctions we rely on explain *why* we had to develop the linguistic distinction.)

¹⁸ See Salmon (2006) and Neale (2007). Complex demonstratives are, in particular, a source of contention.

Notably, we can track individuals without making the proposition dependent on a particular individual, as when we use descriptions *de re*. For instance, in “the president is tall. He could have been short,” “he” designates the individual who is president, and tracks him across counterfactual scenarios. “he” is a singular term, however, and does not track the *role* but the individual, if any, who happens to be president.¹⁹ But “he” does not directly *refer* to that person. If the fragment is evaluated at different circumstances of evaluation “the president” denotes a different individual without anything in the fragment changing its meaning (and even though the sentences expresses the same proposition). “he” would in that case track a different individual. Tracking individuals rather than roles is not a matter of referring to a specific individual, but of using a term that has the characteristic property of singular terms, namely to track individuals, non-descriptively and independently of the roles they play.

Note also that a *definite description*, even if used *de re*, contains no such singular term; it is the pronoun anaphoric on the description that is a singular term. The claim that singular terms are semantically simple does not mean that their contributions can be determined in isolation of their linguistic surroundings. To interpret the sentence just mentioned one also needs to know the antecedent, “the president.” That does not mean that the semantic value of “the president” is part of the semantic value of “he” – the antecedent provides a restriction on where the individual can be found, but is not equivalent to it (indeed, insofar as “he” is a bound variable, “he” does not designate, determine or denote a particular individual at all, but *ranges* over the domain determined by the antecedent and world of evaluation).

¹⁹ I defend this claim in much more detail in Chapter 5.

2.2. Meta-Linguistic Descriptivism and the Syntax of Names

That singular terms are semantically simple, referential expressions is a non-trivial claim. According to *metalinguistic descriptivism* names are, in fact, complex, covert determiner phrases that designate individuals in virtue of their roles, in particular their roles of being the bearer of a particular name. Now, many names are *syntactically* complex. “Alexius Meinong” consists of “Alexius” and “Meinong,” and the name “The Roman Empire” consists of a complex noun phrase. I will, however, claim that if these are *de jure* rigid, referring expressions, they have simple intensions or individuals as their semantic values; though morphologically complex they do form grammatical as well as semantic units (or at least that when we use “The Roman Empire” referentially we use it as a semantic unit), and that there is independent evidence for this. For the purposes of this section, I follow Sullivan (2009) and distinguish between *atomic* and *molecular* expression types on the one hand, and *unstructured* and *structured* expressions on the other. Molecularity does not entail (semantic) structure. While “The Golden Gate Bridge” is (syntactically) a molecular expression by having proper morphological parts, and is – apart from the capital letters – morphologically indistinguishable from a description “the golden gate bridge” there are reasons to carve up grammar in a way that relegates “The Golden Gate Bridge” to the class of semantically unstructured terms. Evidence comes from what we need to do to interpret it. “The Golden Gate Bridge” and “the golden gate bridge” are both noun phrases, but the first is also a proper noun whereas the second is not; the latter is a construct that *contains* one proper noun, “bridge,” adjectives and a determiner.

2.2.1. The Syntax of Proper Names

Insofar as (*de jure*) rigidity is primarily a function of model-theoretic rules associated with names as a grammatical category, we must say something about syntax.

Syntactically, proper names are (or occur as) noun phrases. In natural language nouns (generally) designate roles that individual elements may or may not satisfy. For instance “bachelor” designates a role filled (throughout history) by a variety of people ranging from Kant to the pope. Proper names are noun phrases that refer to individuals that fill such roles. These can consist of a single element, e.g. “Anaximander,” or of two or more elements in the manner that “Alexius” and “Meinong” combine to make up the noun phrase “Alexius Meinong.” Names can also be built up from other linguistic items, such as “Radisson Hotel,” “Lake Eyre” and presumably “Much Ado About Nothing,” but still function as ordinary, referential names. *General nouns* are generally not used in isolation, but combine with determiners to form determiner phrases such as “all bachelors” or “the tyrant.” Such expressions are not referential. On the other hand, even names can contain determiners, such as “The Mighty Handful” and arguably “Sibelius’s Fifth Symphony,” and even occur in plural constructions, such as the family name “the Proudfoots” (which is, interestingly, not “the Proudfeet”). I argue that although it is possible to treat these, or some of these, as definite descriptions, when and if they function as names the semantics are not determined *in part* by the elements they contain. Either they are names and referential, or they are quantifier phrases and not names.

One source of controversy with respect to the syntax and semantics of names is the fact that names sometimes appear in unusual grammatical positions, e.g. verbs (“to

Quine a phenomenon”), adjectives (“a Manstein-Guderian strategy”) or as regular nouns (“we need a new Talleyrand”). In the latter cases they even take determiners (“there’s a Eustace to see you,” “Tate acquired another Mondrian,” “there are ten Xaviers in the telephone catalogue”). Indeed, Burge (1973) takes their occurrences as regular nouns to be primary. On his view an occurrence of a name “Eustace” can be replaced by the nominal expression “person called Eustace,” and Alexius von Meinong is on Burge’s view actually a Meinong; that is, a member of the class of Meinongs (there are presumably other Meinongs as well). What fills the argument place for a name when it occurs free (the most common occurrence in ordinary discourse) is something like a bare demonstrative that gets a denotation in a context. On at least one plausible understanding of Burge’s view names are complex referring expressions.

Assigning a semantic role to the predicate seems correct in cases where names occur in determiner constructions, but by itself the existence of such constructions may not be a counterexample to our view – “Meinong” occurring in the construction “a Meinong” would on our view not belong to the same category as ordinary, free occurrences of the name. Metalinguistic descriptivism, however, is motivated by the idea that ordinary names and determiner constructions involving them should receive a uniform semantic treatment. I will argue that metalinguistic descriptivism is mistaken.

2.2.2. Metalinguistic Descriptivism: Nominal Description Theory

Metalinguistic descriptivism has been defended by Kneale (1962), Loar (1976), Larson & Segal (1995), Bach (2002), Elbourne (2005), Matushansky (2006), and others. To see the

motivation, notice that linguists sometimes analyze names as noun phrases with possibly empty determiner positions. Thus, a name like “Alexius” has the structure:

(1) [DP[DET \emptyset] [N Alexius]]

It seems plausible to interpret the null element as the determiner “the.” Thus, on Bach’s Nominal Description Theory (NDR) the name “Alexius” is taken to be interpreted as “bearer of ‘Alexius’” (Bach, 2002). The class of proper names is different from the class of ordinary nouns, however. According to Bach, learning to use “table” properly requires knowing that “table” expresses the property of being a table, whereas learning to use “Alexius” properly is just a question of being able to spell or pronounce it and know that it is a name *used to refer* by pragmatic mechanisms. As such, knowledge about particular bearers of the name is not part linguistics proper. Although Bach’s analysis may appear to run into the problem that individuals called “Alexius” are generally not *necessarily* called “Alexius,” and thus run afoul of the modal argument, this problem may be avoidable. According to Bach, “Alexius is Alexius” has two readings, just as “the person called ‘Alexius’ is the person called ‘Alexius’” has two readings as a result of the interaction between the determiners, and the modal profile question can be solved by interactions between determiner and modal operators.

As Kripke (1972 [1980], p.69) points out, a metalinguistic description can hardly be the *sense* of a name in anything like the Fregean way. The reason is not that the description is trivial – it is not – but that it is hardly helpful in answering the question of

who “Alexius” refers to by being told that it is the individual called “Alexius.”²⁰

Furthermore, if a sense represents cognitive significance and also determines reference, we get the wrong referent at various circumstances of evaluation insofar as there are worlds in which “Alexius” denotes someone else (i.e. it does not satisfy the *semantic* argument). The metalinguistic description, unless rigidified, does not give us a constant function. But even if it is rigidified we are not out of the water. The descriptions in question are usually improper. There is more than one satisfier of “the person called ‘Alexius.’” With several individuals called “Alexius,” further information is needed to single out one particular denotation in a context. Coming up with a way of providing such information is not a trivial matter if the intended denotation is not immediately available.²¹ The information required to provide an adequate restriction would also be susceptible to the standard Kripkean arguments against descriptivism outlined in Chapter 1. To illustrate the problem, consider what it means to be the bearer of, say, “Thales.” “The bearer of ‘Thales’” is an improper description (there is, for instance, a major electronic systems company called “Thales” as well). We want the philosopher who (allegedly) thought everything was water. But the information we associate with Thales may very well be largely mistaken – he may not have thought that everything was water – so we cannot specify with precision how to make the description proper by restricting the domain, and if we cannot do so, then the NDR analysis would be susceptible to Kripke’s arguments.

²⁰ But see McDowell, 1977, and Kroon, 1989.

²¹ On Bach’s view of definite descriptions (Bach, 1999) sentences such as “Beethoven composed Fidelio” would – given that “the person called ‘Beethoven’” is improper – be literally truth-valueless even though pragmatic mechanisms still makes them interpretable in a context.

The worry, then, is how we ensure that “the bearer of ‘Thales’” has the correct denotation. Now, all theories of proper names have to answer this question, but NDR appears to restrict the range of potential answers too much. It may for instance be tempting to interpret “Thales” as “the bearer_x of ‘Thales,’” where “x” is assigned a name-using tradition in a context. But it cannot be a requirement that one can identify the tradition descriptively to understand the name (that would presumably quickly force one to take *x* as “the tradition linking ‘Thales’ to Thales” and hence a regress). Rather, “x” must be an indexical element that refers to a tradition the current speaker need not know in detail. But in that case, why not take “Thales” itself to be the indexical and hence directly referential element, in which case “the bearer of ‘Thales’” could perhaps be taken to be a metalinguistic *character* and e.g. a causal relation to Thales be invoked as a contextual element that helps assign a referent to “Thales?” This is not meta-linguistic descriptivism, however.

A related worry, which becomes particularly worrisome in the absence of a clear way of making the descriptions proper, concerns what, exactly, the *bearer of property* is supposed to be so that makes it the case that Thales instantiates it. The bearer of “Thales” cannot be merely the person *called* “Thales” by me, for then our explanation is perfectly circular – why is it *Thales* who is called “Thales” by me rather than, say, Epicurus? But it could not be the person called “Thales” by my linguistic community either, unless the linguistic community had an independent way of identifying Thales (and certainly not the person called “Thales” by his contemporaries; Thales was probably not called “Thales” by his contemporaries); and this independent way of identifying Thales should not just associate Thales with a set of descriptions, for in that case Kripke’s original arguments

against descriptivism would still apply. It does not help to appeal to intentions either (“the person called ‘Thales’ that the speaker has in mind”) for it is hardly easier find a determinate answer to who the speaker has in mind than to specify a restriction for the description.²² What the *bearer-of* relation is supposed to be remains rather mysterious, in particular since Bach must be careful to avoid analyzing the *bearer of*-relation as the *refer* relation if he is to avoid explaining the fact that Thales is the bearer of “Thales” in virtue of “Thales” referring to Thales.²³

2.2.3. Metalinguistic Descriptivism, Indefinites and Names With Free Variables

Matushansky (2006) treats the null element of the determiner in (1) as an indefinite, and argues that the semantic value of “Alexius Meinong” is the intersection of the extensions of “Alexius” and “Meinong.” Since the determiner is indefinite, there is no requirement that this intersection is a singleton, so we avoid the problem of improper descriptions, but to avoid the metalinguistic worries Matushansky proposes that in ordinary uses of names there is a tacit element denoting, in context, a naming convention. Thus, to fully understand a name occurrence it is not enough to know the semantic meaning of “Alexius Meinong;” you must also locate the referent by means of the naming convention.

²² He could not say that the philosopher and the electronic systems company do not share a single name, but rather bear different but phonographically indistinguishable names (as in Kaplan, 1990 – this view will be discussed in the next chapter). Individuating names this way defeats the whole purpose of metalinguistic descriptivism – for instance, “there are three Jacks in the room” would then *not* use “Jack” in a manner semantically related to ordinary uses of “Jack.” Worse, it makes understanding “Thales” dependent on knowing who the designation is (fully understanding a name presumably involves knowing which *name* it is); and it would not be sufficient to know merely that it is a name.

²³ Katz defends a slightly more complex view (Katz, 1990, 1994), according to which “Thales” is interpreted as “the thing which is a bearer of ‘Thales.’” See Bach (2002, p.100, ft.40) for an elaborate criticism of this view. In any case, my criticisms of Bach’s view would seem to apply to Katz’s view as well with minor adjustments. See Burgess (forthcoming) for a different take on metalinguistic descriptivism that interprets the view as not being in obvious opposition to Millianism.

Matushansky's explanation involves an uncomfortable circularity, however, which is most obvious in the case of names that are built up of various grammatical elements. The denotation of "The Golden Gate Bridge"²⁴ is not the intersection of the extensions of "the," "golden," "gate" and "bridge" as these are ordinarily interpreted. The syntactic elements must rather have some special extension in the case of "The Golden Gate Bridge" not determined by the conventional extensions associated with the nouns – anything called "The Golden Gate Bridge" should be a Golden Gate Bridge, but it is no requirement that it is a golden gate bridge. But what, then, are the semantic values of the elements of "The Golden Gate Bridge?" In fact, we cannot *determine* this without already knowing that they are part of a syntactic unit, the name. Now, "Bredon Hill" is allegedly a Bredon that is also a hill, so the common noun "hill" seems semantically active. But instead of helping Matushansky's case the example undermines it. "Bredon Hill" is also a poem, and since we cannot determine whether we are talking about a hill or not before we know what "Bredon Hill" refers to in a context, the hill or the poem, we cannot determine the value of the *elements* of "Bredon Hill," including "Hill," unless we know what "Bredon Hill" refers to in the context; so, reference assignment must be prior to composition. Since we need to know the whole to determine the semantics of the parts, "The Golden Gate Bridge" and "Bredon Hill" seem to be semantically unstructured; their meanings cannot be distributed among the proper parts of the name. Hence, "Alexius Meinong" is a semantic unit rather than an expression that determines a semantic value by composing its constituent atoms. Hence names have no proper semantic parts.²⁵

²⁴ The example belongs to Cumming (unpublished, p.27).

²⁵ A similar argument applies to Corazza's idea that what he calls "descriptive names" refers in virtue of being (for practical purposes) *dthat*-constructions (Corazza, 2002); "The University of Miami" should be

There is also a worry that the interpretation of any occurrence of a name must be relativized to a particular name using tradition. Insofar as the tradition is individuated by an indexical element, why does the name-using tradition not do all the work of interpreting the name occurrences? There is no obvious reason why the name cannot be identified with the index that keeps track of this tradition instead of with the indefinite description. This worry is even more obvious when applied to Elbourne's (2005) account. In order to avoid the problem of uniqueness Elbourne suggests a move reminiscent of Burge's. Elbourne argues that an occurrence of "Alexius" can be interpreted as "the individual called 'Alexius' and identical to x ," where " x " is a (free) variable to which Alexius is assigned in a context. But what determines which *variable* is associated with a particular name? Presumably, the *name* occurring in a sentence must do the trick. Thus "Alexius" should e.g. be interpreted using the variable x , and an occurrence of "Aristotle" be interpreted using a different variable, for instance y . But insofar as the values of " x " and " y " are the individuals Alexius and Aristotle, respectively, what are the first parts of Elbourne's analyses, "the individual called 'Alexius'" and "the individual called 'Aristotle,'" supposed to contribute? As long as the variable associated with the name n must be different from any variable associated with a different name the first part of the description appears to play no semantic role whatsoever. What "Alexius" contributes to the evaluation of a sentence in which it occurs is the value of the variable " x " and, it

analyzed as the referential (but semantically compositional) *dthat*[the University of Miami]. If this were correct, however, then since many descriptive names simply do not satisfy their associated descriptions and listeners would have to understand that the elements are parts of a *name* before interpreting them, it seems more plausible to relegate the descriptive elements to meta-semantics, and thus something that is conveyed or help listeners individuate the name-using traditions, rather than as genuinely semantic elements. I discuss partially descriptive names below.

seems, little else, and the descriptions are better understood as doing a metalinguistic rather than semantic job.²⁶

2.3 Deriving New Expressions

If names are structurally simple singular terms we lose the connection between ordinary occurrences of names and irregular occurrences such as “there are ten Xaviers in the phone book” or “the Germans did a Napoleon in Russia.” This is hardly a big loss; irregular occurrences were a red herring to begin with.

2.3.1. Irregular Occurrences

Irregular occurrences of word types are fairly common. Although “to painstake” and “violiny” have no conventional meaning in English, we have a grasp of the intended meaning of “Jack painstakingly worked his way through the dissertation” or “Parry’s scoring is rather violiny.” Our interpretations are perhaps less than fully precise, but drawing on linguistic and contextual resources we are at least able to arrive at the communicative intentions behind the utterances.

Similarly, I suspect readers quickly arrive at a suitable interpretation of “a Soamesian argument” even though “Soamesian” is not a word they have encountered before. “Soamesian” is an adjective, not a name, and the example shows that we are able to devise novel, meaningful expressions on the spot by drawing on general linguistic

²⁶ In fairness, though, Elbourne’s main goal is to provide a unified theory of definite expressions, including pronouns, definite descriptions and names. However, it appears that the free variables in the example mentioned must be interpreted indexically (on pain of regress), and Elbourne’s unified framework does not appear to make room for an explanation of indexicals either.

knowledge. Names can be drawn upon for this purpose as much as other expressions. It is not a big step from this observation to the claim that “a Vlaminck” in “Jack bought a Vlaminck” means a *painting* by Vlaminck. It is not a name at all, but a regular determiner phrase derived from a well-known name by drawing on some kind of linguistic guidelines for noun construction. Similarly, whereas “Hohenzollern” is a name for a family, “a Hohenzollern” is a determiner phrase formed from combining a determiner with a predicate derived from the name, and whereas “Mercedes” refers to a car company (or brand), “a Mercedes” is an indefinite phrase in which “Mercedes” occurs as a common noun rather than a name.

A main motivation for metalinguistic descriptivism is the wish for a unified theory of names, including irregular occurrences. According to Burge

[...] appeal to ‘special’ uses whenever a proper name does not play the role of individual constants is flimsy and theoretically deficient. [Burge, 1973, p.437] [also Bach, 2002, p.77]

But “a Manstein-Guderian strategy,” “Hitler did a Napoleon in Russia” or “I Steven Seagalled him” – all perfectly legitimate if “I know three Eric’s” is – cannot be interpreted the way Burge or Bach proposes either. “Hitler did a the bearer of ‘Napoleon’ in Russia” is not equivalent to “Hitler did a Napoleon in Russia,” and “a [Manstein-Guderian](x) strategy,” where x is a free variable designating a compound of Erich von Manstein and Heinz Guderian, is not equivalent to “a Manstein-Guderian strategy.” So Bach and Burge do not get unified theories either, and must recognize non-standard uses of names. I make a similar appeal to non-standard use for names used as common nouns in determiner

phrases, and see no compelling reason to draw the distinction the way Bach and Burge do, rather than the way I do.

2.3.2 Partially Descriptive Names

Soames has argued that molecular names such as “Mount Everest,” “The University of Miami” or “The Battle of the Ice” have meaningful proper parts even though they are referring names. According to Soames, the parts express certain properties the designation has to fulfill in order for the name to refer to that designation, the descriptive content determines part of the proposition expressed, and the proposition in question is singular but not ‘purely singular.’ Furthermore partially descriptive names are not obstinately rigid; in some possible worlds in which the actual referent exists, the partially descriptive name designates neither the actual referent nor anything else (Soames, 2002, p.120). Notice that meta-linguistic descriptivism would be placed in a similar dilemma regarding these constructions (if there is a dilemma): should an occurrence of “the University of Miami” be interpreted as the description “the University of Miami” or the meta-linguistic description “the bearer of ‘The University of Miami?’” These may be very different beasts, so a meta-linguistic descriptivist would presumably have to rely on distinctions analogous to the ones I am going to make.

Now, “The University of Miami” may be a description. In that case it would have a complex intension determined by the intensions of its parts, and denote what is in fact the university of Miami by virtue of this entity satisfying the conditions compositionally determined by the description. If other places were called “the University of Miami” (there are) we would have to rely on contextual factors to distinguish them. It may even

happen that whereas e.g. “The River Thames” is a name, “The University of Miami” is not, or it may be used as a name and as a description on different occasions. Let us for the sake of argument agree with Soames that it is (used as) a name. On my view it is thus structurally simple. Soames disagrees, in part due to possibly negative answers to questions such as

- i) Could the University of Miami have failed to be a university and still be designated by the name “the University of Miami?”
- ii) Could it be relocated outside Miami and still be the designation of “the University of Miami?”

I do not know the answers to these questions. Thinking that the answer matters is confusing semantic and metaphysical questions. Suppose “the University of Miami” is a structurally simple name. The question of whether it would have continued to refer if the university it refers to were moved is not a semantic question but a metaphysical question about whether the designatum persists; similarly for the question of what would happen if it ceased to be a university. If the designatum ceased to exist it is no surprise that the proper name “The University of Miami” fails to refer to any new entity that may take its place. Hence, “The University of Miami could have failed to be a university” may be false, even necessarily false (since it may be an essential property of the University of Miami that it is a university), but it is not a contradiction *in terms*. Only by assuming that the semantics of names were inextricably tied to the metaphysical properties of their referents would there be a good reason for thinking otherwise. Similarly, in “if Mt.

Everest did not exist, K2 would have been the highest mountain in the world,” “Mt. Everest” refers to Mt. Everest regardless of the fact that Mt. Everest is not a mountain in the world of evaluation. Hence, reference precedes essence also for partially descriptive names, so partially descriptive names are obstinately rigid.

If “The University of Miami” is a semantically simple name, then the fact that the extension of “the University of Miami” is a university in Miami is semantically irrelevant to the fact that “University” and “Miami” occur in the name. If that sounds *prima facie* implausible, I suggest it is because we are often tempted to treat “the University of Miami” as a description, or because there are obvious etymological reasons for the institution receiving that name. Many places called “the University of X” fail to be universities, however. “Patriot University” refers to an institution that cannot reasonably be said to be a university. One might complain that there is a sense in which it is illegitimate for Patriot University to call itself a university, but this is because the name is deceptive in virtue of certain deceptive connotations. The fact that “Patriot University” refers to a determinate entity that we do not think qualifies as a university – “Patriot University is not a university” is strictly true – shows that “Patriot University” refers, and does so independently of any descriptive content.

2.3.3 Names in Language

We have argued that names are structurally simple and that the existence of molecular names does not cast doubt on this conclusion but actually supports it, insofar as we cannot interpret parts of a name without knowing that they are, in fact, parts of names rather than nouns with independently interpretable meanings. Metalinguistic

descriptivism has sometimes been motivated by the idea that names are not really parts of or lexical items in natural languages:

The word 'word' is sometimes used in an extended way that allows one to speak of any and all proper names as words ... Proper names generally [are] not words. [Ziff, 1960, p.86]

Bach points out that

[d]ictionaries are not incomplete for not including [proper names], and your vocabulary is not deficient because of all the proper names you don't know [Bach, 2002, p.82]

I think this point rests on a too Platonistic conception of language, however. Dictionaries do not decide what belongs to a language but *reflect* this, and I doubt that natural languages have determinable, non-arbitrary limits. We make up words of all kinds all the time, and even though most of them do not initiate a tradition this does not mean that such words are not *words*. "To Sutherland Meyerbeer" is a legitimate expression. I am less worried about whether it is legitimate to call it an expression of *English*²⁷, but "to Sutherland" is a verb, and an eminently understandable one, because it behaves like one and has to be assigned a semantic value of the kind that is assigned to transitive verbs.

There is also the worry that

[K]nowing the "meaning" of a name does not require knowing all the individual(s) it belongs to or associating the "right" (substantive) description(s) with it. You just have to recognize it as

²⁷ I should emphasize that this point has nothing to do with the distinction between languages and idiolects as this has been used in the philosophical literature. An idiolect in this sense concerns the *meanings* of words; relativizing the meaning of words to idiolects means relativizing the lexicon of terms to individual users (or groups of users), for instance by way of relativizing senses. This has nothing to do with the question of whether the name itself belongs to the (syntactic) body of expressions that constitute a natural language.

belonging to the category of proper names and know how to apply the NDT schema to it [*ibid.* p.82]

This worry, however, is misplaced. At best it suggests a distinction between character and content for names on the model of demonstratives. Indeed, if one assumes that “Eric” is shared between various bearers then context is needed to supply a value. In that case the character and linguistic meaning, of “Eric” would be “the person called ‘Eric,’” although the semantic contribution of the name would be exactly that of a constant. Understanding “Eric” would thus not require knowing all bearers of that name any more than understanding “he” would require knowing all possible referents of “he.”

2.4 Grammatical Kinds and Semantic Values

In 2.2 we suggested that it would be better to relegate the metalinguistic description “the individual called *n*” to a meta-linguistic level of character. Thus, the contribution of “Alexius” to the proposition expressed by a sentence in which “Alexius” occurs is just Alexius. “The individual called ‘Alexius’” is invoked at most to find a suitable assignment in the course of determining the *truth conditions* of a sentence. The name itself is a complete noun phrase with no proper semantic parts. Variables and constants of first-order logic already behave as semantically simple expressions with individuals as values. We have seen evidence that by using objectual variables and constants we can capture the truth conditions of natural language sentences involving names in a manner that alternative interpretations cannot, which suggests that singular terms in natural language work in a manner that at least roughly correspond to individual terms of first-

order logic. Singular terms pick up individuals and have no descriptive content. Other expressions have semantic values determined differently, but the semantic contributions of other expression types are not the extensions alone. If such expressions are rigid designators of individuals or extensions they are at most *de facto* rigid.

2.4.1 Preliminary Remarks: Compositionality and Semantic Value

Compositionality requires that all elements of a sentence receive semantic values in isolation. Semantic values are subsequently composed to determine the meanings of complex expressions and complete sentences, in accordance with the compositional rules associated with their grammatical categories. In the following “semantic value of *E*” will mean *the contribution E makes to the truth-conditions of (or proposition expressed by) a sentence*. Referring expressions, then, are expressions the semantic value of which depends on the individual referred to. Denoting expressions designate individuals as well, but the individuals are not part of the semantic values. As pointed out above that there is a sense in which all expressions are rigid designators of their *values*; that is, every expression designates a (possibly complex) property distinct from its extension but responsible for assigning the correct semantic value to *E*, and for assigning the correct extension (sets of individuals or *n*-tuples of individuals) to a sentence at a circumstance of evaluation to determine a truth value.

However, although some kind terms, such as “table,” are semantically simple, kind terms rarely if ever occur as noun phrases. Instead, they i) combine with determiners to form compositional determiner phrases such as “the chair” or “every chair;” or ii) occur as part of predicates such as “is green,” in which case their values are functions

from individuals to truth values. Whereas the referents of terms make up a (first-order) universe, predicates denote classes of such referents. When occurring *in* noun phrases (“all green things”) their role is to provide inputs to complex second-order functions, the role of which is to relate members of the extension of a predicate to the extension of some different expression in a manner determined by a determiner.²⁸ The value of the sentence depends on whether the proportion of *F*s determined by *Q* in $[Qx:Fx](Gx)$ is *G*, and the individuals that are in that class will vary across worlds. Hence, such noun phrases are generally non-rigid designators of *extensions*.

To determine the denotation of a definite description (even a rigid description) such as “the queen of Belgium,” one must first, determine the semantic value of “Belgium” and “of” and “queen” and thereby the value of the compositional expression “queen of Belgium;” then add the value of the determiner ‘to obtain the value for the complex expression “the queen of Belgium.” The proper name “Belgium” requires no composition. “Belgium” gets its value assigned “directly.” Furthermore “Belgium” is a *de jure* rigid designator in virtue of having its value determined non-compositionally and certain general facts about natural language interpretation (e.g. an at least *prima facie* isomorphism between syntactic and semantic categories). Hence, “Belgium” is not semantically simple because it is *de jure* rigid; rather, “Belgium” is *de jure* rigid because it is semantically simple, and because the grammatical category to which “Belgium” belongs is characterized by its members receiving their semantic values in this manner.

²⁸ Of course, as described this makes quantifier phrases incomplete expressions, and indeed the quantifiers of standard first-order predicate logic are incomplete. This is avoided in generalized quantifier theory, in which the quantifier phrase involved in “Every *F* is *G*” is interpreted as a type $\langle 1, 1 \rangle$ quantifier.

2.4.2 The Semantics and Grammar of Singular Terms and Predicates

In a standard Tarski semantics the extension of a referring expression is an object, its referent; the extension of an n -place predicate P the set of n -tuples of objects of which P is true; the extension of a sentence a truth value; and the extension of an n -ary sentence connective a function from n -tuples of truth values to truth values. Assuming an extensional language (fragment) the extension of an expression is determined by a single axiom or by extensional composition of the extensions of the elements constituting the expression. The extensions of terms and predicates are given by axioms and the extensions of sentences are functions of the extensions of the various parts constituting the sentence. Terms and predicates play very different roles in determining the extensions of sentences. The extension (I) of an atomic sentence Φ depends on whether the ordered n -tuple of the extensions of the terms in Φ is a member of the extension of the predicate of Φ . Given a predicate P and terms $\alpha_1, \dots, \alpha_n$:

$$(2) I(P(\alpha_1, \dots, \alpha_n) = \text{Truth iff } \langle I(\alpha_1), \dots, I(\alpha_n) \rangle \in I(P)$$

If modal operators are added to the language then whereas the values of terms depend only on whether the axioms – which are insensitive to the world parameter – are satisfied; the extension of a predicate will vary from world to world, and is not essential to its meaning.

In a (basic, static) Montagovian analysis the semantic values of names and sentences are the basic elements by which other expressions are compositionally defined. Singular terms and sentences are *basic* categories. Other categories are *derived*. In

standard notation, terms are expressions of type e and stand (extensionally) for individuals. Sentences are expressions of type t and stand (extensionally) for truth-values (intensionally both may be functions from worlds to extensions, though if terms are referential the function assigned will be constant). Take a simple (syntactically marked) sentence; $S = \textit{sentence}$; $VP = \textit{verb phrase}$:

(3) [_S Jack [_{VP} loves Jill]]

The verb phrase “loves Jill” belongs to a derived syntactic category. Intransitive verb phrases are of type $\langle e, t \rangle$, that is, functions from terms to sentences. Semantically, they are functions from individuals to truth-values. The verb “loves” is transitive, and – disregarding tense or case inflections – belongs to the category $\langle e, \langle e, t \rangle \rangle$. Semantically, *loves* is a function from individuals to functions from individuals to truth-values. We can add determiners:

(4) [_S Some guy [_{VP} loves Jill]]

Quantifiers are second-level predicates. Compared to (3), where the value of the verb phrase operates on the value of “Jack,” in (4) the first-level function “loves Jill” stands for is itself input to the second-level function “some guy” stands for. Quantifier phrases are expressions of syntactic type $\langle \langle e, t \rangle, t \rangle$. Semantically, they stand for second-level functions from first-level functions from individuals to truth-values, to truth-values. Determiners belong to syntactic type $\langle \langle e, t \rangle, \langle \langle e, t \rangle, t \rangle \rangle$ and take predicates to yield

quantifier phrases. To determine the value of $Q\Phi$ one combines the value of Φ with the value of Q , and to obtain an extension one needs a world of evaluation. Whereas the semantic value of an expression of type e is an element of a domain, the value of a quantifier phrase cannot, given compositionality, be an element or even set of elements, but at least a set of set of elements.²⁹

Of course, by asserting that the category e has individuals or constant functions as values and that this is a basic category I beg the question against views that take other types as basic (e.g. Lewis, 1971), or the semantic values of names to be non-constant functions. We have offered at least some support for this assumption, however. Insofar as names require only individuals as values other expressions must have different types of values, and insofar as the grammar assigns semantic values based on *syntactic* category the distinction between referential and other expressions is represented in syntax. Moreover, the class of (semantically) singular terms can be defined syntactically. Singular terms such as “he” and “him” thus belong to the same functional category as names. In Chapter 3 we show that deictic occurrences are singular terms. Since a demonstrative “he” is a singular term all other uses of “he,” including anaphoric uses, should be singular as well, even if their antecedents are not.

2.4.3 Unified Pluralism

On Millian views the value of a name is just the individual, and the value is hence sufficient to guarantee that the name is (*de jure*) rigid. By contrast, I claim that names are

²⁹ See Peters & Westerståhl (2006, p.49–52) for proof that there does not exist a semantics assigning individuals or sets of individuals to phrases that consist of general quantifiers in a systematic (compositional) way.

(*de jure*) rigid because names, in virtue of requirements associated with the grammatical role of names, require (nothing but) objects as values. This provides us with a framework that is as unified as the Millian, but at a different level of generality. Where I think Millian theories tend to go wrong is by assuming that a theory of names must be unified in terms of contributions to propositions; that is, that the unification of the category of names, or singular terms, must be determined by propositions expressed by sentences in which names or singular terms occur. This level is not fundamental. If any level of constraints issues a uniformity requirement it is the level of constraints associated with grammatical category, the expression's functional role. For names and singular terms, the role is to anchor the evaluation of the sentence containing the expression to individuals and make its truth-conditions dependent on how things are with this individual. This is the requirement associated with category *e*, which provides a characterization of the name in terms of its functional role relative to other expressions. The characterization does not itself distinguish constants from variables, or between whether the truth-conditions of a sentence containing the term depends on an already identified individual or not (or whether the individual is a constituent of the Russellian proposition).

Thus, we get a certain form of pluralism about names at the level of propositions – but a unified account in terms of functional or model-theoretic requirements – since there are several ways to derive truth-conditions, or propositions, for sentences containing names consistently with the interpretive requirements. Notice furthermore that by locating the defining semantic properties of names at the level of requirements issued by grammatical category, we may remain somewhat non-committal with respect to theories of propositions. Propositions may be Russellian, but there may be ways to combine direct

reference with unstructured propositions (see e.g. Ninan, 2012), and our view can at least *prima facie* be combined with a truth-theoretic framework.

2.4.4 Montague's Interpretation of Names As Quantifiers

Montague (1973) treated noun phrases consisting of names as type $\langle 1 \rangle$ quantifiers (principal filters on a domain). According to Montague the noun phrase “*Alexius*” denotes in every universe M a set of subsets of M defined, for each individual α , as $(I_\alpha)_M = \{A \subseteq M : \alpha \in A\}$. So *Alexius* (the noun phrase) denotes a quantifier $I_{Alexius}$ of type $\langle 1 \rangle$. The semantic value of the noun phrase is thus strictly speaking not the referent itself. However, there are more general properties that can be used to characterize names independently of the logical representation of names or quantifiers. For instance, names are scope-independent; a name α always sustains the eliminability of λ -expressions ($\lambda x[\Phi x](\alpha) \Rightarrow \Phi(\alpha/x)$ is valid), a fact that would explain why names can always be represented by terms in first-order logic, regardless of what they “really are.”³⁰ This observation suggests that there may be alternative ways of modeling the behavior of singular terms consistent with the assumption that they should behave as individual terms in a proper first-order regimentation, and that they are semantically simple.

On Montague's view names arguably qualify as semantically simple even though their semantic values in a model of language are sets of sets. Since names are still characterized by defining a principal ultrafilter on a domain, the interpretation of the name may still be taken to be constitutively dependent on nothing *but* its extension (e.g.

³⁰ Pasniczek (2002), for instance, provides a way of equating names and quantifiers within first-order logic that retains the difference.

no descriptive information). Thus, for our purposes the Montagovian interpretation may in principle be consistent with the position I defend in this chapter. If this is the case, however, the claim that names function as individual in first-order representations will not necessarily mean that such representations capture all there is to say about names, but that they reflect certain restrictions on how one may interpret them. Perhaps, then, we should be satisfied with committing ourselves to no more than an instrumental view of names that satisfies the following adequacy conditions:

- i) Name occurrences yield object-dependent truth conditions.
- ii) Representing names as constants in first-order systems fully capture their behavior with respect to properties that can be captured in first-order systems.
- iii) Names *can* properly be represented as belonging to the basic grammatical category *e*; that is, representing them as such adequately captures the logical forms of sentences with respect to anything that can be properly represented in first-order (modal) logic.

There are nevertheless some reasons to resist Montague's treatment of names as restricted quantifiers. His motivation is (roughly) as follows.

- (5) Jack thinks that some donkey likes him.

The anaphoric relationship between "Jack" and "him" suggests that (5) should receive a treatment equal to the treatment of "every man thinks that some donkey likes him" in

terms of variable binding. If the extension of “Jack” is the set of all sets containing Jack, then Montague’s approach ensures that (5) is true iff the set of individuals who think themselves to be liked by at least one donkey has Jack as a member. Moreover, quantifier phrases and names exhibit similar functional properties. Both occur both as subjects and as objects of sentences, verbs, and prepositions, and verb phrases combine with both. Montague assumes that if there is a single such rule for forming such sentences there must be a single rule interpreting names and quantifier phrases. If verb phrases denote sets and quantifiers denote sets of sets, a subject-predicate sentence such as (5) is true iff the verb phrase extension is in the extension of the subject. But that means that if the subject “every man” is replaced by “Jack” then “Jack” must denote the set of all sets in which the bearer is an element as well.

One worry, however, is that quantifier phrases themselves do not *prima facie* form a uniform category (although this fact is well-known, it has been used to disassociate referential terms from quantifiers by Neale, 2009, p.402). Remember that quantifiers are supposed to belong to the syntactic category $\langle\langle e, t \rangle, t\rangle$. (3) above had the quantifier phrase in subject position, and that interpretation worked just fine. But how should we treat (3*)?

(3*) [_S Jack [_{VP} loves some girl]]

Transitive verbs are of type $\langle e, \langle e, t \rangle \rangle$. Semantically, they are functions from individuals to functions from individuals to truth-values. But in (3*) the *object* of the verb should be of type $\langle\langle e, t \rangle, t\rangle$, and this cannot work, since we have no available principle for

composing expressions of type $\langle\langle e, t \rangle, t\rangle$ with expressions of type $\langle e, \langle e, t \rangle \rangle$. We could of course say that noun phrases in subject positions are of a different syntactic category than noun phrases in object position (which could be of type $\langle\langle e, \langle e, t \rangle \rangle, \langle e, t \rangle \rangle$), or adjust the category of verbs. Such type-shifts were advocated by Montague himself (1974), who treated transitive verbs as functors, λ -terms over syntactic objects.

It seems at least just as compelling to recognize a level of logical form (*LF*) in the analysis that allows quantifier raising and lowering. (A level of *LF*s may be needed anyway for instance to accommodate covert indexicals or extension-variables (Stanley & Szabó, 2000; Stanley, 2002a).) A rough definition of *LF* could perhaps be that the *LF* of a sentence *s* of a language *L* is the structure that is imposed on *s* in the process of providing a principled truth definition for *L*. Quantifier phrases will simply not appear in argument positions at *LF*; rather, such positions are taken by variables *bound* by the quantifiers. Thus (3*) becomes (3**)

(3**) [Some girl]_x [_S Jack loves *x*]

Variables belong to type *e*. Why, then, should we not associate names with the category of variables (which are, in a sense, the ‘proper’ arguments of verb phrases) rather than with quantifiers? Quantifier movement lets us solve the problem of composition posed by quantifier phrases in object position, but the cost is to introduce an element at the level of *LF*, the variable, that does not belong to type $\langle\langle e, t \rangle, t\rangle$ yet behaves, categorically, in much the same way as names. Montague might presumably interpret variables as quantifier phrases as well, but although this is surely formally possible one may wonder whether the

resulting system will correspond to any practically learnable language and whether it can avoid regress.

Of course, since variables do not, in the case of (3**), belong to any overt syntactic element, Montague could presumably interpret variables as being of type e and retain uniformity of interpretation for all *overt* elements. Even so, adjusting semantic types so that all overt expressions of category e are lifted to the category $\langle\langle e, t \rangle, t\rangle$ appears to obscure his otherwise enlightening analysis of natural language. The semantic type e is suddenly not assigned to any overtly occurring grammatical element at all, and exists in the analysis merely as a technical element to ensure that we can combine the other grammatical elements in the right manner. In short, there are reasons to doubt Montague's own assimilation of names and quantifier phrases. There are admittedly further complications involved here, which we discuss in Chapter 5.³¹ In any case, even if the semantic values of names were quantifiers it would not obviously be a threat to the claim that names are structurally simple, nor obscure the distinction between expressions that function as constants (principal filters) and expressions that function as variables (presumably, then, a range of principles filters).

2.5 Conclusion

In this chapter I argued that singular terms are semantically simple and do not allow any distribution of semantic content between their morphological parts. There is a sharp

³¹ For instance with regard to variable-free semantics (e.g. Barker & Shan, 2008), which avoids variable binding altogether in favor of building semantic values locally, and which are expressively equivalent to analyses involving variable binding (Barker, 2007).

distinction between determiner phrases and singular terms, and views that treat names as determiner phrases are mistaken. While I will argue that the class of singular terms allows variety in terms of e.g. contributions to propositions, semantic simplicity, non-descriptiveness, and the object-tracking (across temporal and modal contexts) feature of singular terms are distinguishing characteristics. The last section opened up for flexibility with regard to how singular terms are represented in a categorial model of meaning. In the next chapter I first establish more precisely what I mean by an expression *functioning as an individual variable or an individual constant*, before I describe further features of singular terms that justify categorizing them with *pronouns*. To sum up where we are:

- i) In Chapter 1 I described some features of singular terms that any adequate account will have to explain, and some constraints on possible explanations.
- ii) In Chapter 2 I explained what distinguishes singular terms from other noun phrases, namely quantifier phrases.
- iii) In Chapter 3 I go into further detail regarding the category of singular terms and attempt to get a clear view of how such expressions work.

Part 2 of the dissertation will explain how we can (and why we should) use the flexibility available to us in our account of singular terms to solve certain familiar puzzles.

CHAPTER 3: THE CONTENT AND CHARACTER OF NAMES

As shown in Chapter 1, we have good reasons to recognize the existence not only of rigid, but *de jure* rigid designators. It is often overlooked (Cumming, 2008, is a notable exception) that even establishing that names are *de jure* rigid is insufficient to show that names correspond to individual constants, since variables are *de jure* rigid as well (though variables are not *de jure* rigid *designators* independently of assignment). I will argue that proper names are sometimes best interpreted as variables. In Chapter 2 I argued that names are “structurally simple” expressions that are rigid in virtue of having (unstructured) individuals as values. Whereas some names function as individual constants and receive values as part of determining the truth-conditions (or proposition), for names that function as variables, while also object-involving, the actual satisfier is irrelevant to determining the truth-conditions or proposition.

In this chapter I first explain what the difference between *functioning as an individual constant* and *functioning as an individual variable* amounts to. I then argue that we have good reasons to recognize some singular terms – pronouns as well as names – that function as individual *constants* (in the sense defined). These are characterized by a distinction between character, or linguistic meaning, and content, the contribution they make to the semantic evaluation of sentences in which they occur. This distinction applies to both names and pronouns, and suggests that names and pronouns share interesting features over and above the requirement that they track individuals rather than roles. I use this observation to motivate the idea that names *are* special duty pronouns (following, loosely, a claim made by Sommer, 1982). In part two of this dissertation I

argue that these characteristics of singular terms are entirely general; anaphoric pronouns are singular terms associated with a character–content distinction as well. Occurrences of anaphoric pronouns should be interpreted as (semantically bound) variables, and the distinction between anaphoric and deictic pronouns is explained by whether the context relative to which the character provides an interpretation is linguistic or extra-linguistic. If names are pronouns, it is natural to ask whether names have anaphoric occurrences as well. I claim that they do, though there are reasons, given in Chapter 6, for why such occurrences are relatively rare compared to anaphoric occurrences of ordinary pronouns. The claim that we cannot provide a unified account of names as constants, and have to recognize the existence of anaphoric occurrences, is defended in Chapter 4.

3.1 Constants and Variables

In the last chapter we left it somewhat open what we mean by saying that *singular terms can function as both constants and variables*. In this section I spell out what the difference amounts to and what it means for natural language semantics. I will merely assume that anaphoric pronouns are bound variables; and provide a more sustained defense in Chapter 5.

3.1.1. Variables and Constants in Natural Language

Take the sentence

(1) He loves her.

In some contexts (1) is interpreted as involving demonstratives and to say of a particular, contextually salient person that he loves some other contextually salient person. In that case, the phrase is a complete, meaningful unit relating two individuals by the binary predicate “loves.” But the phrase (1) can also occur as part of a larger unit, as in

(2) Whenever a woman praises a man, he loves her.

One way to understand how (1) works would be to represent it as an open sentence:

(3) x loves y .

When the occurrences of “he” and “she” are demonstratives a particular, demonstrated, extra-linguistic object is assigned to x and an object is assigned to y . When (1) occurs in (2), the pronouns are bound by some form of antecedent variable-binding operator.

Thus, pronouns appear to occur both in bound and unbound position. I will claim that names, like pronouns, can occur both “bound” and “unbound” as well and, indeed, that the difference between names and pronouns is minimal and primarily pertaining to focus and topic effects that ensure that names are less often obviously bound. This claim calls for a general definition of the notions of “individual variable” and “individual constant” that can be applied to ordinary language semantics.³² Note that we assume

³² Note that I am not trying to provide an exhaustive general definition of *variable* or *constant*, but a definition suitable for modeling a difference in natural language. The metaphysics of variables is a

objectual quantifiers and hence that variables are rigid (designators relative to assignments). A first attempt may be

I) An individual constant is a term whose value i) is an individual of the domain; ii) is determined as part of identifying the intended model for a sentence. A (bound) variable, then, is a term whose value is not part of identifying the intended model for a sentence, though any value assigned must be an individual of the domain.

A perhaps more general definition could be:

I*) A constant is an individual term whose satisfier (an individual) contributes to determining the truth conditions of a sentence, and not just its truth-value at a circumstance of evaluation. A (bound) variable is a term whose satisfier is received as part of determining the truth-value of a sentence, not the truth conditions. Thus, a variable can have different satisfiers in different contexts. A constant cannot (once we have fixed an interpretation).

Let us, at least as an illustration, identify the truth conditions of a sentence *s* with *the possible worlds in which s is true*, and consider a sentence “Aristotle is *F*.” If “Aristotle” is a constant, then the sentence is true in all worlds in which the individual, Aristotle, is *F*. Consider “the author of *Nicomachean Ethics* is *F*” under the assumption that its logical

different, and much harder, issue, and I am not optimistic that *variable* can be defined in any non-circular or non-regressive manner.

form is [the x : Author[...] x] (x is F) and that the determiner implies uniqueness. “ x is F ” is true in all worlds in which some individual satisfies F that is also the (unique) satisfier of “Author[...](x).” Saying that *a constant contributes to determining the truth-conditions* means that in order to determine the truth-conditions of “Aristotle is F ” we already have to determine who Aristotle is – not just any assignment of a value to “Aristotle” will do the trick. Saying that *the variable contributes to determining the truth-value* of a sentence though not the truth-conditions, means that whether “ x is F ” is true depends on there being an individual that satisfies the open sentence relative to a world. Determining whether the sentence is true requires finding such an individual. But *which* individual this is, is irrelevant to determining the set of worlds in which the sentence is true (its truth-conditions). Any satisfier would do, and which individual satisfies it would presumably vary from world to world. (The variable is nevertheless rigid since “ x is the F but x could have failed to be the F ” requires that any assignment to the first occurrence of “ x ” must be the same as the second occurrence even though the second conjunct is evaluated at a different world than the first conjunct.)

3.1.2. Free Variables

Though variables and constants serve different purposes they play the same grammatical role, namely to contribute individuals to the semantic evaluation. Other expressions (quantifiers, predicates, determiners, prepositions) are distinguished from singular terms already at the level of what types of contributions they make to semantic evaluation. The difference between variables and constants is not defined at this level, but at the level of how values are assigned. Perhaps it is helpful, then, to distinguish individual terms from

other types of expressions by defining them as (free) variables. Though the interpretation of free variables is somewhat contentious, we can say that a free variable ranges over the individuals of a domain. Some terms, constants, have their values already determined prior to determining under what conditions a sentence is true. Other terms are not assigned individual values in advance of evaluating the sentence for truth, but are instead bound by quantifier phrases that express a condition for what would be an acceptable value for deeming the sentence containing it true.

I usually assume a Tarskian framework (unless otherwise specified). A *constant* is a term whose value is determined by the interpretation function and identified in the process of determining the intended model for a sentence. A *variable*, on the other hand, is a term whose value is not determined in advance of determining a model. The value assigned to it is not its semantic content; rather, values are assigned to variables as a step in obtaining inductive definitions of satisfaction for quantified sentences. The variable, however, *requires* or *calls for* an individual. If free variables are the ‘basic’ singular terms, we can say that some variable occurrences are assignment invariant, and ‘attached’ to a given value in a semantic evaluation so that a proper evaluation returns that value regardless of circumstance of evaluation (Note that a language may contain other variables and constants as well, such as second-order and plural variables and constants; we will leave these aside.)

Tarskian semantics provides no semantic value for free variables per se; they only receive values relative to assignments. Thus, whereas the value of a constant is given by the interpretation function and the value of a variable by the assignment function, it is unclear what the value of a (*free*) *variable* could be independently of assignment when

contributions of singular terms vary according to whether they are constant or bound. If natural language contains (non-constant) free variables their values should be determined independently of assignment, since assignments of values are an artificial step required by formal truth definitions. Surely the value of the variable cannot be just the assigned individual. If it were, we would not be able to distinguish between the demonstrative interpretation of (1), and (1) as it occurs in (3). For the moment I suggest the following definition (to be refined later):

- i) The semantic value of a variable is an identity function on domains (or universes). The variable designates nothing, but ranges over a domain (and different universes with respect to different possible worlds).
- ii) In the case of constants the input (individual) is already determined (by the interpretation function). Hence, the value of the free variable in its guise as a constant becomes a constant function from indices of evaluation to individuals.
- iii) The semantic value of a variable in bound position is determined as follows: In “ x is F ,” “ x ” is an identity function over a domain. This is its contribution to the open sentence. The semantic value of “ x is F ” thus becomes a function f from possible assignments A (outputs of the identity function) to closed formulas (or propositions) under the assignment of the value of $f(A)$ (“ F ” itself is as usual understood as a function from individuals to truth-values, and propositions can – at least for simplicity – be identified as functions from indices to truth-values). $\exists xFx$ is then true if there is an output f of “ x ” (relative to a domain D provided by the circumstances of evaluation) relative to which Ff is true (or a variant output, if

we follow Tarski closely), and $\forall xFx$ is true if all possible outputs for “ x ” (relative to D) satisfy F .

The crucial question is how these distinctions map onto natural language. The claim that proper names behave as individual terms is an empirical claim, though previous chapters have provided support for the idea. To function as a *constant*, a name must i) contribute no descriptive content to the truth conditions of sentences (or no content apart from the referent that affect the truth conditions of sentences); and ii) the value (referent) must be determined in advance of assessing the sentence for truth at a particular circumstance of evaluation. If a term (occurrence) satisfies these conditions it is directly referential.

A variable, then, is not constant across times and worlds of evaluation. The value that satisfies $\exists xFx$ at a given world-time pair may vary without the truth-conditions of the sentence varying.³³ The distinction is a distinction between whether the value affects the truth-condition of a sentence or only its truth relative to a circumstance of evaluation. Relative to Tarskian semantics the difference is best modeled in terms of constants having their values assigned by the interpretation function. When I talk about a constant receiving its value as part of identifying the intended model for a sentence this should, however, be understood relative to a Tarskian representation and not as a fully general definition of a *constant*, nor as meaning that a first-order interpretation fully captures all

³³ In fact, this claim is imprecise. A variable is insensitive to world of evaluation in the sense that $\exists x(Fx \wedge \Diamond \neg Fx)$ requires that the same value is assigned to both occurrences of x , regardless of whether the worlds assigned to the conjuncts are different (compare $\exists xFx \wedge \Diamond \exists x \neg Fx$). In this sense the value of a variable is also insensitive to the world parameter. In standard Kripke semantics, worlds are also associated with assignment functions, and the value of a variable is of course sensitive to the assignment (that said, it is imprecise to talk about the value of a bound variable being the individual it denotes relative to an assignment; a bound variable rather ranges over a domain of individuals). A constant is, by definition, not sensitive to the assignment parameter.

aspects of how the semantics of singular terms should be modeled for all purposes. The Tarski distinction between terms that get their values from the interpretation function and terms that get their values from the assignment function nevertheless reflect an important natural language distinction.

3.1.3. Some Clarifications

The definition of “variable” and “constant” might differ somewhat from how these terms are used elsewhere. A case in point would be instantial terms or flagged constants in certain proof systems of first-order logic, which would count as variables given our definitions – the quantifier elimination rules require precisely that it does not matter to the truth-conditions of sentences which individual from a domain is chosen; indeed, for universal introduction to be a valid rule it is required that the selection does not matter.³⁴

Compared to the use in mathematics there are some complications. Consider for instance the quadratic equation schema $ax^2 + bx + c = 0$. Here a , b , and c are generally called “constants,” whereas x is a variable. The governing idea is that a , b , and c do not contain modifiable elements – their values are given by a model, in the sense that to determine a value for x the values of a , b and c must be stipulated. The value for x , on the other hand, is determined *in* the model. Hence, a , b and c seem to qualify as constants according to our definitions. However, the notions of “constant” and “variable” in mathematics are only defined relative to models, and the definition – while it captures the difference between variables and constants in Tarskian terms – is relevantly different from what we need when considering natural language.

³⁴ Though see Breckenridge & Magidor, 2012; discussed in Chapter 4.

In a given interpretation (for the schema) $\mathcal{M}a$ would refer to a given number, say, 23, whereas in a different model \mathcal{M}^*a may refer to a different number, say, 14. Natural language sentences are presumably not model-relative in quite this sense. There may be a certain suggestive analogy between a , b and c and *demonstrative pronouns*, insofar as the interpretation of demonstrative pronouns varies with context, yet what they contribute in each particular context exhibits the truth-conditional behavior of a constant. I am wary of pushing this analogy too far, however. As opposed to demonstratives in natural language, there is a clear sense in which the schema has a non-model relative interpretation as well. I.e., since the quadratic equation schema applies to any value given to a , b , and c , regardless of model, and since we can say general things about the schema by quantifying over models, there is a relatively clear sense in which they are variables as well – the *interpretation* of a is non-constant over models. That is, relative to \mathcal{M} , a is a constant, but to interpret the *schema*, we need to quantify over models, which would make a a variable (since it gets a value only relative to a model). Although it is certainly possible to provide a logic of demonstratives (Kaplan, 1977), it is unclear whether the logical properties of the *linguistic meaning* (character) of demonstratives correspond to the quantification over models involved e.g. in proving mathematical theorems.

3.1.4. Conclusion

Variables and constants are both *de jure* rigid expressions, and they are rigid in virtue of properties associated with their grammatical category. A singular term – variable or constant – is one that is satisfied by or designates an individual (or perhaps several

individuals in the case of plural expressions) of a domain, and not a complex compositional structure. It does not follow from these considerations alone that all expressions belonging to the category of singular terms refer, though it *is* required that their satisfiers are the individuals (and only the individuals). Thus, I propose that what characterizes singular terms, and what distinguishes them from other expression types such as quantifiers and predicates, can be captured – at its most general level – by saying that they correspond to the free variable x . In some contexts x could be directly referential – a constant – in others, it could be bound.

The universal theory of singular terms is the theory that associates them with individual variables, not constants. The distinction between variables and constants is one that must be drawn on a case-to-case basis. This is already the case with pronouns, as illustrated by the fact that the pronoun “he” can occur both as a demonstrative pronoun and as an anaphoric (bound) pronoun. In other words, by giving up the Millian assumption that names are always constants, we do not jettison universality in our theory of singular terms, but locate the universal traits at the level that actually does the job of distinguishing singular terms from other types of expressions. In 3.2 I discuss singular terms that (given our definitions) function as constants, including some name occurrences. From these observations I draw some general conclusions about the semantic properties of singular terms and (briefly) discuss the differences between Fregean and Russellian approaches to the role of semantic analysis.

3.2 Names As (Directly Referential) Constants

Insofar as names are singular terms with individuals as values, we can show that certain singular terms, indexicals and demonstratives, function as constants, and provide strong evidence that at least certain name occurrences do so as well.

3.2.1. Indexicals and Constants

The referent of an indexical varies with context, but the indexical has a stable linguistic meaning. In Kaplan's framework its linguistic meaning is its *character* (Kaplan, 1977), a function from contexts to contents. The character of "now" is a function that maps a context C onto the time of C ; the character of "I" is a function that maps the context C onto the agent of C , and so on. A context is defined as a quadruple of time, world, place and speaker. The value of the function defined by the character is, for C as input, the *content*, the contribution of the indexical to the *proposition* expressed by a sentence in which it occurs. That contribution is precisely the individual designated by the occurrence of the indexical or demonstrative. Thus, the time of C is the content of an utterance of "now" relative to C , and the lexical meaning of "now" is relevant (at most) to fixing the reference of an occurrence of "now;" once fixed, only the referent is part of the proposition. Kaplan officially defines a proposition as a function from circumstances of evaluation (world-time pairs) to truth-values, which makes the content of an indexical "now" a constant function from circumstances of evaluation to individuals. (It is possible to recast the view relative to other definitions of propositions as well.)

Indexicals are thus directly referential. If the lexical meaning of an indexical were part of the proposition expressed we get the wrong results. For instance, the character of “I” is – roughly – “the speaker.” If “I” denoted the x such that x is speaking in some world w at t , then the sentence “I do not exist” is true at a circumstance of evaluation $\langle t, w \rangle$ just in case the x that is the speaker at w, t does not exist. Now, if such a speaker exists, the sentence uttered is false at $\langle t, w \rangle$. If the speaker does not exist, on the other hand, “the speaker” denotes no one at $\langle t, w \rangle$, in which case the sentence is, once again, false. Hence, “I do not exist” comes out necessarily false. But “I do not exist” is not necessarily false and, conversely, “I exist” is not necessarily true. Hence, the linguistic meaning is not part of the content of an indexical.

Moreover, not only is the function from circumstances of evaluation constant; it is also *total*. The extension (content) of an indexical is fixed with respect to a context C *prior to* any evaluation of the sentence in which it occurs. The extension is hence independent of the circumstance; it is, in Kaplan’s words, “no more a function (constant or otherwise) of circumstance, than my action is a function of your desires when I decide to do it whether you like it or not” (Kaplan, 1989). But then, an indexical like “I” as used in a context C will denote the agent a of C , C_a , in *all* circumstances $\langle t^*, w^* \rangle$, regardless of whether a exists in t^* and w^* . It is as such an obstinate rather than a persistent rigid designator. In short, the semantic contribution of an indexical is best interpreted as a constant, relative to our definition of ‘constant’ above, since the semantic value is determined in advance of circumstances of evaluation, and since the contribution is the referent itself. Although “he” is notionally a free variable, demonstrative occurrences of

“he” work in the (relevantly) same manner as indexicals and should, at the level of logical form, be represented as a constant.

Certain name occurrences appear to be constants as well, for instance because they appear to be sensitive to Kripke’s *semantic argument*. For such name occurrences the role of any descriptions associated with it is to identify the referent, for instance by identifying the correct name-using tradition that originates in the referent itself. This is illustrated by the fact that which description is associated with the name is irrelevant as long as it successfully identifies the correct tradition. For instance, when introducing the name “Aristotle,” a speaker might have in mind “the last great philosopher of Antiquity” or “the author of *Nicomachean Ethics*,” but which one it is does not matter to the truth conditions of subsequent sentences involving the name. In fact, the description may not even be correct. Someone who believed that Einstein was the inventor of the atomic bomb could conceivably, in the right context, use this (false) assumption to latch onto the correct name-using tradition (Kripke, 1972, p.82). And certainly someone who uses “Aristotle” to refer to Aristotle does not need to be aware of the description (or baptismal act) originally used to instigate the name-using tradition. Perhaps the most obvious examples illustrating why the associated description must be independent of the semantic contribution of a name, are cases of incomplete descriptive knowledge. To expand on another example by Kripke (1972, p.91), speakers may associate with the name “Feynman” just an indefinite description “a physicist,” perhaps in addition to “person called ‘Feynman,’” yet succeed in referring to him. Those descriptions, even in combination, are not sufficient to pick out a unique individual. But it appears that speakers who use the name do *not* intend to use it in a manner that makes it irrelevant

which of multiple physicists called “Feynman” satisfies the sentences they utter, which would be the consequence (if the sentences have a truth-value at all) if the descriptions affected the semantic content. These cases furthermore suggest that speakers do, at least implicitly, rely on some meta-linguistic criterion when interpreting a name, for instance an intention to latch onto a name-using tradition that anchors the use of the name in the intended referent.

It may be argued that treating “Feynman” as a variable bound by a description could yield an adequate analysis. However, since the associated descriptions are not sufficient to individuate the correct individual, this suggestion would require the presence of a lexical element *i* in the binder, such that *i* individuates the correct tradition. Suppose, then, that we represent the binder by the indexed description “the bearer_{*i*} of ‘Feynman.’” What determines which value is assigned to this element? It could not be the associated description “the bearer of ‘Feynman,’” since the problem is exactly that “the bearer of ‘Feynman’” is improper and has several satisfiers and we have no further means of distinguishing between them. In order to give the correct analysis *i* must be considered to *refer* to the correct name-using tradition, namely the one that originates in (the correct) Feynman. Hence, the binder itself would contain a constant, and it seems more parsimonious to let the name “Feynman” itself be the constant that does this work.

3.2.2. Reference and Character: Individuating Names I

Insofar as names have directly referential occurrences, we could, as a start, associate the linguistic meaning of a name with a character, analogous to how a pronoun “he” has a constant linguistic meaning but contributes different semantic values in different

contexts, and hence formally as functions from contexts to contents determined by the standing meaning of the expression. Kaplan (1977) introduced the operator *dthat*, a term-forming expression that combines with a description to form a term.³⁵ The character of the term corresponds to the intension of the definite description it embeds. Thus, the character of “*dthat*(the *F*)” maps any context onto the constant function from world-time pairs to the object that is *F* in *C*. As with indexicals the denotation in *C* is assigned prior to the evaluation of a sentence, so the complex intension determined by “the *F*” is no part of the semantic content of “*dthat*(the *F*).” *dthat* is thus a rigidifier that transposes the descriptive content of a term to a pre-semantic level of character. A name *n* could similarly be interpreted as “*dthat*[the bearer of *n*].” The reference depends on the individual that was originally given the name, and the character or context should in some manner invoke the name-using tradition going back to e.g. a dubbing.

On a demonstrative view of names like this character can, to an extent, replace Fregean senses, but unlike Fregean senses characters are no part of the semantic content of a name.³⁶ Rather, the connection between the occurrence of a name and a referent, a referential chain between the origin of the name and a given occurrence that ensures that the name refers to what it refers to, is pre-semantic and not encoded in the semantics of the name. Since it is determined by conventions external to speakers it may not be

³⁵ There are several complications involved in this move, for instance regarding how exactly the term-former *dthat* is supposed to be understood. On Kaplan’s official view *dthat* does not have a complete character, but gets a character when combined with a demonstration. In that case, a name (analyzed using *dthat*) would strictly speaking not have a constant linguistic meaning (see Salmon, 2002).

³⁶ It is sometimes argued that the character may enter into the evaluation of the name when it is embedded under a hyperintensional operator (Recanati, 1993; 2000). There are ways of making this idea consistent with the claim that names are semantically unstructured. One could argue that the operators in question invoke a mode of presentation for the proposition referred to by the “that”-clause (Salmon, 1986, although Salmon takes such information to be conveyed rather than semantically expressed). In that case it would be the “that”-operator that invoked the complex semantic layer, and not the name. I will instead (Chapter 7) suggest taking the names to be bound by locally defined discourse operators in hypetintensional contexts.

something ordinary speakers will even know. There are still some open options regarding how to understand the link between names and referents, where the choice does matter.

Kaplan (1990) defends what he calls a “common currency” view of names, according to which names are, metaphysically, individuated by their referential links. This makes the character itself object dependent, and two occurrences of a name (the expression) *n* with different referents are occurrences of different names. Alternatively, one can take *n* to be a syntactic type that is used to refer to different individuals in different contexts. *n* would be “*dthat*(the salient individual called *n*)” (Recanati, 1993,³⁷ 2000). Its character is formed from the linguistic meaning “the bearer of *n*,” and the referent is determined by character and context. On both views names are associated with characters and receive contents in contexts, but on the second view the semantic content of a particular name varies whereas on the common currency view context determines *which* name is used. I will in general assume that single names have different referents in different contexts, though it may be possible to adapt what I say to the common currency view.³⁸ Both views accounts for *how* a name gets a content in a context, but the context sensitivity view more intuitively captures the idea that character is the linguistic meaning of a name. In either case, however, we obtain a mechanism for fixing the referent independently of circumstances of evaluation. It should also be noted that in either case are names more closely related to demonstratives than to indexicals (though the

³⁷ Recanati’s earlier view is different. On his earlier view a difference between names and indexicals was that the latter were associated with a particular relation by semantic conventions such that every indexical was associated with a different relation. Proper names, on the other hand, were all associated with the same relation defined by the feature REF. Thus all names would have the same meaning (consisting of REF and the truth-conditionally irrelevant mode of presentation “bearer of that name”), whereas indexicals would be individuated partly by meaning. Proper names were individuated only by their form, and the use of a name in a context served to invoke the name-using convention associated with that form.

³⁸ See Hawthorne & Lepore (2011) for a critique of Kaplan’s view.

difference between demonstratives and indexicals is less on the common currency view). Like a demonstrative a name requires a stipulative act to connect the expression to an individual – linguistic meaning alone is not sufficient to determine a referent.

3.2.3. Reference and Character: Individuating Names II

We can sketch how this could work in a semantic analysis. An expression's character, combined with a context *C*, yields a content, which yields an extension at a circumstance of evaluation. For many expressions, e.g. predicates, character and content coincide (if the values of sentences in *C* are structured propositions both the character and content of a predicate is a property, for instance). For names and indexicals character and content come apart, and content coincides with extension. This leaves us with two main options.

i) On the first option, characters are the linguistic meanings of an expression. The character of an indexical "I" is a description of the kind "the person uttering this token," which combined with a context *C* yields a content *dthat*[the person uttering this token]. For a name the character is the linguistic meaning, "person called *n*," and *C* provides the referent of a given occurrence of *n* represented as *dthat*[the person called *n*] as content. In general, sentences have characters composed by the characters of their constituent expressions, and in a context these determine the proposition, which together with the circumstances of evaluation determine a truth-value. Call this option "R1."

ii) The second option, R2, takes the character to be incomplete (just as the character of "that" is incomplete on Kaplan's view) and in need of a demonstration or name-using tradition to determine a *character*. Demonstratives and (generic) names have different characters on different occasions and need to be supplemented by something

(e.g. a demonstration) to determine a *character*. So there is not one name “Eric” but several; indeed, each *use* has a different linguistic meaning (or is a different stage in the unfolding life of a name). Nevertheless, once the character is settled we proceed as before; a context *C* combines with the character to provide an extension, which is the contribution to the proposition expressed by the sentence in *C*. The main differences between R1 and R2 concern whether the name-using tradition is part of the character or the context, and whether the character is what we can rightfully call the linguistic meaning. Kaplan’s common currency view is a version of R2.

There are at least some reasons to choose R1. First, a name seems to have a complete character; on R2 a character is determined only in context. Second, on R1 a name will have a complete linguistic meaning; on R2 a name is a syncategorematic element that takes a context to produce a linguistic meaning. Now, one standard complaint against Kaplan’s “one-to-one” view is that it generates massive ambiguity (Bach, 1987; 2002), but against R2 in general this complaint is misguided. Names are not massively ambiguous on R2; not only do they not have any *semantic* meaning in isolation, they do not even have a linguistic meaning (apart from a syncategorematic unit that is not ambiguous at all) in isolation. In any case we assume R1 in what follows; the fact that the character is determined by linguistic meaning alone is a distinct advantage when solving some of the challenges in later chapters.

3.2.4. Names, Characters and Senses

In order to fully establish the connection between singular terms in natural language and individual terms in first-order logic, we need values of names to be individuals of the

domain. An alternative is a Fregean view that takes names to correspond *extensionally* to constants (or variables) of first-order logic, but their semantic values – the contributions they make to propositions – to be Fregean senses. The choice between a Fregean and a Russellian view depends to a large extent on what explanatory work propositions are supposed to do, and indeed on the role of semantic theory, for instance regarding the extent to which propositions are supposed to reflect cognitive significance or not.

Russellian theories often take propositions to be structured complexes the constituents of which are the semantic contents of the parts of the sentences that express them.³⁹ The character of a sentence, then, is ‘really’ a function from contexts to structured propositions, and the model-theoretic semantics that Kaplan provides must be viewed as an approximation. Fregean propositions, on the other hand, are composed of senses reflecting the cognitive significance of expressions. For a *predicate* the sense is its linguistic meaning, which together with a possible world determines an extension. For names the situation is more open.

i) One possibility is to view the sense of a name as a mode of presentation, distinct from its linguistic meaning and prone to vary from occasion to occasion and from individual to individual (e.g. Chalmers, 1998; 2002; Sainsbury, 2002⁴⁰). A mode of presentation is nevertheless a mode of presentation *of* an object, and the sense of a *particular use* of a name is rigidified to account for the rigidity of names. So if one participant in a conversation associates Frege with “the person who discovered the

³⁹ Structured propositions have been defended by, among others, King (1995, 2007), Soames (1987, 2010a), and Salmon (1986). Similar considerations are also the background for viewing semantic contents as *interpreted logical forms* (Larson & Ludlow, 1993), or for Cresswell’s account of propositions in terms of structured *intensions* (Cresswell, 1985).

⁴⁰ Though Sainsbury eschews propositions. Instead, the difference is a matter of information associated with the level of axioms in a truth-theoretic framework.

quantifiers_{rigidified}” and another with “the person who invented logicism_{rigidified},” the extension is the same with respect to all worlds. Since senses are modes of presentation and modes of presentation vary between contexts, senses vary between contexts. In other words, senses are associated with specific *uses* of names. The ‘completed’ sense in a context *C* determines the extension in *C*.

ii) Alternatively, the sense is the linguistic meaning (Balaguer, 2005). In that case the sense of “Frege” would be (a rigidification of) “the contextually salient person called ‘Frege.’” In that case sense alone cannot determine reference; rather, the referent is determined by the combination of sense and context (as opposed, for instance, to a predicate, for which the sense is constant and determines an extension at every circumstance of evaluation).

In either case, a context is needed to determine reference. There is a suggestive correspondence between these Fregean views and the two Russellian views described above. Senses appear to correspond closely to characters on Russellian views. For instance, on R1 two utterances can have the same character yet differ in content; on the second Fregean view two utterances can have the same sense yet different truth conditions. Furthermore, Fregean truth-conditions, defined by the extensions of the elements of a proposition, correspond closely to Russellian propositions. Balaguer (2011) argues that Fregean and Russellian views are isomorphic up to a point where no available semantic intuitions distinguish between them; the distinction concerns mainly what one takes propositions expressed by sentences to be, and this is mainly a result of what one takes to be the theoretical role of the notion of *what is said*, i.e. whether it should reflect semantic meaning or truth conditions of sentences (Russellian), or cognitive significance

(Fregean). Russellian views allow characters to at least indicate cognitive significance, but since character is no part of semantic content a Russellian account does not provide a *semantic theory of understanding*. Thus, which account we choose depends at least in part on whether a semantic theory should be a theory of understanding.

The analogy should not be taken too far. Balaguers conjecture relies on a crucial conflation. On a Fregean view the sense of a name, its contribution to a proposition, is a rigidified description. Balaguer, however, describes a view like R2 as follows:

Thus, just as the character of ‘he’ is something like *dthat[male: x]*, so the character of ‘David’ is something like *dthat[bearer of ‘David’: x]*. Thus, on this view, the character of ‘David’ is incomplete. But particular uses of ‘David’ have complete characters, and different uses of ‘David’ can have different characters [*ibid.*, p.59]

But this is wrong in a crucial respect. On a Russellian view, *dthat[bearer of “David”]* is the *content* – or equivalent to the content – of a particular use of “David,” not the *character*, as Balaguer suggests. The *character* of a (use of a) name is rather the *non-rigidified* description “the bearer of *n*” – indeed, the character of “he” is the non-rigid “the most salient male” and the character of *dthat[bearer of “David”]* is the non-rigid “the bearer of ‘David.’” *Dthat*-constructions are already *terms*, and hence directly referential, and Kaplan introduced them precisely not to provide the character of natural language occurrences of the demonstrative “that” but to represent the *content*. As opposed to what Balaguer suggests, Russellian theories *do not require rigidification* at any level of semantic analysis. The expression is already directly referential (on the view we have defended because of a requirement associated with grammatical category). The character provides a means for locating the correct assignment in to a context, but rigidity

and direct reference are already guaranteed by linguistic requirements and not introduced by any implicit or explicit rigidifying operation.

Thus, what I have claimed is a significant drawback with Fregean analyses still distinguishes Fregean from Russellian views, despite Balaguer's conjecture. On Fregean views the descriptive sense of a name is a *rigidified* description, and we need an explanation not only of the mechanism but of why those senses are *rigidified* rather than non-rigidified. The answer that they have to be rigid to ensure the correct modal behavior is dangerously *ad hoc*. Russellian views raise no similar worry, since the modal behavior of names falls out of the referential requirement, which is not defined in modal terms.

One prominent Fregean view that deserves mentioning takes as its point of departure the hypothesis that the meaning of a name is sustained in a language in virtue of a causal link of transmission stretching back to the original introduction of the name. The Fregean view in question attempts to incorporate this tradition or link into the meaning of the name itself (Lewis, 1984; Kroon, 1987; Jackson, 1998a, 1998b). *Causal descriptivism* is the view that the meaning or sense of a name is given by an actually-rigidified description of this process, or of a description along the lines of "the object standing in relation R to my current use of the name n ," where R is usually taken to be a causal relation. Elaborate discussions and criticisms of the view is given in Soames (2002) and Everett (2005), and I will not repeat the details. In general, if a principle such as R gives the meaning of n it implies that in order to understand n a user must – if perhaps not (implausibly) have any knowledge of the causal source or history of his or her use of the name – at least be tacit possession of a theory of reference, which is highly implausible. Furthermore, it seems difficult to spell out R in any particular case without

invoking some indexical relation between the speaker and the origin of the name, which would threaten to make meanings essentially private.

The existence of some such link between a particular use of a name and its origin is plausible.⁴¹ It is also reasonable to assume that this link can in principle be described. It is, however, implausible to think that the description of the link is part of the *meaning* (or even cognitive significance) of the name. Instead, the causal history is pre-semantic or foundational – a relation *R* must hold for a name to have the meaning it has. Foundational facts about how a word came to refer to what it does, or the facts in virtue of which it refers to some individual, should not be confused with semantic facts about the word's meaning. Instead, a name-using tradition (which may be invoked e.g. by a tacit intention to use an expression with the same meaning as one's linguistic surroundings) is part of the extra-linguistic context that helps determine the meaning of a name occurrence, and not something anyone needs to know in order to know the meaning of the name.

3.3. Names As Pronouns

In 3.2 I argued that names sometimes function as constants. When they do, their interpretation corresponds to the interpretation of a pronoun (“he,” “she,” “that,” “it”)

⁴¹ Although I am sympathetic to inferential role semantics (IRS) in general, IRS should be viewed as an account of the grounds or foundations of the meanings of (uses of) expressions, i.e. why an expression has the meaning it has. Even Millian views need a theory of these foundations. Greenberg & Harman (2006) suggest that IRS should be viewed as a theory of how meanings are determined, rather than a theory of what the meanings of expressions are, and this view is consistent with Millianism (even in the case of names). There are furthermore good reasons for why some factor external to semantics or inferential roles must go into determining the meaning of an expression (and in particular names). For instance, we need a way to defuse Putnam's “just more theory” charge to solutions to Putnam's Paradox (concerning, roughly, the possibility of distinguishing the intended model for a language from non-standard ones), anchoring semantic content to extra-semantic features of the world appears to be essential.

used demonstratively. Pronouns are associated with a distinction between character and content. To grasp the proposition expressed by “he is *F*” one needs to know who “he” refers to in the context of utterance, though one can (pragmatically) get a rough approximation of what is said by knowing the linguistic meaning of “he;” that is, from knowing the character. Similarly, to grasp the proposition expressed by “Jack is *F*” one needs to know who “Jack” refers to in the context of utterance, though one can get a rough approximation of what is said – and know *that* the proposition expressed is singular – by knowing the linguistic meaning of “Jack;” that is, from knowing that “Jack” is a name and the character “the bearer of ‘Jack.’” The character is no part of the semantic content of the utterance, however.

The similarity with respect to the character–content distinction suggests going one step further and trying out the hypothesis that names *are* a special kind of pronouns (an idea suggested by Sommer, 1982). Like “he,” a name “Jack” is a tag that we can fasten on to a referent in a context. Of course, there are some significant differences between names and ordinary pronouns such as “he,” and we will explore these in Chapter 6 – uses of names do for instance, often invoke a name-using convention linking the name to its referent; pronouns do not. I suggest that most of the differences can be explained by the simple fact that names are more specialized and have more specialized characters. The very reason we use names is to make more distinctions and keep track of more individuals than ordinary pronouns allow us to do, to the extent that particular names are – at least often – associated with specific, single individuals. Coining a name means adding a character to our linguistic resources that can serve the purpose of individuating as many individuals as we wish, and keep track of them in discourse. Another suggested

difference is that pronouns also occur in anaphoric position. Part 2 of this dissertation is devoted to showing that names occur anaphorically as well (as well as the claim that anaphoric pronouns are singular terms).

On all occurrences, whether anaphoric or deictic, names and pronouns are associated with a difference between linguistic meaning – the character, which is used to provide an interpretation of the term in a context – and content. Even in anaphoric position, the character plays a significant role, namely in helping to determine the proper antecedent. The difference between deictic and anaphoric occurrences is that in the former the expression picks out an individual from an extra-linguistic context, whereas in the latter the relevant context is linguistic. In *all* cases the proffered content of a name or pronoun is that of a free variable. If the relevant context or contextual parameter is extra-linguistic, the character provides an object, ensuring that the free variable is interpreted as a constant. If the relevant context is linguistic, the character connects the expression to an antecedent, which then works as a (semantic) binder. We provide the details of these processes in Chapter 5. In the remainder of this chapter I provide an outline of features that suggest classifying names with pronouns.

3.3.1 Names and Pronouns As Definite Expressions

Pronouns are definite expressions. Now, the correct characterization of definiteness is a matter of discussion (see Abbott, 2001, for an overview), but I will maintain, following Heim (1983), Geurts (1997), and others that definiteness is at least associated with a familiarity constraint (whether or not definiteness is *also* associated with a uniqueness or maximality constraint). Definite expressions constitute what Gundel et al. (1993) call a

“givenness hierarchy.” Let “discourse referent” (*dref*) for the moment be understood simply as *the object talked about*. Roughly, drefs that are uniquely identifiable (at least relative to the context of a particular discourse), for instance by being introduced by an indefinite description and proper contextual restrictions, occupy the bottom level. These drefs can be identified by thematic roles, and be picked up by *definite* descriptions. Definite descriptions have the longest reach. A sufficiently specified definite description can pick up drefs from far earlier in a discourse or even, possibly, across discourses (as when a discourse starts with “do you recall the strange, one-eyed guy?”). By contrast, *activated* drefs are those that are retrievable by demonstratives (if salient extra-linguistic items) or discourse anaphora. On the top of the hierarchy there are drefs that are in-focus, i.e. objects mentioned in the subject position of a sentence.

The reach of names in this hierarchy is more complicated. There seems for instance to be certain familiarity constraints associated with whether to use a full name, e.g. “Tim Minchin,” or whether “Minchin” or “Tim” will do. To felicitously use a name in a discourse the audience must at least somehow be made aware of the individual one is talking about. In discourse initial position – when names are introduced to a discourse – they are often accompanied by descriptions (e.g. “according to Adel Dabecchi, the coordinator of TNC, [...]”) that help audiences identify the referent (or employ the correct name-using tradition). In such cases the description should probably occur in the semantic analysis, perhaps in the following manner:

(4) [The x : Coordinator x]($x = \text{Dabecchi} \wedge [\dots]$)

There are other options for felicitous name introduction as well:

- i) We may use an *indefinite* description, as in “I have a co-worker called ‘Jack.’
Jack is a fine fellow.”
- ii) We may use appositive constructions, such as “we hired a new coworker today,
Jack. Jack is a fine fellow.”
- iii) Tagging, or formal introduction (if the individual is present) as in “that guy
over there is Jack. Jack is a fine fellow,” or “this is Jack.”

In these cases the name is introduced by drawing on already existing links between the name and a dref (or demonstrated individual). On other occasions names are expected to pick up individuals assumed to be familiar to discourse participants, but which have not previously been in focus of a particular discourse, without any explicit extra information. Since ordinary pronouns, e.g. “he” do not allow such uses there are obvious differences between names and ordinary pronouns. Indeed, pronouns used as so-called “pronouns of laziness” often go proxy for names in sentences, since interpreting such pronouns seems to require less cognitive effort than interpreting a name. This is a matter of the nature of the character, however. The character of “he” is “most *salient* male,” which ensures that whatever is predicated of the referent is predicated of the in-focus dref. The character of the name, “the bearer of *n*,” ensures that a name requires somewhat more effort to identify a dref, but the additional cognitive cost is a function of the *benefit* of using names, namely to circumvent in-focus or activated drefs to pick up other, more distant drefs (while ensuring that one is picking out the bare individual rather than the individual

in virtue of a thematic role, as a definite description may do). Such navigation between drefs will plausibly require extra effort.

Another reason to treat names as definite expressions is the fact that names cannot occur in so-called Milsark constructions, which is a standard test for definiteness.

Indefinite expressions allow a “there is/are” preface whereas definite expressions do not; i.e. (5)–(7) are felicitous, whereas (8)–(10) are not:

- | | |
|---------------------------------------|-------------------------------------|
| (5) There is a fox in the hole | (8) *There is the fox in the hole |
| (6) There are some foxes in the hole | (9) *There is every fox in the hole |
| (7) There are a few foxes in the hole | (10) *There is Jack in the hole |

In conclusion, names share with pronouns a familiarity constraint on felicitous use.

3.3.2. Names As Anaphora

The definiteness of names has given rise to at least two influential attempts to construct unified theories of names, pronouns, and even definite descriptions. Elbourne (2006) proposes a static framework where definite expressions are uniformly treated as definite descriptions in a fine-grained situation semantics. According to Elbourne an E-type pronoun has the syntax and semantics of the determiner “the,” which is followed by a phonologically deleted noun phrase. Thus, in “a man entered. He was tall,” the pronoun is [he man], and the deleted noun is present at the sentence’s logical form as “the man.” Names, on the other hand, work the opposite way. “Jack,” on Elbourne’s view, is [the Jack], where the occurrence of “Jack” is (usually) interpreted as $\lambda x(\text{Jack}(x))$ and

understood as “the person called ‘Jack.’” Elbourne treats all definite expressions as E-type pronouns (descriptions). We have already discussed some worries pertaining to such views. “Jack is called ‘Jack’” will have the logical form “[the Jack(x)] is called ‘Jack,’” which is true if the value for x is defined and the value is $\lambda x(\text{Jack}(x))$ in a world w . This sentence, however, is true if the definite description is defined, and cannot be false, even though it is supposedly contingent that Jack is called “Jack.” Since Elbourne lets the descriptive content affect the sentence’s truth conditions the sentence will not behave as it should in modal contexts (see also section 2.2).

Geurts (1997) treats names as presupposition triggers that are semantically bound in a dynamic framework. He defends a version of Discourse Representation Theory in which merging contexts either provides a binder, or the dref has to be accommodated in virtue of a presupposition. (Pronouns and names could also be taken to be just bound variables, if we instead adopted a framework with no intermediate representational layer, such as DPL). Geurts nevertheless treats names as definite descriptions, albeit not Russellian restricted quantifiers, and names and pronouns are taken to be anaphora. In that case, however, indexical pronouns and true demonstratives fall outside the scope of the theory. Indexicals and demonstratives do not appear to occur bound (Zeevat, 1999; Maier, 2009).⁴² Of course, some expressions that are often characterized as indexicals seem to have bound occurrences, such as “actually” or “here” in hypothetical or quantified contexts (“if Obama hadn’t won in 2008, the actual winner would have been Republican”). It is unclear whether “here” or “actually” are indexicals in these cases,

⁴² In “I tended to my garden” the pronoun “my” (the possessive construction is often treated as “the garden of I”) might *prima facie* appear to be bound, but could also – probably better – be explained as a case of co-reference.

however, rather than anaphoric on some situational or contextual variable. In any case, to achieve universality Geurts would need to find bound occurrences of unambiguously indexical expressions such as “I,” but there does not appear to be any.⁴³ There are other attempts at constructing a unified theory; Maier (2009) employs a type of layered Discourse Representation Theory in which referential terms trigger rigid presuppositions. Even this theory, however, requires – and Maier admits as much – certain violations of central principles in special contexts due to extreme pragmatic pressure.

As before, I suggest that the shortcomings of these attempts stem from a wish to cash out the uniformity in terms of *actual* semantic contributions rather than in terms of the level of grammatical requirement allowing a certain pluralism regarding the actual semantic values depending in part on whether the relevant context is linguistic or extra-linguistic. I have argued that the unified account of singular terms is a matter of requiring individuals, and that this can be satisfied by expression occurrences that nevertheless make somewhat different semantic contributions. Admittedly, however, what I offer is, as opposed to Elbourne, Maier or Geurts, a unified view of singular terms rather than a unified view of definite expressions.

⁴³ Partee (1989) suggests that sloppy uses of “I” appear to be bound (“I’m the only one around here who will admit that I could be wrong”); see Maier (2009) for a discussion and rejection of the claim – going into the rather complex details surrounding these cases would take us a little too far afield. Schlenker (2003) suggests that shifted indexicals in attitude reports could be construed as binding indexicals, but these examples (from Amharic) could be even better explained by liberalizing Kaplan’s famous prohibition of monsters (operators that shift contexts); again, see Maier (2009) for a discussion.

3.4. Conclusion

If names are pronouns, or best classified as pronouns, we obtain a unified theory of singular terms that is more flexible than standard direct reference theories. The primary difference between different noun phrases is not a difference in the propositions expressed by sentences containing them but rather a function of a ‘model theoretic’ constraint on what types of values various expressions can take or have assigned to them. Singular terms do not determine their values by satisfaction of any criterion but have them assigned, since they do not allow descriptive conditions or any truth-evaluably relevant distinction between intension and extension. Although the values of singular terms are insensitive to descriptive conditions, we still have flexibility regarding *how* values are assigned. The prime example of a singular term in the unified theory is the variable, which refers directly relative to an assignment but does not refer (but ranges over a domain) in bound occurrences. Referential names and pronouns have values assigned by the interpretation function (and, as such, their values partially define an intended model). Formally, variables have values assigned as well, but bound variables in natural language do not *refer* to the values they receive relative to an assignment (the exact interpretation, using identity functions over domains, is developed in Chapter 5).

In Chapter 4 I argue that not all name occurrences can be interpreted as constants, on pain of yielding the wrong truth conditions of sentences in which they occur. Although many name occurrences satisfy Kripke’s semantic argument against descriptivism, not all do – yet all name occurrences are *de jure* rigid designators and semantically simple. In Chapter 5 I argue that anaphoric pronouns should be interpreted

as variables, and sketch a formal account. In Chapter 6 I then show how we interpret proper name occurrences as variables semantically bound by antecedently introducing (or presupposed) binding expressions and explain how and why proper names and pronouns are associated with somewhat different discourse properties.

Given that names and pronouns are, indeed, associated with different discourse properties one might of course claim that I have already given an obvious difference between names and pronouns, and hence a good reason to resist classifying them together. There are two things I would point out in response, however:

First, the differences are much smaller than one may initially think, and a result of the fact that we have a larger stock of names and can therefore apply them more discriminately than we can pronouns. The very fact that we have a much larger – in principle perhaps infinite – stock of names means that the interpretation of names will automatically be associated with somewhat heavier cognitive demands. But this is not a semantic difference, nor a difference in grammatical role but a difference concerning how we manage our vocabularies. This difference is sufficient to explain their different roles in the givenness hierarchy. The fact that we specially coin names for individuals means that names can reach further back in a discourse, and track drefs even across discourses (as long as the dref is already familiar to the discourse participants). This is why names are useful, and this is the purpose for which we use them. As I argue in Chapter 6, other apparent differences (possibly even the *de se* readings that ordinary pronouns may allow but names do not) can be explained in terms of this practical difference. Among the most significant such differences are i) the fact that names are associated with name-using traditions. A name-using tradition is, on this view, a tradition that maintains an anaphoric

link, potentially across contexts, and possibly by non-semantic (e.g. causal) means as long as there is a suitable (causal) source. ii) As opposed to “he” a name must be explicitly attached to a referent (or dref); “he” attaches to a dref or referent in virtue of the “most salient” inflection in the character. These two differences are closely related. A *baptism* is a speech act that assigns a referent or dref to a name and, under appropriate conditions, establishes an anaphoric link terminating in this referent.

The question of classifying names and pronouns together is, in the end, largely a terminological issue. I am interested in the obvious similarities, such as the fact that both expression types are singular terms as well as the consequences of this (semantic simplicity, *de jure* rigidity, the character–content distinction, correspondence to free variables and – in different contexts – to bound variables or constants). In other words, I think the classification is well motivated and have significant explanatory power.

CHAPTER 4: NAMES AS VARIABLES, NAMES AS CONSTANTS

Like pronouns names are definite expressions associated with familiarity constraints. The proffered content of a name, like a pronoun, is a free variable. In this chapter I argue that, like pronouns, names function both as constants and as variables. Note that throughout the chapter it is crucial to distinguish the *expression* “Vulcan” from the *occurrence* of a name “Vulcan” in a discourse. The expression has a lexical meaning (“the thing called ‘Vulcan’”), but it is *occurrences* of names that receive interpretations in contexts drawing on character context, and any name-using convention available in context. “Aristotle was a shipping magnate” and “Aristotle was a philosopher” involve two occurrences of the same expression, yet the semantic contributions of the occurrences are very different.

The familiarity constraint means that it is as infelicitous to open a discourse with “Jack is a fine fellow” if the audience has no method of finding out which Jack one is talking about, as to open the discourse with “he is a fine fellow” if the audience did not know who one had in mind. Once a dref is sufficiently established, however, a (standard) pronoun remains the preferred way of retrieval; it is, for instance, infelicitous to use “Jack thinks Jack is *F*” when what one has in mind is “Jack thinks he is *F*.” Pronouns pick up the most salient dref by default, whereas names can circumvent drefs to retrieve drefs further back in a discourse. Thus, pronouns require less cognitive effort. On the other hand, names are preferred when several drefs belonging to the same predicate extension are salient. In fact, an important use of names is exactly to keep track of otherwise unmarked drefs, as in “there are two unknown murderers here; let us call them ‘Jack’ and ‘Jill.’” Such cases show one important application of anaphorically used names.

Although presupposition and linguistic context play important roles in resolving the semantics of name occurrences, it is not a given *how* they come into play. Just as the different ways in which pronouns work do not show up at the level of syntax, the difference between what I call *deictic* and *anaphoric* names does not show up at the surface level. Hence they exhibit no difference in how they interact with other elements of a sentence, for instance. There are semantic differences, just as there are semantic differences between deictic and anaphoric pronouns, but by itself the distinction poses no threat e.g. to compositionality. Context, including linguistic context, determines whether the occurrence should be interpreted deictic or anaphoric, but since the expression in either case only takes the kind of value associated with elements of syntactic category *e*, individuals of the domain, we avoid threats to compositionality. It is, however, a mistake to look for a unified way in which these expressions draw on linguistic contexts, antecedents, and presuppositions to determine a contribution to a proposition. Instead, presuppositions can sometimes be a matter of meta-linguistic assignments (demonstratives, indexicals, some name occurrences), and at other times a matter of semantic binding (anaphora, names).

4.1. Introducing Names

In Chapter 3 we argued that names sometimes work as constants. The most clear-cut cases, however, are cases of indexicals and demonstratives. Now, there are two ways to

introduce a name into language.⁴⁴ First, names can pick up individuals introduced by some sort of ostensive, baptismal act. Since demonstratives are directly referential, and any associated descriptive information helps focus the demonstration and is not part of the meaning of the (occurrence of the) demonstrative, a name introduced in this manner has the same content as the demonstrative, the referent. The relationship between the name and the introducing demonstrative is, as the relationship between two occurrences of demonstratives or indexicals, a relation of co-reference and not a matter of binding. Since the demonstrative used to introduce it is structurally simple, rigid, and directly referential, the introduced name – being a singular term – will inherit those characteristics.

Second, names can be introduced by description. In this case, I will argue, the name is anaphoric on the introducing description, and just like an anaphoric pronoun it is semantically represented by a variable bound by the description. The introducing phrase is a denoting phrase, and not referential. Names are definite expressions and thus used to sustain anaphoric chains, but since the denoting phrase is insensitive to what individual actually satisfies its predicative material (regardless of whether the denoting phrase is used *de re* or *de dicto*), then so is any definite expression we use to keep track of that individual. Hence, to understand a sentence containing a name introduced to pick up the referent introduced by the denoting phrase should not require identifying this individual either, any more than an anaphoric pronoun – a singular term as well – does. Compare:

i) Let us call the inventor of the zip, whoever he was, “Julius.” He was smart.

ii) Let us call the inventor of the zip, whoever he was, “Julius.” Julius was smart.

⁴⁴ I emphasize that I am here talking about the original introduction of a name into *language* in the form of a dubbing, which must be distinguished from introducing names to particular *discourses*.

These two sentences should be truth-conditionally equivalent. In the first, however, the occurrence of “he” in the second sentence is not referential, but bound. I claim that in the second sentence, “Julius” is a bound variable as well, rather than a referential term.

When the antecedent of a name is a demonstrative, the character (“thing called *n*”), which yields a referent in combination with a context. Other associated descriptions at most help to point out the correct assignment in a context. The individual assigned will thus in part determine a sentence’s truth-conditions. In the case of descriptive names, the name is equivalent to a discourse anaphoric pronoun, and the truth-conditions of sentences in which it occurs are determined independently of the individual. The character is the same, but the context is linguistic and at most it serves to restrict the class of possible assignments. The designation, if any, is determined relative to a circumstance of evaluation. In 4.2.2 I argue that the distinction between names is based on semantic rather than epistemic considerations.

In general, section 4.2 argues that we need to recognize a distinction between deictic and anaphoric names, and in 4.3 I motivate the claim that names sometimes function as variables. I restrict the discussion to names occurring in the contexts in which they are introduced. As opposed to pronouns, names are able to generate referential traditions extending beyond the context of introduction; these are discussed in Chapter 6.

4.2. The Problem of Descriptive Names

I will, to begin with, focus on the relatively clear cases of what Evans (1979, 1982) calls *descriptive names*, such as “Julius,” introduced by the stipulation

(1) Let us use “Julius” for whoever invented the zip (if anyone did).

Another well known example is “Newman,” introduced by the stipulation “Let us call whoever is the first person born in the 22nd century ‘Newman’” (Kaplan, 1968, 1979). I take it to be relatively clear that descriptive names are names. As their grammatical role suggests, they are singular terms, and have no semantic structure; their role is to supply individuals to the composition of semantic values for sentences in which they occur.

“Julius” is not, for instance, equivalent to the description “the inventor of the zip,” since “Julius invented the zip” is contingent whereas “the inventor of the zip invented the zip” is not, hence “Julius” is rigid. Furthermore, “Julius” does not appear to be equivalent to an actually-rigidified description either. There is nothing modal involved in (1), and it seems to be possible to introduce names in this manner without linguistic resources to formulate modal claims at all. Nevertheless, understanding a sentence containing “Julius” does not require identifying the individual who is actually the inventor of the zip.

The claim that we need to recognize a category of descriptive names distinct from ordinary names has been recognized by Evans (1979, 1982), Davies & Humberstone (1980), Devitt (1981), Davies (1981, 2006), Forbes (1989), McCulloch (1989), Recanati (1993), Sutton (2001), Stanley (2002), Berger (2002), Reimer (2004), Cumming (2008)

and Kanterian (2010), though there is little agreement on the exact nature of the difference. The existence of a relevant difference has been denied by Donnellan (1979), Bach (1987), Salmon (1989, 1990), Soames (1998), Jeshion (2001, 2004), and Sainsbury (2005). One reason to reject a distinction between descriptive (or “anaphoric”) and ostensive (or “deictic”) names is that the distinction might seem to reflect merely a difference in epistemic status. In the case of (1) we do not have (*de re*) epistemic access to Julius, whereas in the case of deictic names we do, at least at the time of introduction. Russellian views tend to be wary of allowing epistemic distinctions to underlie semantic ones, and although *de re* knowledge might be required to fully grasp a proposition expressed by a sentence containing a singular term, Russellians tend to reject the idea that meaning and understanding are constitutively related. Hence, the fact that users of “Julius” are not acquainted with a referent is not necessarily an obstacle to the name being directly referential per se. In this section I argue that the difference between anaphoric and deictic names does not rest on an epistemic distinction. Our epistemic access to certain states of affairs may explain why we need descriptive names, but the difference is semantic. The rest of this chapter is intended to defend this claim.

4.2.1. Julius and the Semantic Argument

As Evans points out we do not use “Julius” in a way that makes the following true (Evans, 1979, p.192):

- (2) If you had invented the zipper you would have been Julius
- (3) If Julius had never invented the zipper, he would not have been Julius

According to Evans we do, however, admit that (4) is true:

(4) If we were to discover that McTaggart invented the zip, then McTaggart would be Julius.

If Evans is right about the falsity of (2) and (3), and we are able to make counterfactual claims about Julius, then “Julius” is sensitive to the Kripke’s modal argument, rigid, and not an abbreviation of a (non-rigid) description. If (4) is true, however, “Julius” is insensitive to Kripke’s *semantic* argument against descriptivism and thus not directly referential. Compare the name “Thales,” which is presumably associated with the description “the philosopher who claimed that everything is water.” The semantic argument purports to show that “Thales” does not mean the same as that definite description (or any cluster of descriptions). If we were to discover that Thales did not claim that everything is water, but that some other philosopher did, “Thales” would not change its reference to this other philosopher. “Julius,” however, is stipulated to be *whoever* invented the zip. If we were to discover that Judson did in fact not invent the zip, then “Julius” would not refer to him, but whoever else in fact *did* invent the zip (if any). Our “Julius”-discourse is tailored to this understanding of “Julius;” that is, our discourse is non-committal with respect to who, in fact, invented the zip.

This is not to deny that “Julius” can come to be associated with Whitcomb Judson when we learn that Whitcomb Judson is, in fact, the inventor of the zip, and hence come to resist assignment changes. My claim is only that insofar as “Julius” *at any point* allows

(4) to be true, “Julius” is not directly referential *at that point*. At least in the context of introduction the name is obviously tied to the description, and someone who doubted that Julius invented the zip in a context where Julius was stipulated to be the inventor, would seriously have misunderstood the point of the stipulation. Several authors (see in particular Reimer, 2004) have offered arguments to support the distinction between descriptive and ordinary (deictic) names for instance by appealing to intuitions about how we would evaluate claims involving e.g. “Jack the Ripper” when used with the intention of picking out someone who is in fact innocent of the murders, how the relation can change over time, and how intuitions about sentences containing “Jack the Ripper” intended to refer to the innocent person are affected when we learn that this person is, in fact, innocent. I do not go into these arguments in detail.⁴⁵ It is unclear what happens once a name gains common currency, especially if the original description is lost. Such intuitions are easily colored by other phenomena, such as reference change, introduction of new tags for familiar objects, the possibility of one name (expression) referring to several individuals, and the existence of pragmatic mechanisms and factors in discourses concerning descriptive names. A convincing case for the distinction between anaphoric and deictic names can be made on independent grounds.

It might be suggested that the intuitive truth of (4) is due to the fact that the modal notions invoked are of a different kind; that is, that (4) invokes an epistemic rather than

⁴⁵ For instance consider “Neptune,” introduced as a name for the purported planet that caused the perturbations of the orbit of Uranus (at least according to the way philosophers of language usually tell the story). Of course, if we were to discover *now* that Neptune is not the cause of those perturbations we would hardly claim that the planet we currently think is Neptune isn’t really Neptune. This shows that “Neptune” is, as of today, an ordinary non-descriptive name. But had anyone made such a discovery right after Leverrier discovered the planet currently named “Neptune,” it is not unlikely that one would have concluded that Leverrier had not discovered Neptune, but something else; see Chapter 6.

metaphysical possibility (Chalmers, 1996). But even though we would naturally employ the stipulation (1) only if we were ignorant about Judson, “Julius” appears to be insensitive to the semantic argument because “Julius” is semantically stipulated to do so by (1), not because we are in an impoverished epistemic situation with respect to the actual inventor. In other words, if there are lessons for epistemology here, they seem to derive from a semantic distinction rather than vice versa, and the peculiarity of “Julius” appears to be a semantic phenomenon.

4.2.2 Semantics Vs. Metasemantics

The last observation requires some elucidation. One reason for rejecting our form of semantic pluralism about names, especially among defenders of a Russellian framework, is that the distinction between descriptive and other names may appear to anchor a semantic distinction in an epistemic one, namely whether we have or lack *de re* knowledge of or acquaintance with the referent. This worry is misplaced. I will argue that the distinction between a descriptive (or anaphoric) and a non-descriptive (or deictic) name does not rely on whether we have *de re* knowledge or not of the designatum, or whether uses of the name are properly causally linked to the designatum.

What we are concerned with is what a name contributes to the truth-conditions of (or proposition expressed by) a sentence in which it occurs. This question must be sharply distinguished from attempting to provide a *theory of reference*. A theory of reference is a theory that purports to answer the question “in virtue of what does a name refer to what it refers to”. A theory of reference is not a semantic but a foundational theory. In the case of the occurrence of a name such as “Aristotle,” for instance, a *causal*

theory of reference (e.g. Kripke, 1971; Devitt, 1974, 2012) says, roughly, that “Aristotle” refers to the individual that is the causal source of that use of the name. Alternatively, one may argue that “Aristotle” refers to Aristotle in virtue of Aristotle satisfying some description (perhaps at the level of character) associated with the name in that particular context. Yet again, it has been argued that it is the existence of an informational link (Evans, 1982), or the availability of *de re* knowledge of the referent (McDowell, 1984), that ensures the assignment of a referent. Finally, according to Horwich (1997), there is no general theory of reference to be given over and above the deflationary schema “*n* refers to *n*.”

I wish to remain neutral with respect to this question. The relation plays no semantic but at most foundational role. My concern is with the truth-conditional contributions of different kinds of names. Importantly, however, treating the question of a theory of reference as orthogonal to the issue at hand places some restrictions on what I can use as evidence for my proposed pluralism. Evans (1982), for instance, claimed that the different contributions of “Julius” compared to, say, “Obama” were grounded in the fact that speakers have different informational links. A use of “Obama” is anchored in Obama in virtue of an informational link to Obama, whereas a use of “Julius” is not – “Julius” is introduced in a context where no proper informational link is available. Thus the semantic difference between the names is sustained in virtue of the absence of a foundational link. In other words, Evans defends the existence of a difference between descriptive and deictic names by appeal to what he claims are different relations of reference; that is, different ways of being directly referential in virtue of different meta-

semantic links. Drawing a semantic difference between various names is for Evans hence justified by non-semantic, and not semantic, considerations.

Since I wish to remain non-committal with respect to a theory of reference (or epistemic or metaphysical issues in general) I cannot argue for the existence of a distinction in this manner. In this chapter the purpose is therefore to show that the difference is justified on *wholly semantic grounds*, namely by reference to the truth-conditions associated with sentences in which these names occur. This is not to deny that epistemic or cognitive needs explain why the distinction exists, nor that we employ special semantic mechanisms in the case of “Julius” *because* we are in a particular epistemic situation with regard to the designatum (perhaps the different semantic resources we employ can even be explained in virtue of the existence or not of informational links to objects that could serve as referents). After all, it is hardly surprising that our cognitive situation should determine which semantic mechanisms we employ or why these distinctions exist in the first place. It is, for instance, primarily a function of our cognitive situation that we employ definite descriptions rather than demonstratives when we do, as shown by the givenness hierarchy of Gundel et al. That does not mean that the semantic distinction between demonstratives and descriptions reduces to epistemic or metaphysical distinctions, but rather that language use reflects cognitive and communicative needs. The primary difference between different kinds of name occurrences may nevertheless be wholly semantic in character.

I argue that the difference between names tracks semantic differences between the expressions used to introduce them, which are sometimes demonstratives, sometimes quantifier expressions. Names are definite expressions, and when the drefs they pick up

are introduced through semantically different types of expressions, the semantic contents of the name occurrences will – as shown by the case of deictic compared to anaphoric pronouns – reflect these semantic differences in some manner. The fact that our epistemic or cognitive situation explains why we used a description rather than a demonstrative to introduce a name in a particular situation does not mean that the difference between the names is epistemic and not semantic. Since names are interpreted in part by drawing on some kind of relation to earlier occurrences (whether these are binding relations or metasemantic links) the introduction – if only the introduction to particular discourses – will matter for how subsequent occurrences are interpreted.

One observation that significantly supports, through purely semantic considerations, the claim that names can occur non-referentially is the fact that names sometimes occur in bound position.

(BN) If a child is christened “Bambi,” and Disney Inc. hears about it, then they will sue Bambi’s parents. (Geurts, 1997, p.322)

In (BN) the used occurrence of “Bambi” cannot be referential if (BN) is to receive the correct truth conditions. Rather, “Bambi” is a bound variable, equivalent to an anaphoric pronoun. Although explicitly bound uses of names are rare, (BN) shows that names *can* legitimately occur bound. So even purely semantic observations strongly suggest that the quest for a unified theory of names has to be given up (this and similar examples will be discussed in section 4.3). The question, then, is not whether names sometimes occur non-

referentially, but how common such occurrences are, and I aim to show that descriptive names such as “Julius” are – based on semantic considerations – not referential.

We have already (cf. (2) and (4)) mentioned the different behavior anaphoric and deictic names exhibit in modal contexts. In 4.2.3. I use the observation that descriptive names give rise to *a priori* contingent truths to argue that they cannot be treated as Russellian, directly referential expressions. The epistemic concerns, however, are not intended to make the distinction we wish to draw an epistemic one. Rather, the *a priori* character of the truths in question is intended to serve as *evidence* for the existence of a semantic distinction – the problem is, to emphasize, that treating descriptive names as directly referential would entail that sentences have truth conditions that these sentences simply do not have, or express propositions that they do not express.

4.2.3. Julius and the Contingent *A Priori*

If “Julius” is directly referential, then by knowing the general proposition *the inventor of the zip invented the zip*, which is surely *a priori* if anything is, one can by (1) infer the singular proposition *Julius invented the zip*. Since “Julius” refers to Whitcomb Judson, and if names are directly referential, then knowing that *Julius invented the zip* is equivalent to knowing that *Whitcomb Judson invented the zip*. And if this is the case, then since *the inventor of the zip invented the zip* is *a priori* and the introduction of “Julius” a mere stipulation, one can know *a priori* that *Whitcomb Judson invented the zip*. One would of course not know that Whitcomb Judson was *called* “Whitcomb Judson,” but this is no part of the singular proposition *Whitcomb Judson invented the zip*. Hence, if descriptive names were directly referential, we would be able to infer all sorts of (indeed

any) singular propositions *a priori*. But *Whitcomb Judson invented the zip* is a contingent proposition about the world that cannot be verified *a priori*. Hence “Julius” cannot be directly referential. We can make the problem a little more precise. Given that Whitcomb Judson = Julius, we can infer (6) from the true proposition (5):

(5) *Julius invented the zip* is knowable *a priori*

(6) *Whitcomb Judson invented the zip* is knowable *a priori*

“Julius invented the zip” and “Judson invented the zip” express the same proposition, so if either is knowable *a priori* then both are.

Notice that the point (and the following argument) does not rely on *Julius invented the zip* being *a priori* in any very strict sense. For our reductio to go through it is enough that our epistemic access to the truth of *Julius invented the zip* (if anyone did) is different from the access we have to *Judson invented the zip*. Notice furthermore that I do not deny that there may be *a priori* contingent truths. Even if such truths exist, *Whitcomb Judson invented the zip* should not be among them. It is, however, easy to overlook the problem this poses for Russellian theories. Superficially, one might mistake it for a variant of Frege’s puzzle, and insofar as there are promising Russellian solutions to Frege’s puzzle (e.g. Salmon, 1986; Richard, 1990; Crimmins, 1992; Soames, 2002) one might think an easy solution is available.

The problem is not a version of Frege’s puzzle, and existing Russellian solutions to Frege’s puzzle cannot provide a solution to this problem of the contingent *a priori*. A careless formulation of Frege’s puzzle may construe it as the problem that since Hesperus

= Phosphorus, and *Hesperus = Hesperus* is knowable *a priori*, it follows that *Hesperus = Phosphorus* is knowable *a priori*. But Russellian solutions to Frege's puzzle do not reject this conclusion; indeed, they endorse it. After all, the proposition *Hesperus = Phosphorus* is the very same proposition as *Hesperus = Hesperus*. So if one of them is *a priori*, so is the other. The problem Russellian theories have to deal with in the case of Frege's puzzle is how even logically omniscient rational agents could fail to assent to the proposition *Hesperus = Hesperus* under one "mode of presentation." The solution is that a subject's grasp of a proposition is not entirely transparent. Russellian solutions to Frege's puzzle usually invoke pragmatic factors or modes of presentations under which propositions are grasped, perhaps to the extent that "believes that" expresses a three-place relation between an agent, a proposition, and a mode of presentation of that proposition (Salmon, 1986); that the truth-conditions of attitude ascriptions involve linguistically enhanced propositions (Richard, 1990); that "that"-ascriptions semantically brings a character or certain contextual factors to the evaluation of a sentence in addition to the embedded proposition; or other pragmatic mechanisms (e.g. Soames, 2002).

In other words, Russellian solutions to Frege's puzzle exploit the semantics of "that"-constructions or belief contexts to solve a puzzle that concerns the substitutivity of identical propositions in hyperintensional contexts. All Russellian solutions accept that "Hesperus is *F*" and "Phosphorus is *F*" express the same proposition. In fact, most Russellian solutions even accept that "Jack believes that Hesperus is *F*" and "Jack believes that Phosphorus is *F*" express the same proposition or at least necessarily have

the same truth-value.⁴⁶ They accept the latter since it follows from the identity of *Hesperus is F* and *Phosphorus is F* and Leibniz's Law; the only way to block the inference appears to be to claim that "that" or "believe that" is or involve (perhaps covert) indexical or demonstrative elements that can be used to pick out certain other aspects of a particular context of utterance.

(5) and (6) do not involve hyperintensional contexts – "knowable" is not an attitudinal operator – but ascribe the same property to the same proposition (and not to beliefs or belief contexts). Neither involves the potentially demonstrative "that." So on a Russellian view (5) and (6) (and not only the propositions embedded in (5) or (6)) should express the very same proposition.⁴⁷ But (6) is false – this is not something any remotely plausible theory could disagree with – and (5) is true. Hence, (5) and (6) cannot express the same proposition. Since the only difference between them is the occurrence of "Julius" and "Judson," respectively, either or both names cannot be Russellian names.⁴⁸

⁴⁶ In fact, a major source of justification for Russellian theories in general is the fact that there *are* situations and readings of "Jack believes that Hesperus is *F*" in which "Hesperus" can be substituted with "Phosphorus" without a change of truth-value. There is no situation in which "Whitcomb Judson" can be substituted for "Julius" in (5) without a change in truth-value.

⁴⁷ Indeed, claiming that any of the factors at work in resolving Frege's puzzle is at work in (5) or (6) seems tantamount to admitting that these factors are truth-conditionally efficacious in non-hyperintensional contexts. Admitting this is tantamount to rejecting Russellianism, and rejecting it for descriptive names only is tantamount to denying that descriptive names are directly referential.

⁴⁸ This is not to deny that "Julius" exhibits behavior *close* to an ordinary Russellian term. Notice that "Hesperus = Phosphorus" is necessarily true. If "Julius" is a name, should not "Julius = Judson" be necessarily true as well? There is a clear sense in which the answer should be "yes," and I discuss the issue in more detail, and devise a framework that precisely predict the differences and similarities between "Judson" and "Julius" in modal contexts, in Chapter 7. The short store is that "Judson = Julius" is indeed necessarily true, but that "Julius," as opposed to "Judson," does not satisfy the semantic argument. So, if evaluated at a different world, "Julius" may be someone else (a different individual, *i*, satisfies "inventor of the zip"), and relative to that world "Judson = *i*" is necessarily true.

4.2.4. Donnellan's Solution

The only way to avoid this conclusion and block the reductio of Russellian views is to deny that (5) is true. Donnellan (1979) argues that although someone who introduces "Julius" by (1) can know *a priori* that "Julius invented the zip" expresses a truth, she cannot know *which* proposition the sentence expresses. This is presumably what we would say about less controversial *a priori* contingent truths involving indexicals. For instance, one can arguably know *a priori* that any utterance of "I am here now" expresses a truth, but one knows this in virtue of knowing the character of the sentence. The proposition expressed by any utterance of "I am here now" is only available to anyone who knows the context. This proposition is contingent, but is also not knowable *a priori*.

According to Donnellan (5) is false since grasping the singular proposition expressed requires *de re* knowledge of the referent, so *Julius invented the zip* is not knowable unless *de re* knowledge of Julius is obtained. Instead, speakers have trivial, *a priori* meta-linguistic knowledge that "Julius invented the zip" expresses a truth in virtue of the fact that "the inventor of the zip invented the zip" expresses a truth if anything satisfies the description. Although Donnellan does not discuss what such knowledge amounts to in detail, it must amount to more than knowing *de re* of the inventor of the zip that he could have failed to have been the inventor, which can be cashed out in terms of quantifier scope (the producers of "Julius" could certainly have *de re* beliefs about Julius in *this* sense). Plausibly, *de re* knowledge involves recognition-based or information-based beliefs (Evans, 1982);⁴⁹ in order to know *de re* of *o* that *o* is *F*, one must have some

⁴⁹ Of course, a Russellian theory of names following Donnellan should be sharply contrasted with Evans's Fregean theory; whereas Evans takes "Julius" to be referential he defends a rather robust view of reference as a relation that can be realized in a variety of ways, that is, he presents a rather sophisticated theory of the

other means of recognizing *o*, not dependent on the fact that *o* is F^{50} (a causal link may be sufficient). The upshot is that instead of a semantic solution to the effect that “Julius is F ” expresses a proposition insensitive to the semantic argument and thus independent of the assignment to “Julius” that makes it true, we get an epistemic solution. Because of our epistemic relation to Whitcomb Judson we do not know what “Julius is F ” expresses (though we know *a priori* that it is something true); the content of the sentence is object-dependent, and access to the content requires empirical work.

The solution has the awkward-sounding implication that someone who utters “Julius invented the zip” does not know what she is saying, but this awkwardness could perhaps be explained away by pragmatic means. More worrisome is the apparent commitment to *necessarily* unknowable propositions, for instance when the introducing description itself involves a possibility-operator. Consider:

That “Aristotle” is a rigid designator can be illustrated by the fact that whereas Aristotle taught Alexander, he did not necessarily teach Alexander. Relative to a possible world in which someone else, call him “Baristotle,” taught Alexander and Aristotle did not, “Aristotle” would still refer to Aristotle, not to Baristotle.

If “Baristotle” is directly referential then Donnellan’s purported solution entails that it would be impossible to grasp the proposition expressed by “Baristotle taught Alexander.”

relationship between meaning and understanding on the one hand and reference on the other. On Russellian views, this variety of reference is relegated to pre-semantics, that is, as conditions on what it takes for a term to have the reference it has and on what it takes to grasp a proposition involving the referent; hence, it is the relationship between understanding on the one hand and meaning and reference on the other that is complex; the (semantic) reference relation itself is not.

⁵⁰ Strictly speaking, what Donnellan needs to provide is merely a criterion of the *absence* of *de re* belief. I submit that such a criterion can still be formulated using Evans’s notion of recognition-based knowledge. An agent fails to have *de re* knowledge of *o* if the agent can identify *o* just by the satisfaction of the description used to introduce “*o*” or information that is entailed by this description (or by the description in conjunction with whatever else the agent knows). Donnellan’s formulation is somewhat weaker than this, but as formulated here the requirement immediately blocks Jeshion’s Oldman example (Jeshion, 2001).

Notice that Baristotle is not a fictional character – if he were, then the sentence “Baristotle could have taught Alexander” would be false. Of course, if fictional characters are possible objects, then Baristotle is a fictional object who could, indeed, have taught Aristotle. But this just makes the problem worse. Instead of providing a solution to how to understand “Baristotle” – it does not; *Baristotle taught Alexander* remains unknowable – it would generalize the present problem – how to understand what we are saying – to all fictional characters as well.

Second, the analogy with contingent *a priori* demonstrative or indexical sentences does not hold. Donnellan’s position is probably best understood as associating “Julius” with an indexical or demonstrative using Kaplan’s *dthat*, a term-forming operator that combined with a definite description “the *F*” forms a term, *dthat*(the *F*), the semantic contribution of which is (only) the denotation of “the *F*,” and transposes the description to the level of meta-semantics or *character* for the resulting term. This description is then supposed to provide a referent for ‘Julius’ to the proposition, even though the referent is epistemically unavailable. But how could one fail to know the proposition expressed? Since the referent of “Julius” *is* the inventor of the zip the Julius-case is not analogous to cases of defective contexts (such as contexts in which “you” is uttered without an addressee, in which case the proposition itself is defective). Furthermore, the Julius-case cannot be a case where the context itself is not sufficiently determinate. If it were, the upshot would just be that the demonstration would make it semantically indeterminate what “Julius” referred to, and that is not the result Donnellan wants, for that seems to suggest that “Julius” is, once again, more like a variable.

In the case of *indexicals*, if *S* utters “yesterday” without knowing what day it is today, the proposition expressed is fully determinate. In fact *S* will have access not only to the character but grasp the proposition *de se* (if not *de re*); in the case of Julius, however, nothing similar becomes available. Now, the Julius-case appears to be closer to utterances involving complex demonstratives than utterances containing indexicals. If *S* hears “that window is a window” (assuming that “that window” is a singular term) she may know *a priori* that the sentence is true if it expresses a proposition at all, even if *S* does not have access to the object demonstrated. But it is doubtful that someone *uttering* the sentence could lack access to the full context – and referent – and still express a determinate proposition. Yet this is supposedly the case with Julius, so even compared to complex demonstratives the analogy does not hold.

Third, since this is a case where speakers have full semantic understanding of the sentence and complete meta-linguistic knowledge of its truth-conditions, Donnellan must reject the following principle (adapted from Jeshion, 2001, p.129):

Accessibility of Content (AC): For all expressions E in L , and all sentences s in L expressing a proposition p , if an agent S understands all the expressions E_s involved in s , then if S was apprised of all the relevant contextual information, S could have an attitude containing p as its content by understanding s .

Now, Russellian views of semantic content reject a strong association between semantic content and cognitive significance, but AC is much weaker than an identification of semantic content with cognitive significance. AC concerns the relationship between full

linguistic understanding and the ability to form attitudes *at all* based on that understanding, and need not entail that we have *full* grasp⁵¹ of or acquaintance with the proposition under all its aspects or modes of presentation; it requires merely that the proposition can be the object of propositional attitudes at all, under *any* mode of presentation. In the Julius case, however, Russellians must deny that we grasp the proposition under *any* mode of presentation – if we did, it would be enough to know *a priori* that Whitcomb Judson invented the zip. Grasping the proposition under a mode of presentation implies grasping the proposition *simpliciter*. Since we know that the proposition is true *a priori* in the case at hand, if we can grasp the proposition under *any* mode of presentation, we know that *Whitcomb Judson invented the zip* is true *a priori*.

Fourth, we do make claims about Julius, express beliefs, draw inferences, make conjectures, and understand what other discourse participants are conveying with sentences containing “Julius,” and we can even give systematic accounts of what goes on in such discourse, yet propositions apparently play no part. As such we would be dangerously close to having a situation in which we are able to make all the moves discourse participants generally do without understanding what we are saying or having

⁵¹ One possible way out may be to argue that propositions are only assented to under modes of presentation (following Salmon, 1986). In that case the context of introduction for “Julius” entails that the proposition expressed by “Julius invented the zip” can only be known (in that context) under a certain mode of presentation, one that does not give rise to the problem that a singular proposition may be knowable *a priori*. In other words, pragmatic factors (contributing to these modes of presentation) ensure that the proposition cannot, in fact, be known *a priori*. This does not so much entail rejecting AC as to claim that understanding a name comes in degrees. It is a little hard to see how understanding the name could come in degrees when the only contribution to the proposition is the referent (certainly there can be no such thing as a partial understanding of a name insofar as a name has no descriptive content as part of its meaning). In other words, Salmon’s view suggests that grasp of *propositions* is a matter of degrees. I cannot do full justice to this view here, but the systematic manner in which these modes of presentation come into play in the case of Julius should be a reason to be wary of excluding these effects from semantic analysis.

beliefs about the contents of the sentences.⁵² Moreover, in order to make (2) and (3) true, discourse involving “Julius” cannot be interpreted merely in virtue of the meta-linguistic truth that “the inventor of the zip invented the zip”; one must also take into account the fact that “Julius” is supposed to work as a *name* with respect to modal contexts (Recanati, 1993, calls uses of descriptive names “non-literal” uses of names). Syntactically, (modal) claims involving “Julius” are *de re* (though given the stipulation “Julius” – as it figures in the interpretation of the *discourse* – is not sensitive to the semantic argument). Hence, the discourse involves the mechanisms of discourses involving ordinary names, but not only do we have to rely on *ad hoc* mechanisms to explain the information conveyed, participants must also understand each others as not providing or obtaining information in terms of the sentence’s literal meaning. Since we can nevertheless interpret the discourse by suspending the idea that “Julius” is assignment-insensitive, we can systematically interpret the discourse without referring to the propositions, and do so without losing anything relevant or add any serious complexities in the process. So why not take the content actually available at face value to be the semantic content?

To conclude, whereas Evans lets non-semantic considerations determine semantic distinctions in the case of descriptive names, Donnellan lets our epistemic situation (with respect to the referent) determine the extent to which we have access to the semantic content of our sentences. A worry, then, is that we are therefore forced to interpret much discourse non-literally that we intuitively seem able to interpret literally. The benefit is a

⁵² There is no possibility for a more general skeptical argument about language here, however. Discourse involving “Julius” is characterized by the fact that participants who are aware of the status of “Julius” do not form singular thoughts, and they are intuitively *aware* that they do not. Donnellan’s view presumably grants this, but assumes that this knowledge is not sufficient to affect the meaning of the sentences the participants utter (though it is sufficient to affect what we convey by uttering them).

unified, Russellian theory of names. My suggestion is that our epistemic situation with regard to an individual at most informs which semantic resources we employ to pick it out (though what resources we use is reflected in later use of the name). The semantic content, however, is accessible at face value. We ensure this by letting “Julius” be interpreted as a variable bound by the antecedent given in (1). In that case our discourses about “Julius” are fully literal, and fully consistent with the Russellian idea that names track individuals independently of whether the individual satisfies any descriptively expressed properties. Furthermore, by associating anaphoric names with bound variables, we are appealing to mechanisms natural language semantics needs to recognize in any case in order to interpret anaphoric pronouns.

4.3. Names As Variables

The fact that certain name occurrences appear not to be directly referential is not itself sufficient to establish that they are variable occurrences. There are, however, considerations in favor of this idea. First of all, names are singular terms, and insofar as this is partially a function of grammar, it would be surprising (and *ad hoc*) if some of them were really covert quantifier phrases. We have supposed that some singular terms (anaphoric pronouns) are best treated as (semantically) bound variables, and that treating singular terms as variables is perfectly consistent with the idea that they are i) non-descriptive, and ii) require individuals (and only individuals) as values. We will discuss this claim, and how to interpret anaphoric pronouns, in some detail in the next chapter.

Moreover, insofar as some singular terms, anaphoric pronouns, already correspond to variables there seems to be nothing objectionable with just stipulating in a context that a given name is to function in the same manner. Consider a situation in which I wish to talk about some chairs and tables in a room I have no demonstrative access to, and suppose I do not even know how many chairs or tables there are in the room. I am nevertheless able to stipulate that an arbitrary chair is to be called “Hans” and an arbitrary table “Jenny,” and then go on to make claims about Hans and Jenny under the assumption that the names pick up arbitrarily selected chairs and tables in the manner a pronoun “it” anaphoric on “a chair” would do. The fact that “it” would be ambiguous between Hans and Jenny in the following discourse even gives me a reason to introduce such names as specialized anaphoric pronouns. Merely *stipulating* that a term is a *constant* seems to be a much more problematic matter, since – given the definition of a constant – identifying the truth conditions of subsequent claims would require that I could assign a referent to the expression; that is, if I am to fully know the truth conditions of my subsequent claims, and if I am to be able to determine the intended model for the claim, I would, as part of that process, have to identify the value.

Second, some names seem to resist Kripke’s semantic argument. If an expression was assigned a value prior to the evaluation of sentences in which it occurs such reassignment is impossible without changing the meaning of the term, in accordance with our definition of “variable” and “constant” in Chapter 3. Formally, the value of a constant is part of defining the intended model, so such expressions are immune to the kind of reassignment that (4) appears to call for. Third, the close relationship between pronouns

and names, and the fact that pronouns do work as both constants and variables, suggest that treating descriptive names as variables is at least an option worth investigating.

In fact, doing so will elegantly provide a solution to the problem of the contingent *a priori*, as well as solutions to several other well-known problems surrounding the semantics of proper names. I show this in Chapter 7. In this section I provide further evidence that some names, in particular descriptive names, should be interpreted as variables. Names can have non-definite and non-uniqueness-implying antecedents (in fact, our example of “Jenny” and “Hans” above already suggests this), and there is evidence that names can occur bound.

4.3.1. Names with Indefinite Antecedents

If names can also be introduced without a uniqueness implication they cannot be analyzed as (rigidified) definite descriptions even on Fregean views of names. On Fregean views names are functional constructions that have individuals as values, but a well-defined function cannot have more than one value (relative to an evaluation). Names introduced by indefinite expressions would have non-definite and arbitrary extensions. Since they are singular terms, they appear to demand being treated as variables (rather than descriptions). Consider first:

Suppose, during World War I, that the Germans had decided to make an attack during the stalling. In that case, they would have easily overrun the first defense lines. Suppose one soldier, let us call him “Hans” had, accidentally, made a reconnaissance mission (or even just peeked up from the trenches once). Then the Allies would have lost the war and Hans would have become a national hero ... [you may freely add further claims about Hans]

This is a counterfactual scenario rather than fiction or pretense. “Hans,” however, is not introduced by a definite description. There is no “one” Hans; any German soldier could have played the role of “Hans.” Neither is Hans a hypothetical figure; “Hans” is any *actual* German soldier appearing in a counterfactual situation. Who fills the role of Hans varies between the worlds (of evaluation) in which the scenario obtains (though “Hans” is nevertheless *rigid*, since whoever “Hans” designates with respect to a world does not have the properties used to introduce the name necessarily, cf. (2) and (3)).

Cumming (2008, p.535) provides another example:

(7) There is a gentleman in Hertfordshire by the name of “Ernest.” Ernest is engaged to two women.

(7) would usually be uttered with a specific individual in mind, but it might not (imagine contemplating information from a census list). But even if uttered with a specific individual in mind (call him “ α ”), (7) would presumably be true even if α was a confirmed bachelor but his neighbor, also named “Ernest,” was engaged to two women. Similarly with (8) and (9) (adapted from Geurts, 1997, p.322, and Cumming, 2008, p.536, respectively):

(8) Perhaps Mary has a son called “Ernest” and perhaps Ernest is the thief.

(9) [fortune teller:] You will meet a man named “Ernest.” Ernest will bring you happiness.

In (9), any happiness-mongering man called “Ernest” who meets you will make the fortune-teller’s prediction count as true. (8) involves a modal operator that appears to require that the proposition expressed is a general one.

If (7)–(9) have general truth-conditions, which is hard to deny, then a direct reference interpretation is doomed. Whereas ordinary names may be interpreted as constant functions from worlds to individuals, “Ernest” cannot, since, as Cumming points out, there could be two gentlemen called “Ernest,” α and β , who are engaged to two women. If w is a world in which this is the case, then $f(w)$ must denote either α or β (on pain of being a function at all). But (7) should be true at w , not undefined. Consider also:

(10) There is a gentleman in Hertfordshire by the name of “Ernest.” Ernest is engaged to two women and is tall.

(11) There is a gentleman in Hertfordshire by the name of “Ernest.” Ernest is engaged to two women and is not tall.

If α is tall in w , then (10) should be true at w , and if “Ernest” is a proper name, $f(w) = \alpha$ (i.e. α is the witness to the existential quantifier). If β is not tall in w , then (11) should be true at w , and if “Ernest” is a proper name, $f(w) = \beta$ (and β is the witness to the quantifier). If “Ernest” were directly referential, then (10) and (11) could not both be true with respect to w , even though they should be. Note that classical descriptivist views fare no better. Since they would also predict that only one of (10) and (11) could be true, the denotation would still be determined by a function from contexts to a *non*-constant

function from worlds to individuals. But “Ernest” cannot be associated with any (well-defined) function.⁵³

It might appear to help to make the input to the functions something more fine-grained than worlds. In situation semantics (Barwise & Perry, 1983) one could provide an analysis of names in terms of functions from situations rather than worlds to individuals. This move is familiar from analyses of incomplete definite descriptions. For instance, with respect to “Jack fetched a book from the shelf and a pen from the drawer, and put the book on the table” any book from the shelf would make the second conjunct true. But even if definite descriptions imply uniqueness, one could (putatively) obtain the right truth-conditions by relativizing uniqueness to the possible situations pairing Jack with an (arbitrary) book from the shelf and existentially quantify over these situations. For (7) relativization to situations enables us to associate well-defined functions with the names. In general, any occurrence of a name, regardless of the quantificational force of its antecedent, could in this manner be analyzed as picking out a unique individual by a well-defined function relative to each situation, and we could say that “Ernest” refers to the individual who verified the sentence with respect to a situation. Hence, even if both (10) and (11) are true, and the two occurrences of “Ernest” have different referents, they are still both referential since we need only refer to one at the time.

But the move appears to get the cart before the horse with respect to direct reference views, since on direct reference views names are referential (*de jure*), not “satisfactional.” According to the idea at hand the referent of “Ernest” is dependent not

⁵³ “Ernest” could in principle be associated with a choice function defined for the introducing predicate, but insofar as this would treat the denotation of “Ernest” as entirely arbitrary (and the sentence as existentially closed) it would precisely amount to rejecting the idea that sense determines reference; the difference between “Ernest” and ordinary referential names would still remain.

only on the introducing antecedent but on who satisfies the second conjunct of the second sentence, so the referent of “Ernest” cannot even be *determined* by the antecedent. There may be multiple situations that satisfy the first sentence, so the antecedent cannot determine the referent. Satisfaction of “Ernest” must instead be relativize to the various consistent extensions of the (minimal) situations that satisfy the antecedent. Insofar as definite descriptions are uniqueness implying, the names in (10) and (11) cannot have their referents even *determined* by definite descriptions. There is furthermore ample evidence that descriptions cannot be referential in anaphoric construction. To obtain the right interpretation of (12) or (13), for instance, we need to quantify over situations, in which case the definite descriptions cannot be referential:

(12) Every note was even more gorgeously played than the note before it.

(13) If I buy a goat and a syringe, Harry will vaccinate the goat with the syringe.

Both (12) and (13) can presumably be made sense of using quantification over situations or events⁵⁴ (though (12) is tricky). In (12) the quantification is universal; it requires that every situation consisting of a temporally ordered pair of notes is such that the later note is more gorgeously played than the previous. Conditionals such as (13) are traditionally interpreted as universal quantifications over (minimal) situations. (13) says that every situation consisting of me and any goat and syringe I bought is true if embedded in a

⁵⁴ (13) in particular may be used as an argument against the uniqueness implication of definite descriptions to begin with since it should be true, not false, if I buy two donkeys and Harry vaccinated both. Yet if definite descriptions are uniqueness implying, then this situation would make the antecedent true and the consequent false (since there was no unique goat that Harry vaccinated). We may disregard this complication, however (see Haraldsen, forthcoming), since it would surely make the situation worse, not better, for the Russellian.

situation extended by Harry and the syringe and his vaccination of the goat in the situation. “The goat” and “the syringe” would have unique denotations relative to each situation but neither is *referential* since the truth-making assignment to the expressions would change from situation to situation, and assessing the sentence for truth requires evaluating all situations. Thus, the occurrence of the definite descriptions cannot be interpreted as referential since it must determine an individual for each situation whereas assessing the sentence depends on *all* the selections/assignments.

To turn this observation into a full-blown objection to the claim that some names are not introduced by referential descriptions, and even to descriptivist views of names, we need to find descriptive names introduced in conditionals or universally quantified sentences as well. We have introduced one such example, (14), already:

(14) If a child is christened “Bambi,” and Disney Inc. hears about it, then they will sue Bambi’s parents. (Geurts, 1997, p.322)

(15) Every woman with a husband called “John” and a lover called “Gerontius” takes only Gerontius to the Rare Names Convention (Elbourne, 2005)

One may obtain well-defined functions from situations to individuals in both cases, but that does not make “Bambi” or “Gerontius” referential; “Bambi” denotes *any* potential child(ren) named “Bambi.” At best “Bambi” as it occurs (used) in (15) picks out a single individual in each situation; but (15) requires that all of them satisfy the consequent ((15) is trivially true if there are none). Evaluating the sentences requires quantifying over situations. Hence, these names are not referential on any account, but bound variables.

Situations in which proper names are relatively obviously bound are rare (perhaps outside of mathematics). This is not surprising given that names have different activation conditions than pronouns – even in (7) it would be more natural to use “he” than “Ernest,” given that activation and givenness constraints (Gundel et al., 1993) call for using the ‘lighter’ pronoun⁵⁵. In other words, well-known constraints on discourse semantics generally disallow names having binders explicitly occurring in the same sentence, regardless of the question of direct reference. However, given that the constructions are legitimate, it is possible to use names in a non-referential manner and, indeed, in a manner that seems to require the variable interpretation⁵⁶. Hence, names *can* function as pronouns (and hence variables, and thus exhibit the behavior exhibited by (2) and (4). This treatment fits “Julius” far better.

4.3.2. Alternative Accounts: Evans and Berger

The idea of associating descriptive names with anaphora has been defended by Evans (1982) and Berger (2002). Neither Evans nor Berger associates such names with variables, and neither succeeds in both equating descriptive names with anaphora and achieving a defensible view of anaphora. Evans takes descriptive names to be referential terms, the values of which are given by rigidifying the introducing description, since they “belong to that category of referring expressions whose contribution to the truth conditions of sentences containing them is stated by means of the relation of reference” (1979, p.180). Evans operates with a much more liberal notion of reference than

⁵⁵ See Zeevat (1999) for a discussion. A possible reason is that names do not allow for *de se* readings (see Cresswell & von Stechow (1982) and Chierchia (1989))

⁵⁶ Geurts provides some admittedly less compelling examples as well. “In the US Leslie can be both a boy and a girl” seems to equivocate on whether “Leslie” denotes the name or its bearer.

Russellian views, however; with reference being whatever relation obtains between expressions and objects that makes the following principle true (Evans, 1979, p.184):

(16) If $R(t_1, \dots, t_n)$ is atomic, and t_1, \dots, t_n are referring expressions, then $R(t_1, \dots, t_n)$ is true iff \langle the referent of t_1, \dots the referent of t_n \rangle satisfies R

Ordinary names are Russellian singular terms, the understanding of which requires having an object-dependent, *de re* thought of the referent. On his truth-theoretic view ordinary names are associated with truth-theoretic axioms of the form (17)

(17) The referent of “Aphla” = Aphla (1982, p.49).

However, Evans takes the existence of descriptive names to show that there exists a variety of ways that singular terms can be related to their referents by R . For descriptive names the axiom from which its truth-theoretic contribution is defined can be stated as:

(18) $\forall x$ (the referent of “Julius” = x iff x uniquely invented the zipper) [*ibid*, p. 50]

(18) does not yield an object-dependent sense, but the reference relation in (17) and (18) is taken to be the same, and the principles use this relation to state the semantic contributions of the name.⁵⁷ Evans’s view of anaphora occurring outside the scope of

⁵⁷ Evans is thus not able to distinguish *de jure* and *de facto* rigidity; actually-rigidified descriptions would count as referential, i.e. satisfying (16) on his view. Now, the difference between ordinary names and

their antecedent quantifiers is exactly parallel. But donkey pronouns are not referential. If they were, “every farmer who owns a donkey beats it” would be false if any farmer owned more than one donkey. Furthermore, “it” is not sensitive to the semantic argument, for if we evaluate the sentence with respect to a world in which John owns a *different* donkey that he beats, it should still come out true. Insofar as names *are*, the analogy between descriptive names and anaphora breaks down (we discuss these matters further in Chapter 5).

Berger (2002, 2006) distinguishes between what he terms *F-type* and *S-type* (“focusing-type” and “satisfying-type”) names. The referents of *F-type* names are determined by speakers’ focus in the manner of demonstratives; associated descriptions serve to facilitate focus and do not affect the semantic content. *S-type* names, such as “Julius,” are determined by descriptions used attributively, but are nevertheless rigid designators of the denotations of the associated descriptions. The same idea underlies his account of discourse anaphora, in which he applies an enriched first-order language to accommodate intersentential binding.⁵⁸ Roughly, discourse anaphora are interpreted by

descriptive names occur at the level of axioms, and (presumably) not at the level of the truth-theorems derived from these axioms. It may thus be tempting to think that for ordinary names the *interpretive* axioms, that is, homophonic axioms of the kind “‘Hesperus’ refers to Hesperus” display the sense of the name whereas the extensionally equivalent axiom “‘Hesperus’ refers to Phosphorus” does not. Since descriptive names such as “Julius” do not have interpretive homophonic axioms (“‘Julius’ refers to Julius” does not display the sense of ‘Julius’), the difference occurs at the (meta-semantic) level of sense, and the reason “‘Julius’ refers to Julius” is not an interpretive axiom is shown by the fact that “Julius” may be an empty name without that affecting its meaning. But this does not quite seem to be the line defended by Evans, however (Kanterian, 2010, develops a view that locates the difference at the metasemantic level); rather, Evans claims that the assertoric content of a descriptive name is given by the description introducing it. Evans would presumably *deny* that the truth theorem “Julius is *F*” is true iff the inventor of the zip is *F*” is derivable from the axiom (18). The thing to notice is that in (18) the quantified statement is *used* – it provides the assertoric content of “Julius;” it is not supposed to provide a substitution rule. Hence, taking Evans to associate “Julius” with a *character* “the inventor of the zip” and as equivalent to *dthat*[the inventor of the zip] is not too far off.

⁵⁸ The interpretation of sentences containing anaphoric pronouns require binding by the antecedent quantifier outside the sentence containing the pronoun, and Berger achieves this by introducing a special

indexed variables bound by quantifier expressions retrieved from previous discourse, augmented with mechanisms to ensure the proper truth conditions for more complex anaphoric relations.

However, Berger is never explicit on the exact role of the descriptions in play, and wavers between treating them as part of semantics and as metasemantic presuppositions. Although *S*-type names correspond (roughly) to descriptive names Berger allows that “a speaker may never be aware of what actual description, attributively used, was initially employed to fix the reference” (*ibid.*, p.8). Since speakers can refer without knowing the description, and may associate the referent with different descriptions, the distinction between *F*-type and *S*-type names seems non-semantic, in the sense that the semantic contributions of the names to propositions are the same. But in that case the parallel between such names and anaphora, which are (semantically) variables, breaks down.

Admittedly, Berger says that if an antecedent quantifier expression “some *F* is *G*” is false, then a sentence anaphoric on expressions in this sentence, such as “she is *H*” cannot be true. In fact, he says that “the truth of the [antecedent sentence] is a *presupposition* for sentences anaphoric upon expressions in it to have a truth-value” (*ibid.*, p.204; my emphasis). In that case it may seem as if the antecedent works as a reference-fixer for subsequent anaphoric expressions. But this is demonstrably incorrect, even

clause in defining what it is for a formula in a sequence (i.e. a discourse) to be true relative to assignments of values to variables. Notice, however, that Berger provides a formal framework for such anaphora but does not attempt to give an account of the mechanisms by which natural language sentences determine the formal counterparts that are supposed to represent their truth conditions. As Soames (2006) points out, this suggests that the clause may be superfluous if there was a translation algorithm from natural language to the formal system that could produce the desired results directly, in which case intersentential binding could be eliminated (this is surely what happens in Salmon’s (2006) similar system, which is explicitly inspired by Berger’s). An obvious worry is that the clauses would be very complex and interdependent (in “every boy loves some girl. She dismisses him,” the translation clause for “him” would depend not only on the antecedent, but also on the translation clause associated with “she”), which at best is a threat to compositionality.

according to Berger's own official view. "She is *H*" should be interpreted as $\exists x_1(Fx_1 \wedge Gx_1) [Hx_1]$, and this sentence is false, not truth-valueless. Thus, if Berger views anaphora as referential terms that get their referents fixed by presuppositions generated by their antecedents, then his view is no more successful than Evans's. If they are bound variables, then they are not parallel to Berger's official treatment of *S*-type names, for which the antecedent descriptions are meta-semantic presuppositions rather than semantically efficacious binders.

4.4. Conclusion

In this chapter we have argued that although some name occurrences are referential and best interpreted as constants, other occurrences are best treated as equivalent to anaphoric pronouns. In the next chapter we will then show that anaphoric pronouns are best interpreted as (semantically) bound variables. Whereas deictic names appear to be best understood as semantically equivalent to and co-referential with the demonstratives introducing them, anaphoric names appear to be best understood as semantically equivalent to anaphoric pronouns, and are, as anaphoric pronouns, related to their antecedents by binding. This interpretation is in line with the unified framework presented in Chapter 1–3; names are singular terms, with no semantic structure, and do not allow e.g. scope variations.

Now, since "Julius" is a variable bound by a wide scope quantificational antecedent it is rigid, and "Julius could have failed to invent the zip" is true. Nonetheless, "Julius" does not satisfy Kripke's semantic argument. Although Judson invented the zip,

and Julius is identical to Whitcomb Judson, the non-modal sentence “Julius invented the zip” is true if evaluated at a world in which someone else, McTaggart, invented the zip and Judson did not. “Julius” is hence not immune to reassignment (we return to this point in chapter 7). Now, some take the possibility that Julius is McTaggart strictly as an epistemic rather than metaphysical possibility (Soames, 2005). But this is problematic. The fact that “Julius invented the zip” is unaffected by the semantic argument appears to be an upshot of semantic distinctions rather than vice versa. “Julius” is insensitive to the semantic argument because “Julius” is stipulated to make sentences containing the name true or false regardless of which individual in fact invented the zip. As in the case of “Hans the chair” and “Jenny the table” above, the name appears to be designed to avoid any commitment to any particular person, despite the intention to make further claims about him or her. Even though our cognitive situation is presumably the reason we use an anaphoric name rather than a demonstrative (no demonstration of Judson was possible in the name-introducing situation), the distinction is semantic, because it depends on the semantic distinctions explicit in the introduction.

Treating names as bound variables raises some obvious worries, however. In the present chapter we have restricted our discussions to names occurring in the contexts in which they are introduced. Names, however, might survive beyond the context of introduction. In the case of deictic names this has been accounted for in some detail (Kripke, 1972, for instance). In the case of anaphoric names, however, the existence of name using traditions may seem much more problematic. This worry will be dealt with in chapter 6, where we also discuss various other questions raised by our semantic pluralism about names.

4.5. Appendix: Arbitrary Names

One of the most intuitive arguments for the idea that descriptive names differ from ordinary names is the fact that the transition from “someone invented the zip” to “Julius invented the zip” looks like a straightforward case of existential elimination. The resulting singular term would in that case be akin to an instantial term or a “flagged constant” in certain formal proof systems. At least insofar as the antecedent description may be indefinite, interpreting the term as a directly referential name would thus appear to correspond to a rather basic fallacy.

Instantial terms, or arbitrary constants, are terms introduced by removing the quantifier head and selecting an arbitrary member from the class of the predicates in the range of that quantifier head. They are not ordinary constants. They can for instance not figure outside of subproofs, and there is a (logical) guarantee that the value of the term satisfies the introducing description (any doubt about this would be a category mistake). Instantial terms are, grammatically speaking terms, however, and should thus semantically speaking be terms as well. Instantial terms are also used in reasoning in natural language, and they seem to be prime examples of descriptive names.

In order for the rules of Existential Elimination and Universal Introduction to be applicable it is crucial that instantial terms are interpreted as arbitrary, and existing analyses of such terms (King, 1987, 2004; Fine, 1985) have treated them as very different from ordinary referential terms. Breckenridge & Magidor (2012) attempt to apply what amounts to an epistemic solution, according to which instantial terms are referential, but

we cannot in principle know which objects a given instantial term refers to – after all, they are stipulated to be arbitrary (Breckenridge & Magidor take reference to be a brute semantic fact, so the fact that we do not know which object an instantial term refers does not mean that we cannot refer to it). Note also that their account is supposed to capture the semantics of instantial terms as they figure in ordinary language reasoning, not to make substantial claims about logic.

Universal Introduction from $F(a)$ is warranted when an instantial term is used if we can demonstrate that an arbitrary F has some property. If “ a ” is an instantial term introduced by quantifier elimination on a formula $\forall(\Phi(x))$, in order to infer that $\forall x(\Psi(x))$ it is not enough that a in fact satisfies Ψ . We need to *prove* that a satisfies Ψ – and insofar as a is a completely arbitrary member of the extension of Φ , proving that a satisfies Ψ is sufficient to prove that all members of the extension of Φ satisfy Ψ . Universal Introduction in ordinary reasoning will thus be what Prawitz (1965) and Fine (1985) call “an improper rule of inference.”⁵⁹ Of course a will presumably have several properties not shared by other members of $\{\|\Phi\|\}$, for instance G . However, since a is arbitrarily selected “ Ga ” would be an unknowable and unprovable truth, and one that could not be legitimately inferred.

An obvious problem with Breckenridge & Magidor’s account is how to deal with reductio proofs with existential premises; that is, where the existential premise is known to be false. Although Breckenridge & Magidor recognize the problem they “suggest that

⁵⁹ Typical examples of improper rules of inference are Conditional Introduction ($\{\phi \mid \psi\} \mid \vdash \phi \rightarrow \psi$) and Necessitation ($\{\emptyset \mid \phi\} \mid \vdash \Box\phi$). In the case of the former, for instance, one supposes ϕ , proves ψ , and infers $\phi \rightarrow \psi$ while discharging the assumption. What we are doing (at least in practical reasoning) is not straightforwardly inferring $\phi \rightarrow \psi$ from the premises, however; instead, we seem to infer $\phi \rightarrow \psi$ from the *fact* that we have shown that ψ follows from ϕ .

the problem [...] is part of the more general problem of empty names, which would need to be addressed independently.” But the problem is not quite independent; their view effectively rules out solutions that interpret empty names as non-referential, as well as theories that take empty names to refer to abstract objects (since these exist). In fact, even appeals to *gappy* propositions (Braun, 1993, 2003) are ruled out – on such view sentences containing empty names do not express complete, truth-evaluable propositions, but convey certain general propositions in virtue of descriptions (contextually) associated with the name that can be used to interpret the speaker rather than the sentence. Applied to instancial terms if they are referential, it would entail that reductio proofs with false existential premises are literally meaningless. Hence, Breckenridge & Magidor’s account is not innocuous with respect to how empty names should be treated.

More importantly, instancial terms cannot *quite* work the same way ordinary names do. Now, the necessitation rule can be interpreted much the same way as universal introduction. Let Φ be a formula. If one can prove that Φ holds for an arbitrary world w , then – since w was arbitrary – it follows that it holds for any world, i.e. that $\Box\Phi$ (assume S5 to avoid unnecessary complications). What, however, would Breckenridge & Magidor be committed to if the proof of Φ involved instancial terms? Take a logical truth:

$$(16) [\exists x(Fx \wedge Lxa) \wedge \forall y(Fy \rightarrow Gy)] \rightarrow \exists x(Gx \wedge Lxa).$$

In the course of a (conditional) proof of this sentence we must select an arbitrary F , say b , that is related by L to a and prove that b is G , and therefore that there is a G that is related by L to a . There is a restriction on the selection, however; we must choose an element

from the class of F s (since existential introduction to $\exists xFx$ is a valid step at any point within the subproof). Let a be Gareth Evans, F be *is a philosopher*, G be *is smart*, and L be *agrees with*). Then the proof may go as follows:

PROOF: Assume that at least one philosopher agrees with Evans and all philosophers are smart. Let “Jack” be an arbitrary philosopher who agrees with Evans. Since Jack is a philosopher and (given our assumptions) every philosopher is smart it follows that Jack is smart. So, it follows that at least one smart person agrees with Evans. Given that it follows from our assumptions, we can infer (by conditional introduction) that *if a philosopher agrees with Evans and all philosophers are smart, then at least one smart person agrees with Evans* is a logical truth.

The argument should be rather trivial. Jack was arbitrarily selected from the class of philosophers who agrees with Evans. The problem with respect to Breckenridge & Magidor’s solution is that the class of philosophers who agree with Evans is not constant across worlds. In fact, if Jack is the same individual across possible worlds, which he is if “Jack” is referential, then the proof is illegitimate – there are worlds in which Jack is not a philosopher (whoever he may be, and regardless of whether he is a philosopher at *some* arbitrary world). We were supposed to choose an arbitrary world, and now this requirement seems to run into conflict with the requirement that Jack is an arbitrary philosopher in that world (indeed, if false existential assumptions can be accounted for by whatever solution one might find to the problem of empty names, then worlds where there are *no* philosophers are less worrisome than worlds in which there are some, but Jack is not among them).

Now, whatever world we select for our premises will give us the correct conclusion, so we will obtain the right conclusion regardless of which world we select. The worry is that we cannot evaluate the *proof itself* relative to other worlds; that is,

PROOF is not a proof relative to all worlds⁶⁰ if “Julius” refers to the same thing at all worlds. PROOF is thus not a *proof*, but at best a *proof schema*; we can run syntactically indistinguishable proofs with respect to each possible world, but at each world we would have to make a new selection from the philosophers in that world – if “Jack” denoted the *same* object, the proof would be inadequate with respect to a world w in which Jack is not a philosopher, since with respect to w the existential elimination would be just as unwarranted as the inference from “there are philosophers, therefore Napoleon is a philosopher” would be with respect to ours. For PROOF to be a proof and hence actually warrant necessitation “Jack” must be an arbitrary philosopher with respect to each world; therefore, if “Jack” is referential and nothing is a philosopher with respect to all worlds in which there are philosopher, “Jack” must refer to different objects in different worlds.

Contrast this with “Evans.” Insofar as names are directly referential rigid designators, “Evans” receives the same value regardless of world. The value of “Evans” is given by the valuation function for the model (of the proof). “Jack” is not. “Jack” gets reassigned for each world with respect to which we evaluate the proof, as an arbitrary member of the domain of philosophers, a domain that varies between worlds. This means, of course, that “Jack” behaves somewhat like a variable; “Jack is a philosopher” is true if there is any assignment to “Jack” that satisfies “is a philosopher” (or for which a variant

⁶⁰ The worry is particularly pertinent since Breckenridge & Magidor distinguish sharply between improper and proper rules of inference. Necessitation is an improper rule. After all, it is generally not the case that if ϕ is true, then so is $\Box\phi$. Rather, we can infer the latter since there was a *proof* that the former was true. So the proof itself (or at least the fact that ϕ was proved) plays a role at least in the semantic rule of necessitation as it figures in ordinary reasoning. Fine illustrates it by distinguishing two roles that every line in an argument or proof should simultaneously serve (1985, p.71): “On the one hand, the statement made, call it φ , will have been *inferred* from the suppositions used, call them Δ . On the other hand, it will have been *demonstrated* that φ can be inferred from the supposition Δ . Although the two goals are distinct, it should be noted that in achieving one I have, in effect, achieved the other. If I succeed in inferring φ from Δ , then that very inference can serve as a basis for demonstrating that φ is inferable from Δ .”

assignment satisfies “is a philosopher,” and the assignment is relative to the world of evaluation). Thus, instancial terms are different from ordinary names. Since “Jack” is not immune to reassignment when the whole proof is evaluated at a different world (although “Jack” is rigid relative to any particular assignment), “Jack” fails Kripke’s semantic argument, and “Evans” does not.

CHAPTER 5: PRONOUNS

In the present chapter I briefly investigate the properties of pronouns. Pronouns are, like names, singular terms and have very little (if any) descriptive content. Yet pronouns are able play different semantic roles; they occur both as demonstratives, in which case they are directly referential, and as anaphora, in which case (I argue) they must be interpreted as (equivalent to) bound variables. As such, pronouns yield prime examples of the semantic pluralism I suggest we find in the case of names. I will first argue that pronouns must be treated as bound variables (relative to our definition of “variable” in chapter 3), and then discuss how we may achieve binding across sentence boundaries.

5.1. E-type Pronouns

Pronouns that may occur as demonstratives can also occur in anaphoric position:

- (1) Jack beats the donkey he owns
- (2) No farmer beats the donkey he owns

In (1) “he” is a *pronoun of laziness*. It is shorthand for “Jack,” and refers to Jack. In (2) “he” is not referential. To interpret (2) one may treat “no farmer” as a restricted quantifier and “he” as a bound variable. Variable binding is subject to rather severe syntactic constraints, most importantly that a bound element must be c-commanded by its antecedent. Since the syntactic structure of (2) is (somewhat simplified) [S[_{NP} No

farmer][_{VP} beat [_{NP} his donkey]]], in which “he” is c-commanded by “no farmer,”⁶¹ this constraint is satisfied. *Discourse anaphora* are different:

(3) Jack owns a donkey. He beats it.

In (3) “it” is not c-commanded by “Jack.” If “a donkey” could bind across sentence boundaries, then (4) should be felicitous, but it is not:

(4) Jack doesn’t own a donkey. He doesn’t beat it.⁶²

In fact, it should be equivalent to “Jack doesn’t own a donkey that he beats.” But (4) does not have this reading. The issue arises even in the context of single sentences where the antecedent fails to c-command the anaphoric pronoun:

(5) Every farmer who owns a donkey beats it.

⁶¹ Strictly speaking, if variable binding is happening here it may be preferable to say that pronouns of laziness, are bound, e.g. by i) treating “Jack” in (1) as a type ⟨1⟩ quantifier with binding potential in the manner of Montague; or ii) associating some coordination condition with the definite “he” that relates it to an antecedent, e.g. [The x : *Salient male*(x) \wedge $x = y$] where y is a free variable, in this case taken by Jack.

⁶² In Geach (1967) discourse anaphora are syntactically bound by their introducing quantifiers. There is a multitude of reasons why this cannot work, in addition to examples like (4). First, compare

(4*) Only few infantrymen survived the Russian campaign. They returned to France.

(4**) [Only few x : Infantryman: x] [x survived the Russian campaign \wedge x returned to France]

If “they” is bound by “few infantrymen,” then (4*) should be interpreted as (4**). But (4**) is not equivalent to (4*), since (4**) is true in a scenario in which most infantrymen survived but few returned. Second, if Geach were right, then quantifiers need to (potentially) scope whole texts. But even if the meaning of sentences depends on previous moves in the discourse, Geach’s treatment entails that the second sentence in (4) does not receive determinate truth conditions before the text is complete (i.e. subsequent occurrences of anaphora help determine how to interpret earlier sentences containing the relevant anaphora), which is absurd. The *semantic* binding discussed later avoids these problems.

Since the syntactic form of (5) is, roughly, [_S[_{NP} Every[_N farmer who owns [_{NP} a donkey]]] [_{VP} beats it]], there is no (systematic) way to get the deeply embedded “a donkey” – which appears in a subordinate position – to c-command “it.” The problem becomes vivid when we represent (5) in predicate logic:

$$(5^*) \forall x \forall z ((\text{Farmer}(x) \wedge \text{Donkey}(z) \wedge \text{Own}(x, z)) \rightarrow \text{Beat}(x, z))$$

How do we get from (5) to (5*), its logical form, in a principled manner? First, the indefinite in subordinate position must rise to take wide scope over the sentence, and must, in some *ad hoc* manner, acquire universal force. And even then “it” would not be a variable bound by a prenex existential quantifier “a donkey,” as syntax suggests.

5.1.1. Evans’s Referential Treatment of E-Type Pronouns

Evans (1977, 1978, 1980) argued that E-type pronouns, pronouns not syntactically bound by their antecedents, are referential; “it” in (3) is a rigid singular term the reference of which is fixed by the description “the only donkey owned by Jack.” However, “it” in (3) should not be referential unless “a donkey” is. This can be illustrated by (6); the second sentence of (6) will on Evans’s view have (roughly) the logical form of (6*):

(6) If Jack owns a donkey, he beats it.

(6*). John beats *dthat*[[the x : donkey(x)][John owns x]]

A similar mechanism must be involved in (5). The indefinite does not refer, but the E-type pronoun refers to the object that verifies the clause in which the antecedent occurs. The proposal entails that (5) is false if there are farmers who own more than one donkey, regardless of whether they beat them or not (Neale 1990, chapter 6), and that (6) is false if Jack owns more than one donkey. It is also often pointed out (McKinsey 1986; Soames 1989) that if we evaluate (6) at a world in which John owns a *different* donkey that he beats, it should still be true; i.e. if the donkey he owns in the actual world is Jerry, and in world w Thomas (and he does not own Jerry in w), and he beats Thomas in w (but not Jerry), then (6) should be true in w . Hence, “it” in (5) and (6) is not referential.

5.1.2. D-Type Pronouns

Neale (1990; see also Davies, 1981) argues that discourse anaphora go proxy for definite descriptions formed from the noun phrase of the generalized quantifier expression in their antecedent. In (3) “it,” which Neale calls a “D-type pronoun,” goes proxy for “the donkey John owns” and the second sentence is equivalent to “Jack beats the donkey Jack owns.” Definite descriptions are uniqueness-implying generalized quantifiers:

$$(7) \text{ “the } x: Fx \wedge \Phi(x)\text{” is true iff } |F| = 1 \wedge |F| - |\wedge \Phi(x)| = 0$$

For (5), however, the uniqueness condition gives the wrong truth-value. Neale therefore introduces a “numberless” determiner “whe,” which incorporates a *maximality requirement* as a plural counterpart of the uniqueness requirement for singular definites:

(8) “ $\text{whe } x: Fx \wedge \Phi(x)$ ” is true iff $|F| \geq 1$ and $|F| - |\wedge \Phi(x)| = 0$

Anaphoric pronouns, then, sometimes go proxy for numberless description, in particular if the antecedent clause combines quantifiers with different quantificational forces.⁶³

If D-type pronouns are semantically equivalent to descriptions we would have counterexamples to our hypothesis that singular terms are *de jure* rigid in virtue of being structurally simple. Discourse anaphora appear structurally simple, and belong to the same grammatical category as *demonstrative* pronouns, yet on Neale’s account they are quantifier phrases. There are independent reasons to think that discourse anaphora are not abbreviated definite descriptions (at least if definite descriptions are uniqueness-implicating). For instance, D-type approaches appear to violate what Moltmann (2006) calls the *Same Value Condition* (SVC). To interpret an anaphoric pronoun it is not enough to use a description retrieved from the environment of the antecedent – the pronoun must have the same extension as the antecedent:

(9) Last year, John met a student who would become his first assistant. Bill met him then too [*ibid*, p.211]

(9) does not have the reading that John and Bill both met their respective future assistants last year. Instead the pronoun must have an interpretation according to which the

⁶³ A problem with this solution is that if pronouns anaphoric on singular existential quantifiers can be both numberless and definite descriptions we need some way of determining which reading is the correct one (see King, 2004, for discussion and objections); sentences generally appear not to have both.

extension corresponds to the *extension* of the antecedent. (9) involves a temporal context, but the same holds for intensional (non-attitudinal) contexts in general.⁶⁴

Above I showed that Evans wrongly predicts that “it” in (6) is a rigid, referential term. The argument was a version of Kripke’s semantic argument, and it establishes at most that discourse anaphora are not *referential*. Neale’s account predicts that anaphora are (generally) non-rigid, not only non-referential. If they were variables they would be rigid but not referential. So consider the following example, due to Salmon (2006, p.133):

(10) A comedian composed the score for *City Lights*. He was multi-talented.

(11) A comedian composed the score for *City Lights*. That he was multi-talented is a contingent truth.

(11) confirms that “he” in (10) is rigid and cannot be replaced by a description formed from the antecedent and still maintain the (modal) truth conditions of the whole discourse fragment. Similarly, consider the difference between the following pairs:

(12) A man composed the music for *City Lights*. If Copland had accepted the task he would have been someone else.

⁶⁴ A related problem is the so-called Bach-Peter sentences, of which (A) is, in fact, an example: (A) “*If a farmer owns a donkey, he beats it.*” In (A), interpreting the pronouns using definite descriptions (interpreted by (7)) recovered from the antecedents will either yield the wrong truth conditions or generate regress. If “he” is “the farmer who owns it,” we still need to interpret the pronoun. If “he” is “the farmer who owns the donkey the farmer owns” we have not yet obtained the right interpretation since “the farmer” is not properly restricted; but if “the farmer” is “the farmer who owns the donkey,” then “the donkey” is not properly restricted, and so on ad infinitum. If “he” is “the farmer who owns a donkey,” however, then the problem moves to “it.” If “it” is “the donkey he owns,” we have another “he” to take care of, and a new regress. If “it” is “the donkey a farmer owns” (A) gets the wrong truth conditions.

(12*) A man composed the music for City Lights. If Copland had accepted the task the composer of City Lights would have been someone else.

The second sentence of (12) is a contradiction. (12*) may have a reading on which the second sentence is a contradiction, but it surely has a reading on which it is true as well, which suggests that anaphoric pronouns are not abbreviations of definite descriptions.

D-type defenders may perhaps argue that the descriptions in question take wide scope over sentences in which the pronouns occur. Wide-scope readings may be achieved by quantifier raising at the level of logical form. However, raising the quantifier leaves a bound variable in place of the original syntactic occurrence of the quantifier. That is, “he would have been someone else” is left as an open sentence “*possibly*: someone(y)[$y \neq x$]” where [The x : composer(x)] is a prenex quantifier binding x . Now, variables are often seen as structural features that help us represent the logical form of a sentence, and as occurring in the representation of rather than in the sentence semantics itself. However, in (12) the second occurrence of “ x ” seems to correspond to an *overt* element in the syntax. Indeed, the second sentence seems best treated as the open sentence itself; but then “him” is “ x ” at the level of logical form, bound by a prenex quantifier. The quantifier must be recovered from the context to interpret the sentence, but the fact that “he” demands a wide scope quantifier suggests that “he” is equivalent to the *variable*, not the quantifier phrase, and that binding comes from somewhere else.⁶⁵ Worse, wide scope is not enough, since the quantifier representing “he” must be able to escape from *any* context, not just the context of the c-commanding unit in which it occurs. To get the correct reading of

⁶⁵ Though see King (1987, 1994, 2004) for a different take.

(12), for instance, the quantifier must scope *the entire conditional*, not just the consequent, which is syntactically problematic at best. Furthermore, if pronouns are covert descriptions they do not display any of the scope interactions with other operators we would expect; there is no narrow scope reading of

(13) Every farmer who owns a donkey owns exactly one saddle for it.

5.1.3 Situations

One way of avoiding the problems is to use situational semantics or event semantics (Heim, 1990; Ludlow, 1994; Elbourne, 2001, 2005). Situation-based accounts take *situations* to be introduced by implicit situation quantifiers associated with the context of the antecedent. Thus, (5) is interpreted as universally quantifying over situations:

(5*) Every minimal situation s in which a farmer owns a donkey is extendable to a situation s^* in which the farmer of s beats the donkey he owns in s .

Situations are world-parts, but achieving a precise understanding of what a *minimal* situation amounts to is difficult. Appeal to minimal situations are nevertheless necessary in order to guarantee uniqueness of denotation. Still, the approach seems well motivated for conditionals, and avoids the regress created by Bach-Peters sentences. However, nothing in (5) suggests a universal situation quantifier apart from the fact that we need it to obtain the right result. To turn this worry into an objection, consider

(14) Every student who had an umbrella brought it to school.

The most natural reading of (14) is a “weak reading” that does not require that every student who owns more than one umbrella brought every umbrella to school. But interpreting (14) in situation semantics is difficult:

(14*) For every student x , if there is a minimal situation s such that x has an umbrella in s , then can be extended to a situation s^* such that x brings the umbrella in s to school in s^* .

In (14*) the variable s in the consequent should be bound by an existential situational quantifier in the antecedent. The obstacle to achieving this is precisely the problem of donkey sentences above making a reappearance in quantification over situations.

In addition, situations do not give us sufficient information to ensure the uniqueness of the individuals. In (15) the pronouns cannot be replaced by (uniqueness-implicating) definite descriptions even when relativized to minimal situations:

(15) Everyone who buys a sage plant buys two others along with it. (Heim, 1983)

Nor do situation-based approaches straightforwardly give us the correct truth conditions for non-persistent quantifiers:

(16) If a farmer has exactly one donkey, he beats it.

(16) is true if any minimal situation in which a farmer owns exactly one donkey is extendable to a situation in which he beats it. But if a farmer owns n donkeys, there are n minimal situations in which he owns exactly one donkey. Thus (16) cannot easily be distinguished from the claim that every farmer beats every donkey he owns.⁶⁶

5.1.4 Definiteness and Semantic Simplicity

Since pronouns are definite expressions replacing them with definite descriptions is often possible. The pairs (10)–(11) and (12)–(12*) show that such substitution is problematic in intensional contexts. The situation is thus parallel to that of names. In non-intensional contexts replacing a name with an extensionally equivalent description is possible, but such substitutions fail in intensional contexts. For names I have argued that semantic simplicity predicted rigidity. Since pronouns are syntactically simple as well, we should *prima facie* expect them to work the same way. That a definite (as opposed to indefinite) description can often be invoked to fix the interpretation of the pronoun does hence not mean that the contribution of the pronoun *is* the contribution of the description; the correlation is already predicted by the fact that both are definite expressions associated with familiarity constraints. Thus, the fact that pronouns in non-modal contexts can often be replaced by definite descriptions carries little weight (since the substitution fails in modal contexts), and less than the fact that they are syntactically singular. To be sure, the definite description does help us interpret the variable (after all, a definite description will

⁶⁶ I have not provided an exhaustive discussion of a situational approach to donkey anaphora here, of course. Heim (1990) discusses solutions to (15). Kratzer (1989) suggests a solution to the problem of non-persistent quantifiers by (essentially) providing persistent interpretations for them. Ludlow (1994) suggests modifying the notion of a minimal situation. See Moltmann (2006) for discussion and further problems.

usually represent the variable's antecedent binder⁶⁷). However, such descriptions (as opposed to the binder) is best seen as meta-linguistic or context-modifying tools, the role of which may be to locate the antecedent and coordinate anaphoric chains.⁶⁸

If anaphoric pronouns are variables we retain an appealing symmetry between anaphoric and demonstrative pronouns. Both are singular terms and associated with a character–content distinction. The character of “he,” for instance, is “the most salient male.” Demonstrative occurrences are characterized by drawing on an extra-linguistic context. For anaphoric occurrences the *context* is linguistic, and the role of character is to locate an apposite binder. The character helps interpret the pronoun, but is no part of the semantic content – in “a man loves a woman; she loves him too” the characters of “he” and “she” help coordinate drefs properly, but leave no trace in the semantics.

Some counterexamples to the rigidity of pronouns have been proposed.

⁶⁷ If there were no quantificational difference between definites and indefinites, just a familiarity constraint (see also Ludlow & Segal (2003); Szabó (2006)), the analogy would be obvious. To represent the pronoun using the indefinite instead of the definite expression would engender no truth-conditional difference would ensue – it is the fact that the dref is a *familiar* that requests a definite. Consider: “A man started a brawl at the bar today. [...] He was tall,” where [...] is some intervening discourse. If a discourse participant got lost in the anaphoric network and asked “who is ‘he’?” the correct response is “the man who started the brawl.” If the same question is asked by someone who was not present when antecedent was uttered, and hence lacked access to the relevant dref, the best response is “a man who started a brawl at the bar today.”

⁶⁸ The matter is complicated by compelling evidence that discourse anaphora are not in general *syntactically* bound variables. In English the second occurrence of “his” in “Jack loves his wife” is ambiguous, but in many languages there is no ambiguity here – the bound case is marked by particular reflexive pronouns (“sin” versus “hans” for the possessive cases in Norwegian). Neale (2005), for instance, argues that this distinction counts against treating discourse pronouns as variables (since they are distinct – and marked by case, gender etc. in unbound cases). In the end, I suspect the difference may justify drawing a clearer distinction between syntactically bound elements and discourse anaphora (as e.g. in Dekker, 2012). However, I believe the difference primarily reflects a difference in activation conditions. Discourse anaphora are in any case, as we argue below, not syntactically bound variables, and understanding them on the model of individual variables is, as such, perhaps somewhat misleading. Nevertheless, I think there is a strong case for treating them as semantically as at least analogous to variables in first-order languages, under the important caveat that basic first-order languages do not really have the means to represent semantic binding at all. In 5.4 I offer a semantics for discourse anaphora, and I think the remaining question concerns the extent to which this interpretation also extends to anaphora that appear to be syntactically bound. Whether this promissory note is sufficient to allay Neale’s worries deserves further investigation. See also Dekker (2012) for a discussion of the relationship between pronouns and variables.

(17) Someone broke into the house. Mary believes that he stole the silver.

Moltmann (2006, p.212) claims that it is natural to interpret “he” in (17) *de dicto* (and many authors appear to agree). As long as we avoid confusing a rigid with a referential reading, however, I submit that the pronoun in (17) is rigid (compare “Mary believes that he would have failed if she had installed an alarm”). Similarly, (18) does not have a contradictory reading, whereas (19) arguably does (Evans, 1977):

(18) A man murdered Smith, but Joe doesn’t believe that he murdered Smith

(19) A man murdered Smith, but Joe doesn’t believe that the man who murdered Smith murdered Smith.

Neale (1990) argues that the *de dicto* reading should be theoretically available in (18) but is, for pragmatic reasons, so unlikely that we tend to disregard it. But the *de dicto* reading is not available in (18), and this is quite clear insofar as the *de dicto* reading *is* available in (19). There are cases that seem to require a *de dicto* readings as well, especially if the *antecedent* occurs within an attitudinal context (as in Geach sentences, discussed in Chapter 7). Even if there are such cases, examples involving attitudinal contexts are not obvious counterexamples to the rigidity of anaphoric pronouns. Even proper name rigidity is known to create problems in attitudinal contexts. If no counterexample to the rigidity of anaphoric pronouns is found that involves only temporal or intensional, not attitudinal contexts, I take this to be evidence *in favor* of anaphoric pronoun rigidity.

A possible non-attitudinal example might be

(20) The number of planets is eight. It could have been something else.

(20) might seem to have a reading according to which the second sentence is true iff a different number could have numbered the planets. However, “It” in this case does not seem to be a pronoun anaphoric on “the number of planets” at all. Rather, the pronoun is a pleonastic one, or at most anaphoric on a situation variable. Compare

(21) The president of the US is Obama. He could have been someone else.

(22) The president of the US is Obama. It could have been someone else.

(22) is true; (21) is barely felicitous unless the second sentence is a contradiction. “It” in (22) is not a pronoun anaphoric on “the president” at all. (20) has the same structure as (22), and “it” should accordingly not be interpreted as an anaphoric pronoun here either.⁶⁹

5.2. Dynamic Frameworks

D-type approaches treat anaphoric pronouns as semantically complete expressions. If anaphoric pronouns are treated as variables bound by linguistic context, then sentences in which they occur will not themselves be semantically complete but depend on linguistic context for their truth-conditional completions. Since quantifiers do not syntactically bind

⁶⁹ The situation is admittedly complicated by the inverse copular structure; such constructions are known to exhibit quite different properties compared to ordinary copular sentences (see Moro, 2007).

variables across sentence boundaries, such views require some form of score-keeping that assigns the anaphoric pronoun to its appropriate antecedent. On *familiarity views* of definiteness, indefinite expressions introduce new objects of focus to a discourse, whereas definites pick up objects that have already been made available (when no explicit antecedent is available, one needs presupposition projection; e.g. Heim, 1992).⁷⁰

The score-keeping feature of such theories suggests a revision of the role of *context* to reflect the fact that contributions to a discourse are interpreted in light of previous contributions (Lewis, 1979). Stalnaker (1978, 1998) defines the common ground of a discourse, what participants draw upon to understand new utterances, as a *context set*, a set of propositions that represent the possibilities that remain consistent with the information available at a stage in the discourse. Insofar as information added to a discourse concerns identifying the truth conditions of sentences it can be represented as a set of possible worlds (or functions from worlds), and extending information is a matter of eliminating worlds from the context set (though Stalnaker's diagonalization operation

⁷⁰ Now, definiteness could also be interpreted as a matter of the *unique availability* of a satisfier of a predicate (Roberts, 2003, 2004; Heusinger, 2004). Part of the role of definite expressions would still be to correctly index information supplied in discourse in a manner that ensures that any information given is associated with the right individuals. Notice also that analyzing definite descriptions in terms of the unique satisfaction of predicates goes beyond the widely accepted definition of definiteness in Barwise & Cooper (1981). According to Barwise & Cooper, a type $\langle 1, 1 \rangle$ quantifier Q is definite iff for every set A and every universe M , $(Q^A)_M$ is either empty or a non-trivial principal filter on M . If proper names are definite expressions (which they probably are) the definition must be modified, however. If names are treated in the manner of Montague, then whereas a lexical item “*Alexius*” denotes an individual, *Alexius*, the *noun phrase* “*Alexius*” denotes in every universe M a set of subsets of M (defined, for each individual α as $(I_\alpha)_M = \{A \subseteq M : \alpha \in A\}$). So *Alexius* (the noun phrase) denotes a quantifier $I_{Alexius}$ of type $\langle 1 \rangle$. We can then extend Barwise & Cooper's definition so that a definite noun phrase denotes a type $\langle 1 \rangle$ quantifier that is either empty or a non-trivial principal filter on every universe (see Peters & Westerståhl, 2006). Given this natural modification, however, there is a variety of ways to make definite descriptions satisfy the definition. Russellian approaches qualify; so do referential approaches and most if not all approaches that treat definite descriptions as terms, including interpreting them as ϵ -terms or choice functions. In the latter case the definite selects a principal filter that is picked up by a subsequent definite (see Haraldsen, forthcoming).

provides a more flexible set of resources for updating the context set). Similar ideas are defended in Gazdar (1979) and Kamp (1981).

Although Kaplan's treatment of demonstratives already relies on using *context* to resolve apparently non-compositional semantic phenomena (identical linguistic constructions that have different truth-conditions on different occasions), to account for anaphoric relations context must do more than Kaplan's logical or physical contexts. If updating a context set is just a matter of determining truth conditions, the additional framework we need could be defined on top of traditional truth-conditional semantics. But that may in fact not suffice to keep track of anaphoric relations. Such relations are updated without necessarily affecting information relevant to truth conditions; the linguistic context is itself context-sensitive (Heim, 1982, Groenendijk & Stokhof, 1991, Kamp & Reyle, 1994, Chierchia, 1995), as illustrated⁷¹ by Partee's famous example (Heim, 1982): (23) and (24) have the same truth-conditions, but (24) is infelicitous:

(23) Out of ten marbles, one is not in the bag. It must be under the couch.

(24) Out of ten marbles, nine are in the bag. It must be under the couch.

(23) and (24) thus differ in how they change contexts over and above how information affects truth-conditions. Kaplan-contexts only involve parameters that affect truth-conditional content, and what is at stake here is not information about the world but information about the discourse itself. To capture this, the states representing the available information at a stage must take into account both kinds of information.

⁷¹ See Muskens, van Benthem & Visser (1997) for a survey of arguments to the same effect using e.g. tense and aspect, presupposition, modality, conditionals, counterfactuals, and questions and answers.

On *dynamic views* unbound anaphoric pronouns can (syntactically) be assimilated to variables in formal logic. Mechanisms proposed to accommodate dynamic contexts that provide interpretations or representations for occurrences of such variables, include discourse structures understood as psychological layers of representation (Kamp & Reyle, 1993); a structured set of questions under discussion (Roberts, 1996b); a dynamically updateable global choice function defined for all predicates in a context (Peregrin & Heusinger, 2004; Heusinger, 2004); or sets of information states that encode a set of assignment functions (Groenendijk & Stokhof, 1990, 2000; Muskens, 1996). We divide dynamic theories into representational and non-representational accounts. 5.2.1 and 5.2.2 outline the respective approaches, and 5.2.3. adds some critical remarks.

It is worth pointing out that “dynamic semantics” is often used to denote only the non-representational views of e.g. Groenendijk & Stokhof (1990). These views present a rather radical departure from traditional approaches to linguistic interpretation and take the interpretation of an utterance to be not a proposition, but a function from a context C to an updated context C^* , where C^* is in turn an argument to the next utterance. Sentence meanings are hence identified with *potential operations on contexts*, where a context is understood as a current state of information, and the sentence is not only evaluated truth-conditionally but according to how it alters context. By contrast, discourse representation structures in representational views such as DRT are linguistic, not semantic, objects – they are *representations of* information, not information, and the interpretation of these structures takes the form of standard, static truth conditions. So on these views, *meaning* (if anything can be defined as the meaning of a sentence) is not dynamic.

5.2.1 Discourse Referents and DRT

To get a better grasp of representational views we (re-)introduce *discourse referents* (“drefs”).⁷² Whereas indefinites introduce new objects into a discourse, the role of definite expressions is to direct focus to items already introduced. Karttunen (1976) pointed out that the mechanisms involved in this familiarity requirement are operative even when there is no particular reference:

(25) If you find an error, write it down in your book.

(25) has a universal interpretation, so “it” is not a *particular* error. Instead the indefinite introduces a dref, an element that stands for the “thing,” or repository of things, introduced by the indefinite and picked up by the pronoun.

What drefs *are* depends on the system in which they are incorporated. In dynamic systems the official content of a pronoun is a variable the value of which is (given by) a contextually specified function assigning values to variables. Interpretations of sentences in which the pronoun occurs are constrained by information associated with that variable by the interlocutors. Drefs should thus in some manner encode this information, and a

⁷² Another motivation is related to indefinite descriptions. It is often observed that indefinite descriptions are accompanied by “referential intentions” (e.g. “I had a beer after work” sounds misleading if I had five). It has been suggested that indefinite descriptions are ambiguous between attributive (quantificational) and referential uses (Chastain, 1975; Donnellan, 1978; Fodor & Sag, 1982). But there are syntactic constraints, in particular so-called island constraints, that create problems for the quantificational option, and potential “intermediate scope” readings that create problems for uniformly referential readings (Ludlow & Neale, 1991). A promising move is to interpret indefinites not as quantifiers (nor referential terms) but as open expressions introducing drefs (“an *F*” is “*Fx*”), as in DRT, or to interpret drefs corresponding to indefinites by choice functions or ϵ -terms (Winter, 1997; Heusinger, 2004). Such moves suggest that in a fully adequate theory names and pronouns are not *really* constants or variables, though since sentences involving them will still be (first-order) semantically equivalent to (prenex) existentially quantified sentences, these options would represent refinements of rather than obvious objections to the views defended here.

pronoun links already available information to a new discourse contribution. Officially, then, a dref is a node in the information structure of a discourse, which helps organize the stored information and which belongs to the level of representation rather than the level of objects represented. Hence, a dref is not an individual of a domain, but something that keeps track of (e.g.) potential variable assignments. Its role is to coordinate assignments to ensure that any element assigned to “an F is G ” is the same as the element assigned to the subsequent “it is H ,” and that the fragment is not true if every F that is G is different from every F that is H . A dref may thus be viewed as encoding an assignment function or something (a choice function or Skolem function) that does the work of an assignment function. Across a discourse the drefs – the set of permissible assignment functions – are managed dynamically to reflect the way information changes. Adding information about a dref means constraining the permissible assignments.

In Discourse Representation Theory (DRT) (Kamp, 1981; Kamp & Reyle, 1993) the antecedent is the dref itself (introduced by an indefinite). The rest of the information associated with it is stored in the *discourse structure*, which provides the apposite binder. DRT thus employs a level of representations, discourse representation structures (DRS), which are consecutively constructed and developed in an unfolding discourse. A DRT consists of a formal definition of the representation language, which contains a recursive definition of the set of well-formed DRSs and a model-theoretic semantics for its members, as well as a construction procedure specifying how a given DRS is supposed to be augmented or extended for a given input sentence.

A DRS consists of a universe of drefs and a set of conditions that encode the information already given regarding the drefs. Suppose we have a DRS like:

(26) $[x, y: \text{farmer}(x), \text{donkey}(y), \text{owns}(x, y)]$

(26) represents the information that there are two objects, a farmer and a donkey, and that the farmer owns the donkey. In the model-theory we need an embedding function from drefs to individuals. The embedding function f verifies (26) in a model \mathcal{M} just in case f 's domain includes (at least) x and y , and in \mathcal{M} $f(x)$ is a farmer, $f(y)$ a donkey, and $f(x)$ owns $f(y)$. (26) is read as “a farmer owned a donkey.” The two indefinites introduce new drefs (as well as the information from the predicates). Suppose the next DRS in the discourse is (27) (the subscripts indicate that the variables need antecedents):

(27) $[u_a, v_a: \text{beats}(u, v)]$

Since (27) is given on the background of (26) the variables are (pragmatically) resolved through a merging of (26) and (27) to obtain:

(28) $[x, y: \text{farmer}(x), \text{donkey}(y), \text{owned}(x, y), \text{beats}(x, y)]$

(6) needs a more complex treatment. “beats it” is not c-commanded by the antecedent, and requires a new DRS along the line of (28). To account for the anaphoric relation between DRSs – the antecedent and the DRS for “beats it” – we also need the notion of *accessibility* between DRSs. The result (after variable relations are resolved) is:

$$(29) [{}_1x: \text{Jack}(x), [{}_2y: \text{donkey}(y), \text{owns}(x, y)] \rightarrow [{}_3 : \text{beats}(u_a, v_a)]]$$

The subscripts indicate the relative structure between the DRSs – the manner in which the DRSs are governed by each other – and define the accessibility relations between them: If f , the embedding function, is to verify DRS_1 (scoping the entire (29)), then $f(x)$ must be the individual corresponding to the dref *Jack*, and every extension of f that verifies DRS_2 must be extendable to a function verifying DRS_3 . Hence, DRS_1 must be accessible to DRS_2 , which is accessible to DRS_3 (accessibility is transitive in DRT). u is then linked to x and v to y , and the result is a DRS for “Jack beats every donkey he owns.”

We can define quantifiers and other operators as well, but for details I refer the reader to Kamp & Reyle (1993). Importantly, indefinites are not quantifiers in DRT. Their role is to introduce new drefs (and conditions), and they have no quantificational force on their own. Instead, “a man” contributes what is, in effect, an open formula “ $\text{man}(x)$ ” and introduces a dref that can be picked up by discourse operators of different kinds. The apparent quantificational force of “a man” derives from the environment in which it occurs. In (26) the force is existential since this DRS is verified in \mathcal{M} iff there is a way of verifying it in \mathcal{M} . In (28) it is universal since a condition $\Phi \rightarrow \Phi^*$ is verified in \mathcal{M} iff every way of verifying Φ can be extended into a way of verifying Φ^* .

Two features of the DRT account are worth emphasizing. First, indefinites do not bind anaphora. Instead, anaphora are related to antecedent conditions by the fact that the drefs in the DRS for the antecedent condition are accessible (in the domain of the DRS) for the anaphoric expression. There is no syntactic binding. Binding, requires pragmatic

resolution, and DRT is hence not compositional. Context and pragmatic factors determine the contribution of indefinites and anaphora. Second, DRT is a representational theory. DRSs are not discourse information but *representations* of information – sentences in a discourse get their semantics indirectly via these representations, and the interpretations take the form of (static) definitions of truth. Thus, there is a sense in which DRT is not *semantically* dynamic. The semantic contents for (23) and (24) – their models – are the same, and it is the representations that are dynamically different. Indeed, *representations* play an ineliminable role in DRT, and the status of the representations – which are vehicles for content, not contents – is a matter of debate. (Kamp & Reyle (1994) took them to be mental entities, in which case DRT would be essentially psychologistic.)

5.2.2 DPL: A Brief Outline

In DRT the meaning of a sentence (if any can genuinely be ascribed when representations do all the work) consists in the manner in which it modifies a DRS, and DRSs are bearers of truth conditions. A DRS Φ_i is transformed into Φ_{i+1} when a new sentence is added to the discourse. Semantically, there will be an associated transition at the model theoretic level from $I(\Phi_i)$ to $I(\Phi_{i+1})$. This observation suggests that it is possible to dispense with the level of representations and obtain a dynamic theory in which sentence meanings are sentences' contributions to transitions from one *semantic* entity to another.

In DPL (Groenendijk & Stokhof, 1991) values are assigned to variables directly, not via an intermediate level. An *information state* contains information both about the world and about discourse. The former is modeled as a set of worlds, whereas the latter (drefs) is modeled as sets of assignment functions that can be manipulated by operators in

the semantics. Drefs are also elements of information, and hence part of semantics.

Updating context is a matter of deleting worlds incompatible with the information state at a given stage. Information does not only delete worlds directly, however, but also in the course of keeping track of assignments to variables at the different worlds – these assignments are modified along the way, and a world w is deleted if no assignment satisfies a given contribution at w . The idea, then, is that discourse information is linked to information about the world through possible assignments of objects to variables, namely the objects that relative to a particular world (and the values of other drefs) could be the values of the variable in question. Since more than one extension of a discourse is possible, links can be maintained in several worlds simultaneously. Drefs are partial, indefinite, non-identified information objects, which are picked up by an anaphoric pronoun by aligning the assignment to the pronoun with the assignment determined by its antecedent. *Indefinites* are interpreted in terms of dynamic existential quantification, and the dynamic effect of adding an indefinite is to make a new dref available. Note, however, that in DPL there is no level at which the anaphoric relation is *represented* – all dynamic processes occur at the level of semantics, and updates of context information must be represented semantically. Thus, in DPL a sentence's meaning is identified with its *context change potential*; meanings are *updates* of information states.

Assume that a context is a set of pairs, each of which consists of a possible world and an assignment function. The context comprises the world-assignment pairs where the assignment assigns objects to the variables that satisfy relevant conditions (predicates) in that world. Formally, the interpretation of a formula ϕ relative to a model \mathcal{M} is a *set of pairs of variable assignments* $\langle g, h \rangle$ such that $\langle g, h \rangle$ is in the interpretation of ϕ relative to

\mathcal{M} iff ϕ can be successfully interpreted relative to an *input* assignment g and yield h as a possible *output*. (That is, ϕ is true relative to g iff $\exists h g[\phi]h$.) In the classical version of DPL in Groenendijk & Stokhof (1991) atomic sentences have the same input and output assignments, however, and affect the context by mapping a context to a (possibly) smaller context by retaining the world-assignment pairs that make the sentence true (whereas a negated sentence $\neg p$ eliminates the world-assignment pairs that make p true).

The main difference compared to a standard first-order semantics concerns existentially quantified formulas and conjunctions. If g is an input assignment, then the interpretation of $\exists x\phi$ requires that any assignment k differing from g only in the value for x be tested to see if it serves as an input for the interpretation of ϕ .⁷³ If it does, and h is the output, then h is also a possible output of $\exists x\phi$ on g (if x occurs free in ϕ and ϕ imposes certain constraints on the valuation of x , then the output valuation must also satisfy these conditions). In terms of context updates, then, an existentially quantified sentence containing x applies to a context C by changing the assignment g of any world-assignment pair in C to the alternatives h that make the scope of the quantifier true. These assignments are then possible outputs for the sentence, and remain in C . A *conjunction* is interpreted by applying the conjuncts in the right order. $\phi \wedge \psi$ does not affect context on its own, but composes possible changes introduced by the conjuncts in combination. So if ϕ takes an input g and produces j (where $g \neq j$), and ψ takes j as input and delivers h as

⁷³ Notice the connection to Tarski semantics. In Tarskian semantics “ $\exists x\phi$ ” is true in \mathcal{M} , under an assignment A iff there is some assignment A^* that differs from A at most in the value that gets assigned to x , and ϕ is true in \mathcal{M} under A^* . Implicit in Tarski semantics, then, is a relation $[x]$ that holds between A and A^* iff for all variables y it is the case that $y \neq x$ implies $A(y) = A^*(y)$. So the idea of pairs of assignments associated with sentences is already implicit in Tarski semantics.

possible output, then $\phi \wedge \psi$ accepts g as input with h as possible output. The semantics of $\phi \wedge \psi$ thus involves interpreting ϕ *prior* to ψ , so “ \wedge ” is not commutative in DPL.

It is, however, important to emphasize the crucial change of perspective that takes place when moving to DPL. In DPL meanings are themselves dynamic, or *types of actions* individuated by the changes they bring about. The information state represented by a sentence Fx is a *precondition*, which can be defined (disregarding the world parameter) as a set of all assignments g with the property that $g(x)$ is an object satisfying F . *Meanings*, however, are relations between input states and output states. For “an F ” the (dynamic) meaning of “ $\exists x$ ” is an instruction to replace the old value of x by some arbitrary new value – i.e. the action type it denotes is the action of *resetting* the variable. The predicate “ F ” is a test on the “resetting” of the variable, and the precondition of the sentence is the set of assignments that can be successfully reset with respect to x to an assignment in F (the interpretation of “ F ” is a test that passes on satisfying assignments and throws away others, but modifies nothing); the assignments that pass the test then constitute the updated information state.⁷⁴ “He is tall” is represented as the open sentence

⁷⁴ An advantage of DPL is that it is fully compositional, and can be represented in a (dynamic) Montague grammar (DMG) (see in particular Muskens, 1996). Instead of basic types e and t of classical extensional Montague grammar, DMG has basic types e , t and m , where m is the class of *markers*. Information states pick out entities for markers, and can be viewed as objects of type $\langle e, m \rangle$. Abbreviating $\langle e, m \rangle$ as s (for “state”), we can call objects of type $\langle \langle t, s \rangle, s \rangle$ *state transitions*. Now, in classical, extensional Montague grammar “a man” would be $\lambda P \exists x(\text{man}(x) \wedge Px)$, and verb phrase meanings (for “ P ”) are of type $\langle t, e \rangle$. In Dynamic Montague Grammar (DMG) the translation of an indefinite noun phrase introduces an anaphoric index. “A man” is $\lambda P \lambda a \lambda a^* \exists x(\text{man}(x) \wedge Pu_i(u_i|x)aa^*)$. The variable P in the DMG translation of “a man” has been lifted, relative to classical Montague grammar, to type $\langle \langle \langle t, s \rangle, s \rangle, m \rangle$. Indeed, DMG can be viewed as the result of systematic replacement of *entities* by *markers* and of *truth-values* by *state transitions*. The (verb phrase) meaning for “is F ,” for instance, is a function that maps a marker to a state transition. The state transition for marker u_i will check whether the input state maps u_i to an F entity, and whether the output context equals the input context. The variables a, a^* range over states, and the expression $(u_i|x)a$ denotes the result of resetting the value of u_i in a to x (that is, the old value of u_i is replaced). The anaphoric index i on reference marker u_i is introduced by the translation. (A remaining worry, however, is to specify how the correct markers are assigned; see Muskens, 2011).

“Tall(x).” Even if “ x ” is syntactically free it gets connected to a preceding quantifier through the assignment updates, and is in effect bound by it. We say that “ x ” in that case is *semantically* (but not syntactically) bound.

I refer the reader to Groenendijk & Stokhof (1990) for further details. The system can be extended with generalized quantifiers, plurals, modal and epistemic operators, and other features, and several modifications have been proposed. The outline given here will hopefully suffice at least to suggest the rather novel treatment of anaphora in DPL compared to classical, static systems.⁷⁵

5.2.3 Whither Dynamic Semantics?

An information state contains a set of worlds, so one might say that the state *entails* a proposition (set of worlds) if the state is a subset of that proposition. Nevertheless, sentence meanings are not propositions in DPL. Indeed, it is unclear whether DPL gives us anything that can play the traditional role of *meaning* – that which a speaker expresses and an audience grasps in successful communication. Although sentences can be associated with truth conditions, they receive truth conditions only derivatively – their meanings are their context change potential. In particular, it seems difficult to relate meanings to belief states. At least if belief contents are things that have a semantic

⁷⁵ Donkey anaphora can be interpreted given the observations that for any formula ϕ, ψ , $(\exists x\phi \wedge \psi) \Leftrightarrow \exists x(\phi \wedge \psi)$, and $(\exists x\phi \rightarrow \psi) \Leftrightarrow \forall x(\phi \rightarrow \psi)$, which are classical equivalences, but without the restriction that x not occur free in ψ . Instead, a variable in ψ is semantically or dynamically bound by a previous existential quantifier. (In conditionals the semantic binding gets universal force.) On the other hand, the fact that \neg is an operator that introduces tests without any further dynamic impact means that (for instance) $\neg\neg\phi$ is *not* equivalent to ϕ in DPL. This explains why (4) is (generally) not acceptable. The facts about semantic binding given in the semantic clauses correspond one-to-one to the facts about accessibility (and inaccessibility) in DRT, and DPL is hence a (fully compositional) generalization of DRT without the representational level (Groenendijk & Stokhof, 1991; Muskens, 1996).

structure, they should be information states or things that realize information states. But belief contents are not DPL meanings, *functions* on information states, on pain of regress.

It may be possible to capture the differences between pronouns treated as constants and pronouns treated as variables in some version of DPL (though it is unclear for instance how we achieve anything resembling a distinction between character and content). However, DPL treat anaphoric pronouns as discourse markers, the role of which is to map information to information states, and variables themselves have no semantic values, at least not in a way that would resemble the values of constants – at most, variables are repositories of information. It would for our purposes be at least preferable (though see below) if we could identify specific (if context-dependent) values for pronouns (as variables) that are analogous to the values constant expressions such as indexicals or names. I will argue that it is possible to achieve this in an essentially static system. It may be that the solution can be translated into some version of DPL, though even if we could, hesitations regarding the significant departure from classical semantics might at least confer some value on finding a solution in classical, static semantics.⁷⁶

In DRT the level of representation (DRSs) is associated with a classical model theory. But DRT leaves no obvious room for propositions in any classical sense either – the representational levels are, at least on many versions of DRT, better compared to Davidsonian meta-linguistic axioms in individual speakers' idiolects than to Fregean

⁷⁶ It has been argued that DPL's departure from classical semantics is less than one might immediately think. Van Eijck & Visser (2010), for instance, argue that dynamic semantics provides a generalization of truth conditional semantics rather than a radically different alternative. Sentences' "compositional [dynamic] meanings" are functional relations between information states; classical meanings, on the other hand, are preconditions for success of the discourse actions, or "projections of the compositional meanings." However, it is – as mentioned – hard to see exactly how to square any traditionally oriented view of *propositions* with DPL, at least insofar as propositions are assumed to play the roles they are traditionally supposed to play as the meanings of sentences and contents of belief. For problems related to the treatment of anaphoric pronouns in dynamic systems, see Elbourne (2005) and Moltmann (2006).

senses – but one can at least derive classical truth conditions for particular contributions. DRT is, however, inherently representational. A sentence utterance prompts the audience to modify their representations (DRSs), and it is representations, not sentences, that have a truth-conditional interpretation. For Kamp & Reyle (1993) the representations are psychological entities, a position that raises familiar concerns regarding subjectivism about meaning stemming back to Frege (1884) and Wittgenstein (1954) (for a recent discussion, see Steedman & Stone, 2006). On this interpretation, DRT reduces semantic operations to psychology. The construction algorithms for DRSs model how we *in fact* interpret discourse; they do not issue *normative* constraints on DRS construction, and are at most generalizations over actual practices. There may be naturalistically explainable benefits of following these construction algorithms rather than others, but that a given construction procedure is, say, truth-preserving is at most evidence that we have followed the e.g. evolutionarily most advantageous way of processing sentences rather than the *correct* one. If one, on the other hand, views the correctness or not of an interpretation as something that is constrained by objective *rules* – that a sentence is correctly interpreted if we have given it the *correct semantics*, where correctness is a matter of making the correct moves with regard to the rules of interpretation – then psychologism is out.

It is not inherent in DRT that representations are subjective. They could, for instance, be possibly *abstract* contexts representing the common ground of a discourse. This would not resolve all worries, however. A thorny question for DRT concerns directly referential terms – demonstratives or names – for which truth-conditional contributions depend on the referent, not descriptive information associated with the expression. Kamp suggests employing “external anchors,” where anchors are (formally)

pairs consisting of a *dref* and a referent, to interpret referential expressions. Anchors are no part of the DRS-constructions, since DRSs are linguistic objects, but rather constrain which embeddings of individual representations are allowed in the model to ensure that the truth-conditions are properly sensitive to the referent. A name, then, contributes a *dref* – a representational symbol – to the DRS, and an *anchored dref refers* to an individual.

But first, an occurrence of the indexical “here” must be anchored as well. If an occurrence of “here” gives rise to a *dref*, however, and it is the *dref* that is anchored, then the *dref* cannot correspond to the character of “here” as we would expect but, it seems, to the content. Yet the *dref* is supposed to be the representation-level entity, and should surely correspond to the character, if anything (see also Maier, 2009). And if names are closely related to demonstratives, the same worry applies to names. Second, remember that the representational level of DRT is not just a helpful model-theoretic level but an intermediate level between sentences and models that is essential to the interpretation – resolving anaphoric links, for instance, relies on operations on *symbols* in the DRS, the *drefs*, and the incremental construction of these linguistic objects, not on the model theory. External anchors, however, are no part of the representations, but supposed to constrain mappings of DRSs onto model theoretic entities. As pairs of *drefs* and referents, anchors mix model-theoretic entities with representational ones, and it is at least unclear what to make of such hybrid entities.

One way of possibly avoiding the last worry is to model names as (semantically) bound discourse referents, the occurrences of which trigger binding presuppositions, as suggested by Geurts (1999). Maier (2009) develops a multi-layered version of DRT in which presuppositions triggered by name occurrences are rigid. Thus, Geurts and Maier

avoid hybrid entities, and names are interpreted as bound drefs. However, to account for expressions that are not only rigid but *directly referential* something more is needed.

Directly referential expressions can be characterized by the fact that they satisfy Kripke's semantic argument, and are only contingently related to the presuppositions governing them. Perhaps, as Maier suggests, such expressions may be associated with extreme pragmatic pressure to maintain their status as rigid even in the face of information that contradict the relevant presuppositions, but, in addition to a certain worry of *ad hocness*, it is not clear that this is enough to completely characterize the behavior of "ardently" referential expressions (that are never "just" rigid), such as indexicals.

A different worry is that whereas drefs may be rich representational entities, and properly constrain assignments of values, they are not themselves associated with particular semantic values. Information stored in the discourse structure connects a free variable to an appropriate binder, but in a *model* it presumably receives a value only as a step in defining truth for a DRS. Insofar as pronouns are singular terms, however, one may think that they should have semantic values characteristic of expressions belonging to that category. But a value relative to an assignment is not a proper semantic value for an overtly occurring variable (that value is *not* the semantic value of the expression; otherwise there is no semantic difference between variables and constants). On DRT the values of discourse anaphora correspond to updates of or quantifications over assignments, something that at least *prima facie* appears to involve an appeal to entities (assignment functions) that belong to the abstract tools used to build models, and not something that serves as semantic values for natural language expressions.

There are ways of interpreting discourse anaphora using elements (choice functions or Skolem functions, for instance) that can perform the roles of assignment functions without invoking the model theory itself, and we discuss these below. Such interpretations can at least in principle be captured in static frameworks (at least as long as we recognize a dynamic discourse context or salience structure that keeps track of anaphoric chains). To sum up, although DPL may, in principle, be a suitable framework for the pluralism defended in this dissertation its departure from classical semantics at least suggests looking for a way of interpreting discourse anaphora – as singular terms rather than descriptions – in a static framework. In the next section I first discuss the extent to which we should take the possibility of providing semantic values for discourse anaphora interpreted as semantically bound variables as a success criterion for a theory of anaphora, and discuss alternative static approaches. In 5.4 I outline my own model.

5.3 Notes on the Semantics of Variables, and Static Interpretations

For constants, we can define *reference* by axioms (relative to a sequence s of objects)

$$\text{REF}(\text{“Jack”}, s) = \text{Jack}$$

Individual variables are terms whose values are given by the following axiom schema:

$$\text{REF}(\text{“}x_k\text{”}, s) = s_k$$

Thus, the referent of “ x_k ” relative to s is the k th element of s . As opposed to constants, values of variables are sensitive to choice of sequence, and vary with assignment. The schema suffices for defining truth-conditions for quantified formulas

“ $\exists x_k \Phi$ ” is satisfied by s iff Φ is satisfied by at least one sequence that differs from s at most in the k th position.

But what is the *non-relative* semantic value of “ x_k ?” One cannot say that since the value of a term is the individual it denotes, the value of an occurrence of a variable is the individual in the sequence assigned to x by an assignment. In that case, the contribution of a variable is indistinguishable from the contribution of constant, for instance at the level of the proposition. Not only is grasping a natural language sentence containing an overt variable expression sensitive to the fact that the variable may have several satisfiers, and different ones in different worlds; this sensitivity should be reflected in the semantic content of the variable expression, not only in the model theory.

5.3.1. Do We Need Values for Variables?

What variables *are* is a difficult question (see Fine, 1985, 2003, 2007). Our task does not concern the nature of variables, however, but finding something that can serve as the semantic value of a bound variable. Variables are not referential; they designate no individual but *range* over domains. They are rigid, in the sense that occurrences of x bound by the same quantifier must be satisfied by the same value regardless of the

occurrence of modal operators, but not strictly speaking *designators* (except relative to an assignment). Can we sidestep the question of finding apposite semantic values, then?

First, the problem is not a problem for the view of descriptive name I defend in particular, but for any view that interprets any pronoun occurrences as variables. In other words, an adequate account of variables is needed in any case. Any adequate account must distinguish constants and variables, and can hence be used to interpret occurrences of discourse anaphora and names in a pluralist framework. *Second*, the worry is not necessarily a worry for classical semantics but a potential problem for a referential semantics for a language containing syntactically overt bound variables distinct from referential terms. Notice that Tarskian semantics *does* give us the correct truth-conditions for sentences containing variables, and perhaps no more is needed. One might argue that the idea that a variable needs an independent semantic value is a consequence of adherence to strict classical compositionality, where every syntactic element needs a well-defined value. This commitment is not mandatory (it is rejected by DRT). The notion of compositionality in question would be a very strong one, according to which the semantic value of a sentence is not only a function of the values of its constituent expressions, but a compound of those constituents. *Third*, even a commitment to strict compositionality may not require robust, separate semantic values for bound variables. Variables are, in a sense, “traces” left by quantifier raising, and one could add that we sometimes (especially in stretches of discourse) make them explicit for the purpose of clarifying quantifier scope – it is thus not obvious that we ought to worry too much about giving them separate semantic values since their contributions can be understood in terms of the scope of their associated antecedents. Overt variable occurrences may then be

understood as markers that signal that the sentence concerns an individual that fill the roles denoted by the predicates in a manner defined by the discourse.

One may wonder whether conceding that variables do not have semantic values on their own jeopardize the analogy between variables and constants I have defended, but it is not obvious that it will. A deictic name or demonstrative occurrence needs an overt value to be interpretable because there is nothing else for it to contribute. The analogy may still hold if we view variable occurrences as “fundamental” and constant occurrences as derivative. Remember that the requirement that constants are directly referential is derived from the extensionality of variables (a consequence of the fact that quantifiers are objectual). Variables are required to contribute no descriptive material and range over individuals; constants are special cases since they draw on the model or extra-linguistic context for a sentence to be interpretable. Since there is no more semantic structure to a name than to a variable the only option is then to accommodate it by giving it a referent.

On the other hand, discourse anaphora are exactly *not* in the syntactic scope of the antecedent, and as shown by discourse fragments such as “few philosophers came to the party. They had a good time, though,” the binding inside *c*-scope is significantly different phenomenon from any (semantic) binding outside. This at least suggests that appeals to quantifier traces and scope resolutions are not quite enough; discourse anaphora alter discourse information in a manner that scope resolutions alone cannot explain.

5.3.2 The Semantics of Variables: A Suggestion

An open expression Φx may be treated as a function from any potential value v of the bound variable to the value of Φx under the assignment of v . The open sentence could

then be associated with an open proposition or propositional function, and the semantic value be a function from assignments to closed propositions. Following Salmon (2006) the value of Φ could thus be represented as:

(V): The value of Φ with respect to variable x under assignment s is: (λv) [the value of Φ under s^x_v], that is, the function that maps any element $v \in D_x$ (where D_x is the domain of x) to the value of Φ under the assignment s^* that differs from s at most in that assigns v to Φ .

The open proposition denoted by “ $F(\textit{Julius})$ ” (given that “*Julius*” is a variable) is true in a context supplying at least a world w and an assignment that maps *julius* to Judson just in case Judson is F in w . Insofar as “*Julius*” occurs bound – that there is an antecedent of the kind $\exists x_{\textit{julius}}$ [invented-the-zip($x_{\textit{julius}}$)] – the contribution of the open sentence is a function from objects to truth-values, which again is an input to the binding quantifier phrase, a function from functions from objects to truth-values, to truth-values. According to Salmon’s suggestion sentences are associated with propositions,⁷⁷ a sentence involving a (free) variable is, in effect, interpreted as a closed sentence where the variable expression is bound. What the open sentence “ x is F ” contributes in a context, however, is not just the proposition closed under the assignment of a value to “ x ” – there is no *single* such

⁷⁷ Insofar as the value of the variable is an identity function, there are as many propositions in the domain of Φx (in bound position) as there are individuals of the domain. An alternative is to say that the character of “ x is F ” is, roughly, a function from a context C to a closed proposition in which the occurrence of “ x ” is bound (where C also determines the binder). Its content is a function (determined by the binder in C) from a function from assignments to truth-values, to truth-values. Again, “ x is F ” is itself a “variable expression,” but now we can let it range over a universe of truth-values rather than propositions.

proposition. Under the scope of “ $\exists x$,” however, “ x is F ” ranges over singular propositions (or is a function from assignments to singular propositions), and the general proposition is true if there is an assignment that yields a true singular proposition.

We can also say something about the contribution of the variable itself. Since variables do not *designate* universes of values – the value of “him” is certainly not the class of males in our domain – but range over them, we may say that their values are *identity functions* on said universes. That is, variables denote functions that supply potential values to the function that is the value of Φx , and contribute these functions to the value of the binding quantifier. The binder Q would in that case determine which input to Φx satisfies $Q\Phi x$ in accordance with standard Tarski semantics. If the quantifier phrase is an existential restricted quantifier [$\exists x: \Phi x$](Ψx), there must be an assignment (or variant assignment), an output of the identity function over the relevant domain, that satisfies Ψx and also satisfies Φx .

And now we can capture the difference between variables and constants. Whereas the intension of a bound variable is the function from any potential value of the variable to the extension of the variable under the assignment of that value (that is, an identity function on a domain), constants are constant functions from worlds to extensions. The value of the constant is set by the model. The *value* of the variable expression is not assignment sensitive either, insofar as the identity function is defined for all possible worlds, but which input/output to the function satisfies $Q\Phi x$ may change when $Q\Phi x$ is evaluated at a different world with a different associated domain (e.g. if the world of

evaluation for the *whole* of $Q\Phi x$ is shifted – if an intensional operator occurs *within* Φx , then all occurrences of “ x ” must be satisfied by the same individual for Φx to be satisfied.

We can let any sentence with a singular term denote an open proposition; if there is an appropriate binder available, the open proposition is an input to this binder. If there is not, or if other factors tell us to look to the extra-linguistic context for an interpretation, we can accommodate an extra-linguistic referent, and the term is a constant. We thus retain an appealing symmetry between variables and constants. The difference is that since the value for a constant c is provided by the model, identifying the intended model for a sentence $\Phi(c, x)$ requires identifying (assigning) its referent but does not require identifying the value of x that makes Φ true. One upshot is a difference with regard to tolerating emptiness. Insofar as an interpretation of $\Phi(c, x)$ depends on the value of the constant c , Φ cannot be interpreted if c is empty. On the other hand, the interpretation (as opposed to truth-value) does not depend on what satisfies x . The identity function associated with “ x ” is defined insofar as its range is defined – the domain is non-empty – even if nothing satisfies Φ .

5.3.3. Static Approaches to Cross-Sentential Variable Binding

According to Salmon “he” in (10) is regarded as an occurrence of a variable bound by a wide-scope quantifier that is merely implicit in the second sentence. This yields the right modal behavior. On the other hand, the semantic content of the second sentence in (10) cannot be read off its surface form. Its underlying logical form is that of a closed sentence, and to evaluate the occurrence of the apparently open sentence we must recover the off-stage antecedent quantifier. Unfortunately, we do not have any obvious systematic

way of recovering off-stage quantifier phrases, and hence no systematic way to account for instance for donkey anaphora.

Berger (2002) interprets sentences containing discourse anaphora as being dependent on their antecedent background (A-B) sentences. Syntactically, discourse anaphora are interpreted as *parametric covarying occurrences of variables*, variables the assigned value of which must be the same as the value assigned to the variables in the A-B sentence. The semantic value of a variable occurrence is thus, presumably, a parameter. It is unclear how this approach would deal with A-B sentences containing non-logical quantifiers. Furthermore, Berger emphasizes that his task is to develop an embedding formal language, not to develop “algorithms for embedding parts of natural language into such formal systems” (*ibid.*, p.203). My main worry with his demarcation of the project is that he allows himself to employ modeling resources for which it is doubtful that *any* such embedding could exist., which would make his system unilluminating given that his goal is to understand natural language. Take the anaphoric pronoun occurrence in

(30) He¹ was tall.

The occurrence of “He¹” should be interpreted as a variable parametrically covarying with the variable in the A-B sentence, e.g.

(31) A man entered.

The problem is that there *is* no variable in (31), and hence nothing with which “He¹” in (30) can covary. To be sure, we use variables *in first-order representations* of (31) to mark quantifier scope, but these variables are eminently dispensable. Since the quantifier in (31) scopes over “entered,” one may legitimately talk of “entered” being *bound* by “a man,; but to claim that the *natural language* sentence contains mysterious, covert *variables* seems close to confusing the formal model with the phenomenon (quantifier-scope) being modeled. Now, Berger may grant as much, and concede that talk of *covarying variable occurrences* must be reducible to something that does not require positing variable occurrences as “real” phenomena in (31), but we are not told how to do this. It is hence not clear how or even whether natural language sentences can be mapped into Berger’s formal system. Note – crucially – that the problem is not that Berger presupposes a syntactic level of Logical Form where elements can appear displaced from their surface position; the problem is that the variables are mere helpful devices at the level of logical form and do not correspond to any element on the surface at all.

Moltmann (2006) defends a static framework in which sentences express structured propositions, sequences composed of the meanings of primitive constituents. The propositions expressed by “Bob loves Jill” and “Bob loves every girl” are

(32) <LOVES, Bob, Jill>

(33) <EVERY(GIRL), < λx , LOVES, Bob, x >>

Structured propositions require separate definitions of truth conditions (*w* for world):

$$(32^*) \|\langle \text{LOVES}, \text{Bob Jill} \rangle\|^w = 1 \text{ iff } \text{LOVES}^w(\langle \text{Bob}, \text{Jill} \rangle) = 1$$

$$(33^*) \|\langle \text{EVERY}(\text{GIRL}), \langle \lambda x, \text{LOVES}, \text{Bob}, x \rangle \rangle\|^w = 1 \text{ iff } \text{GIRL}^w \subseteq$$

$$|\lambda x(\text{LOVES}(\langle \text{Bob}, x \rangle))|^w$$

An advantage of structured propositions (over, say, unstructured propositions), then, is that they may in principle be well-defined while truth-conditionally incomplete. On Moltmann's account truth-conditionally incomplete structured propositions can accommodate anaphoric connections, and have truth-conditions only relative to a truth-conditionally completing *background* of (a set of) structured propositions. Incomplete propositions may contain variable-like objects – parametric objects – that contain argument positions that connect them to a background context. Parametric objects are associated with anchoring functions, partial functions from parametric objects to actual objects. Parametric objects thus play a role similar to drefs in DRT, and are interpreted by quantification over assignment functions. In Moltmann's system, however, each sentence in a discourse expresses a (possibly truth-conditionally incomplete) proposition.

It is, however, an open question what such parametric objects *are*. It is tempting to treat them as mere placeholders that are replaced by real objects in the evaluation of sentences. Admittedly, this interpretation may capture an intuitive connection between how we understand expressions that function as variables, which are characterized by the absence of a commitment to any identified individual fulfilling the antecedent conditions, and truth conditions, in a manner that corresponds to how we understand the connection between constant expressions and truth conditions. A discrepancy between constants and variables remains, however. The value of a constant is the referent, and there is no

mediating layer between expression and semantic value. A variable expression, such as “he,” does not express a certain quantification or manipulation of assignment functions (or anchoring functions) on Moltmann’s view, but a *placeholder* that is subsequently interpreted as quantification over assignment functions. But then placeholders must occur at the level of *values*. Hence, we get a third level of placeholders between variable expressions and their values that we avoid in the case of constants. Furthermore, and parallel to the situation with anchored drefs in DRT, quantification over anchoring functions involves mixing model theoretic tools with semantic contents. The role of these model-theoretic resources (assignment functions) is to determine the truth of a sentence *relative* to contents, and they should as such not also be part of contents.⁷⁸

Neither Moltmann nor Berger provides anything that clearly serves the role of semantic values for variables. An opposite worry arises in certain *variable-free* system (e.g. Barker & Shan, 2008). As Barker & Shan show, the semantic contributions of expressions that are interpreted as bound, or as taking wide scope, in non-variable-free systems can be built locally without variable binding at all – indeed, Barker (2007) has shown that adequate accounts of either kind are expressively equivalent. If we employ variable-free semantics we build the values for drefs locally (and directly), and the semantic analysis yields a value for any occurrence of an anaphoric pronoun. However, since the values of anaphoric pronouns rely on the introducing antecedent this object

⁷⁸ More ontologically robust view of parametric objects include accounts that associate drefs with individuation conditions on when they are *shared* by different agents (Zimmermann, 1998), for instance, but these do not in general facilitate isolable semantic values for bound variable occurrences. One option that avoids this worry, and which we discuss below is to interpret variables by elements that can play the role of assignment functions, such as choice functions or Skolem functions.

would be rather complex, and the relationship between anaphoric pronouns and directly referential terms (insofar as these are semantically simple) is lost.

5.4. A Dynamic Pragmatic Account

In this section I outline an approach (I call it UIF) that seems to account for the behavior of anaphoric pronouns, which supplies isolable semantic values for variables, and which can be extended to descriptive names. The account is based on the definition of variables as identity functions in 5.3.2, but rather different from the interpretation of pronouns as identity functions in the variable-free semantics of Jacobson (1999, 2000). Insofar as I have defended the existence of natural language expressions that function as variables, it may appear surprising to employ a system that “eliminates” variables. First, however, the account interprets anaphora in a manner that ensures that in a first-order regimentation, an individual variable would be the only suitable substitution. Second, the values satisfy the constraints on variables identified in previous chapters (referential requirement, semantic simplicity, *de jure* rigidity). Third, elimination is to be expected. First-order variables have no values in bound occurrences, but receive them (at least in Tarski semantics) relative to assignments as steps in truth-definitions for quantified sentences. To provide *sui generis* values for variables it is hence clear that we have to replace them in our analysis.⁷⁹ (a similar distinction between pronouns and variables is drawn in

⁷⁹ Another point is that pronouns may not *literally* be variables bound by antecedent quantifiers in natural language (see Bach & Partee, 1984). Pronouns are often understood as co-indexed with a c-commanding element occurring in argument position, and are taken to be bound by the operator that identifies the pronoun with the argument position, generally an abstraction operator, in which case e.g. the verb phrase “praised his own efforts” should be interpreted as $\lambda x(x \text{ praised } x\text{s efforts})$ (see e.g. Szabolcsi, forthcoming,

Dekker, 2012, who nevertheless – as I do – treat pronouns as closely related to variables, at least with respect to distinctions that can be drawn in a first-order language). I provide a (relatively brief) and intuitive outline in 5.4.1, and a more thorough and formal explanation of the mechanisms in 5.4.2.

5.4.1 Updateable Identity Functions (UIF)

Consider first (34), which is represented as an open sentence with a free variable (35):

(34) He was tall.

(35) Tall(x)

The semantic value of (34) could initially be thought of as a function from values of “ x ” to truth-values. Above we suggested treating “ x ” as an identity function on a domain, a function that maps the domain onto itself. In the context of (31), this amounts to treating the value of “ x ” as an identity function – or more properly, an inclusion function – from the domain D associated with the sentence (the total domain of the world of evaluation) into the extension of the predicate; i.e., “tall” is a function that maps outputs of the identity function on D , to truth-values. The *truth* of “Tall(x)” thus depends on identity mappings from D to the extension of “Tall.” (And note: the function maps the *whole* domain, not just males; the *character* of “he” is “most salient male” – or perhaps “most salient males” – but the restriction to males does not show up in the semantic content.

for a discussion). This complication illustrates another gap between the first-order representation we assumed in previous chapters and the requirements of a proper linguistic treatment, and suggest that the correct analysis is better captured by e.g. the variable-free systems of Jacobson or Barker & Shan.

The character will later play an important different role.) (34) is, however, not yet truth evaluable. We have defined it as a composite function from $[\lambda x. x]$ on D into $\lambda y[y \text{ is Tall}]$ but have not yet said under which mappings it is true.

We can in principle give truth conditions for (35) when bound by a quantifier. In “some man is tall” (35) is bound by $\exists x: \text{Man}(x)$, and the result is true iff there is a partial identity function, or partial inclusion, from the domain of the context into the extension of “tall,” or – more precisely – if there is a function from an identity function from D to $\{\text{TALL}\}$, to a function from a function from an inclusion function from D to $\{\text{MAN}\}$ to truth-values, to truth-values, such that (roughly) $\{|\lambda y[\text{My}]|(|\lambda x. x|)|_D| \cap |\lambda y[\text{Ty}]|(|\lambda x. x|)|_D| \neq \emptyset\}$. If (35) is bound by “ $\forall x$ ” the result is true iff the function from the domain into the extension of “tall” is total. I present this as an optional move, however. In “some man is tall” variables are covert markers of scope not present in the syntax, so we are not obligated to represent them as identity functions or anything that needs a value.

If “ x ” is a *demonstrative*, then the input (and output) of the identity function is fixed. The *character* of the pronoun determines a referent in a context, and the expression is fixed to that input. Since the input is fixed to an extra-linguistic anchor, independently of world of evaluation, linguistic context or semantic operators, the value of the pronoun is a constant function from circumstances of evaluation to individuals.

The advantage of using identity functions is brought out when we apply them to variables that occur syntactically free. Suppose (34) is uttered subsequent to:

(36) A man entered.

Instead syntactically binding occurrences of “x” in an open sentence, previous discourse works as a *domain restriction*. (36) restricts the domain of “x” to the domain of entering men.⁸⁰ In the subsequent (34), the role of “he” is to connect the sentence to previous discourse, and it does so by taking as input *not* the whole domain, but the domain provided by (36). The *character* of “he,” “most salient male(s),” helps determine the correct interpretation, just as it does in the demonstrative case. But in (36) there is a salient *linguistic* antecedent, so instead of anchoring the pronoun to an extra-linguistic object the pronoun is connected to the individuals of the relevant class, namely the class of entering men. The character of “he,” then, can be represented as an identity function from contextually supplied domains to an identity function from domains of the context set to individuals. For *demonstratives*, the context is extra-linguistic, and “reset” the intensions associated with the terms – identity functions from domains to individuals – to contextually supplied values so that the outputs are the contextually supplied individual, regardless of world of evaluation. For *variables*, there is no one such input; the input varies with world of evaluation *of the antecedent*. Thus, the values supplied by the domain of the antecedent “reset” the intension of “he,” which means for instance that despite the fact that the input domain varies with world at which the antecedent is evaluated, the expression is nevertheless rigid relative to any world of evaluation for the antecedent (even if the pronoun is in the scope of a modal operator occurring after the antecedent). Thus, demonstrative and anaphoric pronouns work exactly the same way; it

⁸⁰ It is important, however, that the domain restriction is a restriction on potential satisfiers on anaphora, rather than a restriction of situations. (36) might *also* determine a set of possible situations, but the domain restriction effect (which is presumably also restricted by context; “entered” signals a particular location *l*, and the salient men are the ones that entered at *l*) is a different one. Relative to each situation or world in the context set, then, (36) makes the domain of entering men available for creating anaphoric links.

is the *contexts of utterance* that are different, and that difference has consequences for the truth conditions of sentences in which the pronoun occurs.

The result is that “he” in (34) is interpreted as an identity function from the domain of entering men, the individuals that satisfy the predicates in the antecedent, into the interpretation of TALL, a function from individuals to truth-values. No one individual (or set of individuals) is the value of the pronoun, since the input domains are different at different worlds in the context set. The role of the pronoun is to supply individuals (and only individuals) from an antecedent domain to the predicate of the anaphoric sentence.

Truth conditions and the maximality constraint

What, then, are the *truth conditions* of (34)? The answer depends on whether or not anaphoric pronouns are associated with a maximality constraint (e.g. Kadmon, 1990); that is, whether “he” in (34) requires that *all* entering men were tall (*maximality* is hence slightly weaker than *uniqueness*). I will not try to settle the question here. Suppose anaphora are *not* associated with a maximality constraint. Then “Tall(x)” is true if a partial inclusion function on tall things can be defined with the domain of entering men as input. If we subsequently add “he(x) was also strange,” “ x ” is a function from the updated domain of tall men who entered, and “Strange(x)” is true if there is an inclusion function from this domain into the class of strange things (i.e. that there is *some* tall entering man that yields *truth* for the “Strange” function). If anaphoric pronouns *are* associated with a maximality constraint, then (34) is true if the function from the domain of entering men to the class of tall things is *total*. UIF can accommodate both options, but the results are different.

First, one might wonder whether the absence of a maximality constraint is plausible beyond pronouns with singular indefinite antecedents (if at all). If different quantifier antecedents bring different interpretive requirements to an anaphoric pronoun, we may need discourse to keep track of this parameter as well, and there is no easy way of letting semantic or pragmatic features keep track of the quantificational force of the antecedent in addition to coordinating anaphoric links. I am, however, not convinced that a maximality constraint is required for *any* singular anaphoric pronoun.

Second, if pronouns are not associated with a maximality constraint, then the input domain for identity functions must be continuously updated in a discourse. If (34) is followed by “he was also old,” then the input domain for the latter must be the domain of tall, entering men, not only the domain of entering men. If only the domain determined by (36), *entering men*, were used, then the fragment “A man entered. He was tall. He was also old” would yield truth if two men entered, one tall and young, and one short and old. This is not a legitimate reading, and we must hence let domains be consecutively restricted in discourse. If pronouns *are* associated with a maximality constraint we may interpret all subsequent anaphora relative to the domain determined by (36) and avoid further updates. Third, if there is a maximality constraint, then “a man entered. He was tall” will *not* be semantically equivalent to “a tall man entered” since the former rules out any non-tall men entering and the latter does not. On the other hand, “every man loves his wife. He loves his mistress, too” *will* be equivalent to “every man loves both his wife and his mistress.” If there is no maximality constraint, the situation is the opposite. In many (most?) situations the difference is moot. Even with no maximality constraint, other pragmatic factors, situation restrictions or context restrictions, may intervene and leave

only a single element in the domain for the anaphoric pronouns in either case. On both options, the domain of the function associated with an anaphoric pronoun is the domain left by the antecedent, and context restricts domains associated with a particular anaphoric link represented by a variable.

UIF does not treat open sentences as syntactically bound, but as, in effect, semantically bound. We can embed UIF in context sets defined as maximal sets of worlds consistent with whatever presuppositions are at work (Stalnaker, 1998). Discourse contributions diminish the set – restrict the context – by removing worlds inconsistent with that contribution. The role of an indefinite expression is not only to cut worlds, however, but to initiate anaphoric chains by making restrictions on domains *associated with that chain* available for the interpretation of later contributions containing anaphora. We can think of updateable predicate extensions as separate context parameters. A sentence $Q\Phi x$ updates context by removing the worlds with respect to which $Q\Phi x$ is false, but also makes the domain of Φ available. The domains supplied by Φ vary with worlds in the context set, so we cannot associate the pronoun with a particular referent (even if there is only one available value in the actual world). Subsequent contributions update the domain restriction relative to all (remaining) worlds in the context set. A world w is cut if there is no way of satisfying a contribution containing a pronoun in w .

Predictions

UIF allows us to provide semantic values for all expressions, including anaphoric pronouns. Furthermore, as opposed to DRT and DPL each contribution to a discourse denotes primarily truth conditional content and only derivatively context updates since

updates are determined in part by the truth conditional content. We cannot without further ado associate each contribution with a complete, Russellian proposition, but we may, if we wish, follow Moltmann and associate them with incomplete Russellian propositions that are truth-conditionally complete relative to a background (a domain restriction). In the remainder of this section I present some results achieved in UIF. In the next, I provide a more precise account of the mechanisms.

I) We predict the right results for “few philosophers came to the party. They had a good time.” The interpretation cannot be $[Few(x): Philosophers(x)] (Came\ to\ the\ party(x) \wedge Had\ a\ good\ time(x))$, since this would be true if many philosophers showed up but most of them did not have a good time, which the original fragment is not. Instead, we take “few philosophers came to the party” to restrict the context set by removing all worlds where this sentence is false, and “they” to be an identity function on the domain of philosophers who came to the party (with respect to every world remaining in the context set). Note that we obtain the correct results regardless of whether the anaphoric pronouns are associated with a maximality constraint, as long as we avoid syntactic binding. We also obtain the correct interpretations of (9) and the sage plant example (15).

II) UIF predicts that “no” cannot in general take anaphora. Since “no Φ is Ψ ” is $\{|\Phi \cap \Psi|\} = \emptyset$, the input domain for the anaphoric pronoun is empty and the identity function is undefined. But this is correct. As illustrated by (4), “no” does not license anaphora. Now, certain contexts may appear to license such anaphora, but if e.g. “Harry bought no donkeys; they looked sick,” is legitimate it is because of a scope ambiguity in the antecedent. The second sentence is legitimate if the first sentence is interpreted as

“there are some donkeys, but Harry bought none of them,” in which case the donkeys that Harry didn’t buy are, indeed, available as input for anaphora.

III) We predict that variables are rigid. In $\exists x(Fx \wedge \Diamond \neg Fx)$ the first conjunct makes available a domain of Fs . The variable in the second conjunct is interpreted as a function from the domain associated with the first conjunct into the domain associated with $\neg F$ at the world of evaluation for the counterfactual. In standard semantics each world has a domain, but the intended interpretation of $\exists x(Fx \wedge \Diamond \neg Fx)$ is one relative to which the values of “ x ” are the same (since the second conjunct is satisfied only if some elements of the antecedent domain also occur in the domain of the world of evaluation for the second conjunct), which means that “ x ” is rigid.

IV) We get the correct results for the following:

(23) Out of ten marbles, one is not in the bag. It must be under the couch.

(24) Out of ten marbles, nine are in the bag. It must be under the sofa.

(23) is felicitous; (24) is not, even though they are truth-conditionally equivalent. Our system predicts the correct result. In (23), but not in (24), the domain of missing marbles have been made available as an input domain for the function associated with “it.” We require that anaphora have explicitly given antecedents. Hence, we also predict that (37), but not (38), is felicitous (Elbourne, 2005, p.54):

(37) Every man who has a wife is sitting next to her.

(38) *Every married man is sitting next to her.

Many D-type systems for anaphora interpretation have problems explaining why (37) but not (38) are grammatical (on a reading on which they are truth-conditionally equivalent), since the link between antecedent and anaphora is treated informally. In UIF no input domain for “her” is defined in (38), so “her” is undefined as well. (We deal with donkey sentences such as (37) below.)

5.4.2. A Pragmatic Interpretation⁸¹

(34) should be true iff for all assignments $s_{Dem}, x_{s_{Dem}} \in I(\text{“tall”})$, where D_{em} is the set of individuals satisfying the antecedent, if we associate the pronoun with a maximality constraint; and iff for all assignments s_{Dem} (relative to D_{em}) there is an assignment s^*_{Dem} that differs from s at most in what individual it assigns to x , and $x_{s^*_{Dem}} \in I(\text{“tall”})$, if we do not. Domain restrictions are, in practice, supposed to serve the role of assignment restrictions in the accounts discussed above. Indeed, in the simplest cases the effect is essentially to ensure that predicate extensions are properly related. (34) is true if $\{M \cap E\} \subseteq \{\text{TALL}\}$, given a maximality constraint, or if $|\{M \cap E\} \cap \{\text{TALL}\}| \neq \emptyset$, if not. The feature of natural language we are specifically trying to capture in UIF, however, is how the *pronoun* facilitates the correct way of relating sentences in a discourse.

The domain for the identity function is defined by previous discourse. To obtain a systematic account we let the restrictions take place in a pragmatic layer, the *salience structure* of a discourse. The salience structure is a (partial) ordering (“by salience”) of the predicate extensions of the total domain of individuals, or – perhaps more accurately

⁸¹ In the following I freely mix set-theoretic and model-theoretic notation; this may be remedied if desired.

– of all subsets of the domain (we provide more detail below). The salience structure is updated throughout discourse in different ways to make new sub-domains available as inputs for subsequent sentences. We call the most salient input set at any given stage the *Topic* (T) of the discourse at that stage. *Topic* is, simply, a set of individuals. The *context set* of a discourse is, as usual, a set of worlds w consistent with the information at a given stage, but we add salience structure as a parameter. The salience structure defines a *Topic* for every world, so a context set is a set of pairs $\langle w, T \rangle$. Since the extensions of predicates generally vary between worlds, so does *Topic*. A discourse contribution can update a context both by deleting pairs in which the world is inconsistent with the contribution from the context set (intersecting the set of worlds with the proposition), and update the salience structure to yield a new *Topic* for each world.

For simplicity I disregard time and other parameters in the following (for the most part I also disregard worlds). Second, I will also talk about *Topic*-updates, or T -updates, even though it is strictly speaking the salience structure that is updated to yield a new T at every stage in a discourse. An anaphoric pronoun is defined on the at any time salient *Topic* of the discourse. A third, potentially more significant simplification is that we treat pronouns, interpreted as identity functions, as belonging to grammatical category e (hence, the notation $\|\lambda x. x\|$ for values of $\lambda x. x$). Strictly speaking it may be more appropriate to assign functions to type $\langle e, e \rangle$. If we do, we must adjust other grammatical categories accordingly, and may need to introduce type-transformation rules more or less

parallel to the type-shift rules of Jacobson (1999).⁸² An important motivation would be to attain strict compositionality over the *discourse* (as presented *UIF* only allows strict compositionality for single sentences). Finally, I will assume that anaphora are *not* associated with a maximality constraint unless otherwise specified.

Basic semantics for UIF

The semantics are static. We define a model in the usual manner, $\mathcal{M} = \langle D, I \rangle$. We also need a sequence of stages of discourse, k_0, \dots, k_n . If a sentence is satisfied at a stage it is true. We also have a sequence of *Topics*, T_0, \dots, T_n . The purpose of T is to provide inputs for anaphora. Although all sentences update T , T plays no role in the interpretation of sentences that do not contain anaphoric pronouns. At stage k_0 , $T_0 = D$. As a preliminary, we can associate T_n with stage k_n , but as we will see the relationship is more complicated. At any stage k_n , $T_n \subseteq D$. The truth-definition for an atomic sentence is:

⁸² In Jacobson's (1999, 2000) variable-free semantics, anaphoric pronouns are interpreted as identity functions, though she avoids giving separate values for pronouns altogether. Roughly, when a predicate $\lambda y(Fy)$ combines with a pronoun $\lambda x. x$ the pronoun is either:

i) Subjected to the *g*-rule, a transformation rule for composing functions to yield a phrase of type $[\lambda f_{\langle e, e \rangle}. \lambda z.(F(f(z)))](\lambda x.x)$, which is equivalent to $\lambda z(F(z))$. In other words, "he is tall" becomes equivalent to "tall," but the function denoted by "tall" can be applied to any salient object in the context.

ii) Subjected to the *z*-rule, a transformation rule for composing functions to yield higher-order functions that are appropriate input to quantifiers; for instance, "he suspects he is ill" is (disregarding any semantic complications with attitudinal operators) $[\lambda P_{\langle e, t \rangle}. \lambda x(\text{Suspects}(x, P(x)))](\lambda y(\text{Ill}(y)))$, which is equivalent to $\lambda x(x \text{ suspect } x \text{ is ill})$, and the appropriate input to a quantifier head such as "some man."

When the pronoun occurrence is anaphoric but free, we apply both rules. First, we apply the *g*-rule to the pronoun to yield $[\lambda f_{\langle e, e \rangle}. f]$, which provides an open slot for a function of type $\langle e, e \rangle$. With respect to (34), "is tall," $\lambda z(\text{Tall}(z))$, subjected to both rules, is a function taking $[\lambda f_{\langle e, e \rangle}. f]$ and passing it upward. By passing the slot $\langle e, e \rangle$ upward the denotation of the sentence will be a function taking arguments of type $\langle e, e \rangle$, which is applied to the (salient) function that w.r.t. (36) maps men to the activity of entering. On Jacobson's account anaphoric pronouns make no direct semantic contribution to the interpretation, but initiates type transformations that enable us to bind sentences in a discourse together in the right way. A benefit of the analysis is that it is fully compositional. The semantic object that is the value of a sentence in a discourse will become increasingly complex as discourse unfolds, however, but that is presumably to be expected from a fully semantic interpretation of discourse (i.e. without a pragmatic layer of salience).

$$\mathcal{M}_k \models Fa \text{ iff } I(a)_{T_{k-1}} \in I(F)_{T_{k-1}}$$

T is inert in this case, so the definition reduces to $\mathcal{M}_k \models Fa$ iff $I(a) \in I(F)$. Next, we define atomic sentences with free variables:

$$\mathcal{M}_k \models Fx \text{ iff there is an } x \in T_{k-1} \text{ such that } \|\lambda(x)_{T_{k-1}}. x\| \in I(F).$$

In general, for a relation with n free variables:

$$\mathcal{M}_k \models R\langle x_1, \dots, x_n \rangle \text{ iff there is an } x_1, \dots, \text{ there is an } x_n, \text{ such that}$$

- There is a sequence $s \in T_{k-1}$ such that $x_1 \in s, \dots, x_n \in s$ and

$$\langle \|\lambda(x_1)_{s \in T_{k-1}}. x_1\|, \dots, \|\lambda(x_n)_{s \in T_{k-1}}. x_n\| \rangle \in I(R), \text{ if } T_{k-1} \text{ is a set of sequences.}^{83}$$

- There is an $x_1 \in T_{k-1}, \dots, \text{ there is an } x_n \in T_{k-1}$ such that

$$\langle \|\lambda(x_1)_{T_{k-1}}. x_1\|, \dots, \|\lambda(x_n)_{T_{k-1}}. x_n\| \rangle \in I(R), \text{ otherwise.}$$

For non-atomic sentences (non-commutative conjunction; material conditional):

$$\mathcal{M}_k \models \neg \Phi \text{ iff } \mathcal{M}_k \not\models \Phi$$

⁸³ T may be a set of sequences of individuals rather than individuals. In that case the ordered sequences $\in I(R)$ may or may not be permutations of the sequences in T , but any two elements in a particular sequence satisfying R must also be elements of a particular sequence of T . The motivation is as follows. In “a man loves a woman. She rejects him,” we need to ensure that the men rejected by a woman x are men who loves x (and not some different woman). The second condition applies in cases where no such coordination is needed, as in “a man entered. A woman left. He saw her,” which is true if some entering man saw some leaving woman (given the absence of a maximality constraint). We return to these issues below.

$$\mathcal{M}_k \models \Phi \wedge \Psi \text{ iff } \mathcal{M}_k \models \Phi_{T_{k-1}} \text{ and } \mathcal{M}_{k+1} \models \Psi_{T_k}$$

$$\mathcal{M}_k \models \Phi \rightarrow \Psi \text{ iff } \mathcal{M}_k \not\models \Phi \wedge \neg \Psi$$

Quantifiers can be defined the ordinary manner. We may (to circumvent familiar concerns such as the island escaping properties of indefinites) define “an F is G ” as an open sentence $\lambda z[Fz \wedge Gz](\|\lambda x. x\|)$ for which the input is (always) the whole domain D (although open, the sentence does not contain pronouns, and should hence not invoke T).

$$\begin{aligned} \mathcal{M}_k \models \text{“an } F \text{ is } G\text{”} &\text{ iff there is an } x \in D, \text{ such that } \|\lambda(x)_{D}. x\| \in I(F) \text{ and } \|\lambda(x)_{D}. x\| \\ &\in I(G) \text{ [i.e. iff } \{F_D \cap G_D\} \neq \emptyset] \end{aligned}$$

Suppose the sentence following (36), uttered at stage k_n , is (34), uttered at stage k_{n+1} . At k_{n+1} the relevant *Topic* is T_{kn} , the set of entering men. “he” is a function $\lambda[(x)_{T_{kn}}. x]$. Any world in which $|\{\lambda y[\text{TALL}(y)](\|\lambda(x)_{T_{kn}}. x\|)\}| = \emptyset$ is removed from the context set, and T is updated by intersecting T_{kn} with $\{\text{TALL}\}$ to $T_{kn+1} = \{M \cap E\} \cap \{\text{TALL}\}$. In short, to interpret “he is tall” we draw on the current *Topic*, and subsequently raise the class of *tall, entering men* to salience (replacing the set of *entering men*, though it may of course be the case that $T_{kn+1} = T_{kn}$). If anaphoric pronouns are associated with a maximality constraint, however, then (34) is simply a test – it is true iff a total identity function can be defined for T_{kn} such that the values are in the class of tall things. It updates nothing.

Suppose, on the other hand, the sentence following (36) is (39):

(39) A woman left.

(39) updates *Topic* as well, by raising the class of leaving women to salience. Since (39) contains no anaphoric pronoun the current *Topic* is otiose, and (39) defines a new one.

Salience structure

To account for *T*-updates, remember that *T* is the *most salient class* of individuals of *D*. We thus update the *T* relevant for interpreting a sentence by reordering the salience structure and raise a new *T* to salience. We define the salience structure as the *powerset* of the domain, $\wp D$. Each stage *k* defines a partial order on $\wp D$, and use this order to define the class of individuals to be used to interpret a subsequent free variable.

For instance, (36) is true iff $\{|M \cap E| \neq \emptyset\}$, so all $\langle w, T \rangle$ pairs such that $w \models \{|M \cap E| = \emptyset\}$ are removed from the context set. In addition, (36) updates *T* for all remaining pairs in the context set. (For perspicuity I use *ME* instead of $\{M \cap E\}$). At k_0 :

$$\wp D_{k_0} = \{P_1, \dots, I(ME), \dots P_n\}$$

ME updates the salience structure by pairing $\wp D$ with $I(ME)$. Since $I(ME) \subseteq \wp D$, we get:

$$\wp D_{k_1} = \{I(ME), \{P_1, \dots, I(ME), \dots P_n\}\}$$

In effect, (36) creates what amounts to a (very) partial ordering of $\wp D$ with $I(ME)$ as the most salient class, and $T_{k1} = I(ME)$. Note that salience structure updates do not *remove* anything from a context set, but reorganize (the powerset of) the domain of individuals.

T -updates set T to the most salient domain. The intersection of T with a predicate, as in (34), happens when the contribution is an open sentence Fx , in which case x signals a mapping from T , leaving $\{I(F) \cap T\}$ as the new *Topic*. Suppose (36) is instead followed by (39). (39) updates the salience structure by raising the class of leaving women to salience, replacing rather than intersecting T . It will, however, be useful not to lose track of previous *Topics*. Let us say that “a woman left” is true iff $\{|WL| \neq \emptyset\}$. The salience structure prior to (39) is $\wp D_{k1} = \{I(ME), \{P_1, \dots, I(ME), \dots, I(WL), \dots P_n\}\}$, and to capture the correct ordering, we update this to

$$\wp D_{k2} = \{I(WL), \{I(WL), \{I(ME), \{P_1, \dots, I(ME), \dots, I(WL), \dots P_n\}\}\}$$

Suppose (39) is followed by “he is tall.” Since *leaving women* is currently the salient class, “he” is undefined (given the character of “he”) on this class and requires pragmatic resolution. By keeping a record of previous updates in the salience structure (by modeling *raising to salience* as consecutive orderings) we at least retain an intuitive idea of how to accomplish such resolutions, namely by moving to the second most salient T – the most salient possible antecedent class for the pronoun – in this case *entering men*.

Dynamic Topics

Though the semantics of UIF are static, T -updates are dynamic. As opposed to other dynamic systems, the updates are defined by the *interpretation* of previous sentences, and cannot without further ado be defined “directly” from previous updates (as in DPL). To refine the system and achieve more expressive power, we associate with *each one-place predicate* a pair of *Topics*, $\langle T_n, T_{n+1} \rangle$.⁸⁴ Thus, a sentence denotes a *sequence* $\langle T_n, \dots, T_{n+m} \rangle$ rather than a pair – after all, it is the predicates of a sentence s that defines the output *Topic*; the determiners of s are of course crucial to defining the truth conditions of s , and hence which worlds remain in the context set, but – at least in these simple cases – the operations on *Topic* for “at least one man entered” and “exactly one man entered” are the same, yielding the set of entering men.⁸⁵ Strictly speaking, then, (36) is associated with a pair of updates, $\langle T_n, T_{n+1} \rangle$ for “man” and a pair $\langle T_{n+1}, T_{n+2} \rangle$ for “entered.” As a shorthand, we say that (36) denotes the update $\langle T_n; T_{n+2} \rangle$. The truth of a sentence at a stage k is hence truth relative to a sequence of *Topics*, starting at the final output T at $k-1$.

For any output, if T_{n+1} is the output of $I(F)$ relative to T_n , then either $T_{n+1} = I(F)$, or – if F is associated with a free variable – a restriction $\{x \mid x \in T_n \text{ and } x \in I(F)\}$. (36) is true relative to T iff $\{|M \cap E| \neq \emptyset\}$, or, in first-order terms (if we assume that “an” is the existential quantifier), true relative to an assignment g iff there is a variant assignment g^* that differs from g at most in what element e it assigns to x , and e is a value of $\lambda y[M y \wedge$

⁸⁴ It *may* thus be possible to make the predicates themselves dynamic, i.e. denote *pairs* of extensions, and avoid relativization to *Topics* at a separate level of discourse. I do not explore this option here.

⁸⁵ For “every man entered” we could treat the T -update as the same as well, though since “every man entered” is true iff $\{M \subseteq E\}$, $\{M \cap E\} = M$. It is striking that e.g. monotonicity properties seem to have no impact on updates. We deal with *plurals* below. “No man entered” does not license anaphora since all worlds relative to which T is defined are removed.

$\lambda z[Ez](y)$]. Dynamically, (36) updates T to $T_1 = \|\lambda y[My \wedge \lambda z[Ez](y)]\|$. If (36) is followed by (39), denoting $\langle T_1; T_2 \rangle$, then $T_2 = \|\lambda y[WL(y)]\|$. Suppose (39) is then followed by:

(40) She was quiet.

(40) denotes $\langle T_2; T_3 \rangle$, but the interpretation requires reference to $T_2 = \|\lambda y[WL(y)]\|$, such that there is an $x \in T_2$ and $\|\lambda(x)_{T_2}. x\| \in I(Q)$. In effect, then, the interpretation of Q amounts to $\lambda u[\lambda y[WL(y)](u) \wedge Qu]$, so $T_3 = \|\lambda u[\lambda y[WL(y)](u) \wedge Qu]\|$.

Binary Predicates

Natural language discourse allows multiple drefs to be activated simultaneously as long as there is some way of keeping track of them. This is the role of the character of a pronoun. One may, at least in general, not have drefs denoted by the same pronoun (type) activated at the same time⁸⁶ (indeed, as we see in Chapter 6, this is an important motivation for introducing descriptive names). If we did the character of the pronoun would not suffice to locate the appropriate antecedent *Topic* in the salience structure. As a general rule (and in accordance with the semantics for open sentences above), pronouns are interpreted in light of the current *Topic*, but we may (pragmatically) allow the character to retrieve *Topics* further back in the salience structure if needed.⁸⁷

⁸⁶ We can deal with exceptions (“a philosopher met a linguist. He liked him”), however, by pragmatically retrieving earlier *Topics*. There will be some indeterminacy (the second sentence is ambiguous), though the dependency relation (see below) suggests that “he” is the philosopher and “him” the linguist.

⁸⁷ There is a rather obvious theory of definite descriptions defined in terms of familiarity to be developed here as well. Instead of implying uniqueness, the role of (at least anaphoric) definite descriptions is to index elements far back in the salience structure. Admittedly, since the elements would need to be introduced, *discourse initial* definite descriptions would require a different treatment).

Now, the extension of a binary predicate is a set of ordered pairs. Suppose that, instead of (40), (39) is followed by

(40*) She avoided him.

To interpret (40*) we need to engage in some pragmatic repair. As suggested above, (40*) can access both the most salient *Topic*, $T_2 = I(WL)$ and the slightly more distant $T_1 = I(ME)$.⁸⁸ Hence, we interpret “she” relative to the former and “him” relative to the latter. The choice of antecedents is guided by the characters of “she” and “him,” respectively. “him” is interpreted as an identity function $\lambda(y)_{T_1}. y$, on the class of entering men, which is the input to the class of things being avoided. “She” is a function $\lambda(x)_{T_2}. x$, on the class of leaving women to the class of individuals avoiding outputs of $\lambda(y)_{T_1}. y$. In terms of updates, (40*) denotes $\langle T_2; T_{3*} \rangle$, though since “avoided” is a binary predicate, T_{3*} should be the set of pairs $\langle u, v \rangle \in A$ such that $u \in I(WL)$ and $v \in I(ME)$. To achieve this we need the Cartesian product $\{T_2 \times T_1\}$, and let T_{3*} be the set of resulting pairs $\langle x \in T_2, y \in T_1 \rangle$ that are in the extension of A .

T_{3*} licenses the anaphoric pronoun “they” (as in “they didn’t talk,” which is true if there is no pair of T_{3*} in the extension of “talk”), but (40*) may also be followed by:

(41) He saw her.

⁸⁸ One may alternatively let a *Topic* for each activated pronoun be salient at any time (one T for males, one for females, etc), but we would still need to appeal extralinguistic factors for the proper coordination.

For (41), the input domain is T_{3*} , a set of pairs. The interpretation of “her” must pick the first elements of the pairs of T_{3*} ; that is, the leaving women who avoided an entering man. The domain of “he,” on the other hand, is the second element of the pairs of the domain of T_{3*} , but the legitimate values of “he” are restricted to those outputs y such that there is an x , $\langle x, y \rangle \in T_{3*}$, and $\langle y, x \rangle \in \text{SAW}$ – that is, the members of the output *Topic* are permutations of the elements of T_{3*} . “he” is a *dependent* variable. “he” is associated with an identity function, but in this case its range is restricted to those values v such that for some u , $\langle u, v \rangle \in T$. The easiest way to ensure the correct interpretation is to take the permissible inputs to a binary predicate (e.g. “saw”) to be the pairs of T or the set of inverse pairs (but not both) of T . However, for compositionality reasons we wish to allow the variables to be interpretable one at a time. “saw” is here a function from an identity function from the first elements of T to things seen x , to a function from an identity function from the second elements y of T such that Txy , to truth values.⁸⁹

If the sentence following (40*) were e.g. “he rose,” this restriction would be otiose. In (41), however, the dependency condition requires that the resolutions of one pronoun depends on the resolution of the other, to the effect that the legitimate input to SAW are only of the pairs $\langle z, w \rangle \in I(\text{SAW})$ that are permutations of the pairs $\langle w, z \rangle \in T_{3*}$. Similarly, for

(42) Some girl likes every boy. (43) He likes her, too.

⁸⁹ In “she saw him” we would, similarly, have a function from an identity function from the *second* elements of T to things seen y , to a function from an identity function from the *first* elements x of T such that Txy , to truth-values. The dependence relation is, in other words, symmetric.

The output T of (42) is the pairs of boys and girls defined by the “like” relation such that the first element is a girl x who likes every boy and the second element a boy y liked by x . In (43) “her” is, once again, the girls who like every boy. “he” is a dependent variable. For any x such that x is the output of the “her”-function, y is an acceptable input iff y is paired with x in T (and the output T is the set of such pairs that are also in {LIKES}).

Conditionals may require more resources than we have developed; conditionals with existential antecedents appear to have unselective force. Take

(44) If a farmer owns a donkey, he beats it.

The antecedent updates T , which is used to interpret the consequent. If anaphoric pronouns are associated with a maximality constraint, we obtain the correct interpretation immediately. If not, we obtain the correct anaphora resolution, but may need to introduce universal or unselective quantification over situations or events to obtain the correct generality. Several such extensions would be consistent with our framework.

Since every predicate is associated with a T -update, in (45)

(45) Every farmer who owns a donkey beats it

“owns a donkey” denotes a T -update to a set of pairs of farmers who own donkeys and donkeys they own. The inputs to “it” are these pairs, and “it” is dependent. For any x such that x is a donkey-beating farmer, a donkey y is an acceptable input to the “it”-function iff y is paired with x in T . Intuitively, if we associate anaphora with a maximality constraint,

any farmer x who owns donkeys must beat all the x -owned donkeys, so we get a *strong reading* of the sentence (Kanazawa, 1994). If we do not associate anaphora with a maximality constraint, we may seem to get the *weak reading*, namely that every farmer must beat at least one of the donkeys he or she owns.

However, things are more complicated. The input for “beats it” is a set of “owns”-pairs, and, since the input for “it” is dependent, it must be restricted to pairs that are also ‘owns’-pairs. But there is no anaphoric pronoun for farmer; “ x beats y ” is syntactically bound (in the x -position) by “every farmer who owns a donkey.” Now, farmers who own more than one donkey appear in more than one pair in T . If (45) has a weak reading, we allow some pairs from T that are legitimate inputs to “beats” to fail to satisfy “beats.” But should we? The question amounts to whether the universal binder applies to *every occurrence* of a farmer in an owns-pair supplied by T (that is, every token of a donkey-owning farmer), or whether it is sufficient that every different farmer that occurs in T is paired with some donkey of his in the interpretation of “beats.” That is, does the universal quantifier bind the same farmer as many times as he occurs in a farmer-donkey pair, or only (at least) once per individual farmer? If the former, which seems plausible, we get the strong reading. (Tarski semantics and satisfaction relative to all sequences may suggest the strong reading.) Perhaps the answer is less than fully determinate for natural language. Though the strong reading is plausible in (45) it is less obviously plausible in “every student who had an umbrella brought it to school.”

Plurals and Constants

Let us assume, following Boolos (1984, 1985) that (46) receives a first-order interpretation as (47), where \angle is an inclusion relation and $\exists xx$ a plural quantifier (other ways of treating plurals will presumably yield equivalent results):

(46) Some men entered.

(47) $\exists xx (\forall y (y \angle xx \rightarrow (My \wedge Ey)))$

Plurality should be recorded in the *Topic*, for instance as $\{y \angle xx\} \subseteq \{M \cap E\}$, or by recording it directly in the predicate: Let the plural inflection be of type $\langle\langle t, e \rangle, \langle t, e \rangle\rangle$; that is, a function $\lambda P \lambda x [Px : Px]$, where P is a “plural transformation” that takes a predicate $\lambda x [Fx]$ and turns it into a predicate $\lambda x [Fx : Px]$ (with (47) being the truth conditions).

(48) They were tall.

“They” is an identity function on xx (or $\{\{Mx : Px\} \cap \{Ex : Px\}\}$) that maps *pluralities* of entering men into “tall,” which likewise must be subjected to a plural transformation.

“Tall” is distributive, and is ascribed to the individuals $x \mid x \angle yy$ who entered. (Collective predicates, e.g. “lifted a piano,” are ascribed to the plurality.) Suppose (46) is followed by

(49) Some of them were tall.

“Some of them” is a quantifier of type $\langle t, \langle t, e \rangle \rangle$. The determiner “some” is of type $\langle \langle t, \langle t, e \rangle \rangle, \langle t, e \rangle \rangle$ and “them” is of type $\langle e \rangle$. The relational preposition “of” is thus of type $\langle \langle t, e \rangle, e \rangle$. However, “of” is here used in a partitive sense. Semantically, given our interpretation in terms of pluralities, ‘of’ is $\lambda x x \lambda y [y \subset x]$. The quantifier phrase “some of them” thus maps *elements* of the output of the identity function on pluralities to “tall;” it does not map pluralities into “tall” directly. Such constructions indicate that assigning independent semantic values to anaphora is advantageous. In particular, since “tall” is predicated *not* on the elements of T (the plurality of entering men), but of elements of the pluralities of T (individual men), we should perhaps expect (49) *not* to update T with respect to *the pluralities* (but only for the individual men). If (49) is followed by “some of them were short” the antecedent for “them” would in that case be T , as defined by (46), not (49).

Finally, we should let atomic sentences containing constants provide updates as well. Fa denotes an update $\langle T_n, T_{n+1} \rangle$, where $T_{n+1} = I(a)$. For the subsequent “he is G ” the input T_{n+1} provides the interpretation for “he,” which is an identity function on T_{n+1} into $I(G)$. Thus, “he is G ” turns out to be true iff $I(a) \in I(G)$, which is precisely what we want.

Adding maximality constraints

If we associate anaphoric pronouns with a maximality constraint we obtain a simpler system. We would still need each sentence to update the *Topic* of a discourse, but we would no longer need sentences containing anaphoric pronouns to produce outputs that consist of the intersection of the input *Topic* with the predicate of that sentence. Instead, we let sentences *not* containing anaphoric pronouns (and perhaps sentences not containing definite expressions, including names and definite descriptions) update *Topic*

by substitution the current *Topic* with a new one. Sentences that *do* contain definite expressions, on the other hand, do *not* update *Topic*, but act as tests (i.e. eliminate worlds from the context set). The semantic clause for open sentences would require that the identity function from T is total:

$$\mathcal{M}_k \models Fx \text{ iff for all } x \in T_{k-1}, \|\lambda(x)_{T_{k-1}}. x\| \in I(F).$$

It is otherwise straightforward to modify UIF to associate pronouns with a maximality constraint. A benefit would be that we immediately obtain the correct (universal) reading of conditionals. We would still need a dependency constraint for pronouns when *Topic* is a set of ordered pairs, and it may be less obviously possible to obtain a weak reading of donkey sentences. Otherwise, the solutions and mechanisms would be as above.

5.5 Conclusion

There are alternative accounts that may suit our purposes as well. Interpreting pronouns using choice functions or Skolem functions (e.g. Winter, 1997) provides pronouns with semantic values that can do the work of assignment functions. To interpret a sentence containing an indefinite, a choice function selects a member of the extension class of its noun. The selection updates the salience structure or information states to ensure that the value (which serves as a test) of a subsequent definite expression, e.g. an anaphoric pronoun, coincides with the value of the indefinite. On dynamic views with choice functions neither definites nor indefinites have quantificational force on their own, but

are interpreted using variables over functions that assign individuals to predicates. Since these function variables are existentially closed, definite and indefinite expressions will fundamentally have quantificational force in an interpretation.

Heusinger (2004) and Peregrin & Heusinger (2004), employ the ε -calculus, a conservative extension of the predicate calculus. ε -terms are non-deterministic, and leaves it open which of the objects of a domain a particular term stands for. (ε -terms represent a version of the axiom of choice and can be interpreted as choice functions.) Peregrin & Heusinger interpret the salience structure of a discourse as a “global” choice function f defined for every predicate by “choosing” an arbitrary representative from the extension of each predicate. Indefinite expressions are “local” choice functions defined over the extension of its constituent noun, and update f by making the locally selected representative the value that f yields for that predicate (as well as raising it to salience). Definite expressions, including anaphora, are in turn interpreted by the updated f , which ensures that anaphora designate the selection of the antecedent indefinite. Definites are thus “ ε -tests,” whereas indefinites are ε -updates, and f serves as both a model for discourse salience and to ensure that the context-change potential of expressions interact in the correct manner. Consecutive updates of f enable us to coordinate drefs in stretches of discourse involving multiple representatives of the same predicate. For a mapping of the system into a dynamic framework, see Heusinger (2004).

Although the ε -calculus is expressively richer than the ordinary predicate calculus, $F(\varepsilon xFx)$ is truth-conditionally equivalent to $\exists xFx$, and a sentence $\Phi(t)$ where t is an ε -term defined by a singular existential quantifier over a predicate F occurring in a sentence prior to Φ , is first-order equivalent to $[\exists xFx](\Phi x/t)$. It is plausible to interpret

descriptive names using ϵ -terms as well; that is, as terms interpreted by functions defined over the extension of the name-introducing predicate. “Julius” is then interpreted by the choice function that selects an (in principle) arbitrary member from the extension of “inventor of the zip.” The drawback for our purposes is that it becomes somewhat harder to account for the analogies between the values of constant terms and variable terms.

Regardless of what formal machinery we employ it will need to be sensitive to a distinction in natural language corresponding to the distinction between variables and constants in chapter 3. Of course, to interpret variable-occurrences of *names* most of the complications discussed in this chapter are avoided in any case. Names usually have uniqueness-implying antecedents (though there are some exceptions, such as the “Bambi” example discussed in Chapter 4). The main difference between such names and constant names is sensitivity to world of evaluation. In a dynamic model of discourse with context sets, a descriptive name “Julius” will be satisfied by different individuals at different worlds (though rigid relative to all of them). A constant name, on the other hand, is determined in advance of evaluation, and in advance of identifying a context set. In the next chapters I return to names. In Chapter 6 I describe distinctive properties of variable names, and the differences between such names and ordinary pronouns such as “she.” In Chapter 7 I use the flexibility of the account to solve certain long-standing problems in the philosophy of language, in particular the problem of the contingent *a priori* and the problem of empty names.

CHAPTER 6: PLURALISM IN CONTEXT

Perhaps it is a bit misleading to say that we are defending *pluralism*. The fundamental idea is that a sentence “*n* is *F*” is radically *context sensitive*. It receives an interpretation only relative to a context. From the point of view of grammar, it is an open sentence – grammatically issued constraints on name interpretation require that “*n*” is semantically simple and takes individuals of a domain as values. Context determines the interpretation either by supplying a referent, for instance through a name-using tradition, or by binding. In the latter case the character of the name can (as we will argue) facilitate the projection of a binding presupposition if no antecedent is explicitly available. Nevertheless, names prefer reference over binding. For ordinary pronouns, preferences are reversed, and they prefer binding over accommodation. The efficacious difference is *character*. The characters of ordinary pronouns require the most salient individual (with few additional constraints). The characters of names do not issue a salience condition (which is also why presupposition projection is possible). Importantly, however, it is misleading to speak of two types of names. Rather, there are two types of *context*, linguistic and extra-linguistic, and the semantics of the *occurrence* of a name depends on what context is relevant for the interpretation.

A special characteristic property of names is that they can be used to retrieve referents beyond the original context in which the name was introduced. Demonstrative pronouns cannot retrieve referents once the referent disappears from the physical context, and ordinary anaphoric pronouns cannot sustain anaphoric links outside of a context.

Insofar as names are semantically closely related to pronouns they are best characterized

as “special duty pronouns” (Sommer, 1982). Whereas we have a limited stock of ordinary pronouns to retrieve salient drefs in a discourse, we have a potentially infinite stock of names to tag elements for retrieval, and since we generally avoid using the same name for different individuals, we attain the means to retrieve them over large stretches of discourse or even across discourses. As Kripke (1972) suggested, in the right context a name introduction can initiate a tradition that speakers in subsequent contexts may (usually tacitly) invoke to retrieve the referent assigned to the name with a demonstration in the earlier context.

I will not fully address the question of *how* names retrieve drefs across contexts. To do so would require a commitment to a particular theory of reference or metalinguistic view of what determines that a referential use of a name refers to the individual it refers to that I wish to avoid. I have argued (Chapter 4) that there is ample *semantic* evidence for the distinction between descriptive and referential – or *anaphoric* and *deictic* – name uses, independently of questions involving epistemology, metaphysics, or concepts. Nor do I wish to commit myself to any way of drawing the distinction between semantic reference and speaker reference. The relationship between referential expressions, or referential uses, and semantic structure applies in any case, as does the distinction between deictic and anaphoric uses of names

But if there are two different types of name occurrences, if no referent is supplied in the case of a descriptive name, and if the context of introduction of the name affects the interpretation, how do we distinguish between anaphoric and deictic names in contexts outside of a context of introduction? Indeed, how can a difference in semantic mechanisms invoked in one context affect the semantics of occurrences of names used in

a different context? Relating certain names to anaphoric pronouns raises questions not raised by anaphoric pronouns, which never escape the original context. And even if these worries can be satisfactorily addressed, why are speakers themselves generally unable to recognize the distinction between anaphoric and deictic occurrences?

In this chapter we address these questions. Although descriptive names are much more common than usually appreciated, they rarely escape their contexts of introduction. There are principled reasons for this, which I discuss in 6.2. In 6.3 I discuss what the difference, and the difference between names and ordinary pronouns, amounts to in actual use and why it is little surprise that it may be difficult for speakers to determine whether a given name occurrence is anaphoric or deictic. In 6.4 I discuss anaphoric names that escape the context of introduction and how such names can come to have referential occurrences. First, however, I provide a summary of what we have established thus far and the significance of semantic pluralism. In the following it is crucial to distinguish the *occurrence* of an expression from the expression itself. Name-using traditions are associated with *expressions*; *meanings* or interpretations are assigned to *occurrences* of these expressions drawing, in part, on the associated tradition.

6.1. Unified Pluralism

In Chapter 4 we saw that treating all proper names as referential – constants whose values are determined as a step in determining the truth conditions of a sentence – make the wrong predictions in certain situations, in particular those involving descriptively introduced names. Pronouns occur both as demonstratives and as anaphora. In the former

case they are directly referential and their values (in a context) are determined prior to truth evaluation; that is, the referent is assigned as a step in determining the truth-*conditions* of the sentence containing the name. Since their values do not change with circumstances of evaluation, they are best treated as constants. In anaphoric constructions, the proffered content of a pronoun occurrence is that of a variable. Such pronoun occurrences are singular terms and have no descriptive content, yet the individual that determines the truth-*value* of the sentence (relative to a circumstance of evaluation) is no part of the truth condition of a sentence in which it occurs. My suggestion is that we find the same “pluralism” in the case of names.

6.1.1 The Unified Theory

One obvious worry with advocating a pluralist view is, well, that it is pluralist. When names do not always receive the same interpretation, how do we determine how to interpret a given occurrence of a name (or pronoun)? Moreover, to let the interpretation of a name be determined by context – not only in terms of which individual to assign but in terms of what *kind* of value to assign – seems to lend the view a tint of gerrymandering.

The view we are defending, however, is a unified view of singular terms, and as unified as a theory of singular terms could reasonably be. All occurrences of singular terms belong to a single semantic category, have no descriptive content, and contribute individuals (only) to determining the truth of sentences in which they occur. It looks less than unified only if one mistakes the level at which a theory of names need to establish a unified set of criteria, for instance by requiring unity in terms of contributions to

propositions by requiring that sentences containing names express singular propositions. That would be an untenable criterion for singular terms in general, however – anaphoric pronouns are not referential. It is hard to see the principled advantage of having a unified theory of names at the expense of a unified theory of singular terms, compared to a unified theory of singular terms that implies a pluralist view of name interpretation, when the latter is able to provide a plausible account of how any particular term occurrence should be interpreted, and insofar as we need the distinction between singular terms interpreted as variables and as constants in any case.

It is plausible that a uniformity requirement associated with grammar or other meta-linguistic requirements regarding names (or singular terms) takes precedence over the unity of propositions expressed. It is grammar and syntax that are the first clues to interpretation, and any notion of compositionality requires that we can figure out *how* a singular term is supposed to be interpreted from this constraint. At this level we do have a unified view through the requirement that sentences containing singular terms take objects – and only objects – as values for those singular terms, and this *requirement* is a universal principle. The existence of anaphoric pronouns already shows that there is a certain variety of ways by which names can be assigned to a sentence. The claim that names also have anaphoric occurrences is thus a novel move, but not a threat to uniformity.

Grammatical categories are defined functionally. A singular term belongs to Montague's category *e*, but there are different ways to derive the truth conditions of sentences from this function, depending on whether a particular individual is assigned (constants), or whether the truth-conditions are insensitive to which individual is assigned

(variables). The requirement that the value must be an individual is sufficient to derive the rigidity of terms from compositionality, regardless of whether the propositions expressed are singular or not. At the very least, unification in terms of semantic function should ensure that if we know the language and the grammatical categories of the constituent expressions of a sentence, we should be able to at least provide a *schema* for how a sentence must be interpreted even in a situation in which we do not actually know the meaning of the individual expressions. Our pluralist theory is unified in this sense.

6.1.2. Names and Descriptions

Interpreting anaphoric names requires recovering an antecedent description appropriately related to the description used to introduce the name (we make this idea more precise below). Note, however, that interpretations of names in general rely on descriptions. Even in the case of names that are interpreted as constants, or (proxies for) demonstratives, descriptions are usually required to introduce names to a discourse (e.g. “John Smith, our new co-worker” or “Zachary Taylor, the twelfth president of the United States”). Only when all participants are familiar with the individual can a name be felicitously introduced without an accompanying description or presupposition projection. Regardless of whether an occurrence is deictic and anaphoric, the description is no part of the content but a pragmatic means for coordinating anaphoric links or referential chains. The description is generally only contingently satisfied by the referent or by any particular value for the bound anaphoric pronoun.

The difference between anaphoric and deictic names depends on the role played by this (explicit or presupposed) description, and names inherit their semantic status from

the mode of introduction into a discourse. In most cases the description, together with the character of the name and various contextual factors, retrieve a referent by individuating a name-using tradition extending back to an original demonstration (or baptism), and provides a means for assigning a referent to the name. Since previous uses of the names were demonstrative and directly referential, any (successful) attempt to use the name in the same manner as previous users (which is implicit in the meta-linguistic description “the person others refer to by *n*” or even “thing called *n*”) results in an occurrence co-referential with the original demonstrative.

In the case of a descriptive name there is no referent to inherit from previous use. Since no demonstration of a referent has taken place, the linguistic element from which the name inherits semantic status cannot be a demonstrative. Now, whether a name is anaphoric or deictic depends on factors that may be external to discourse participants; if the occurrence is deictic, it is because there is a successful demonstrative from which the name can inherit its semantic content, and a successful demonstration has taken place (and what it demonstrates) may be partially inaccessible to the discourse participants.

If there must be a successful demonstration from which a name inherits its meaning in order for it to be deictic, it follows that empty names are anaphoric names if meaningful. An unsuccessful demonstrative, for instance a demonstrative use of “he” when there is no salient male, yields a sentence that is literally contentless. Importantly, however, if “he” is a failed demonstrative and someone (Joe) introduces the name “Jack” by “he is Jack,” there are two options. Either subsequent sentences containing “Jack” are literally meaningless, or alternatively, we can engage in *repair* and presupposition accommodation – instead of treating “Jack” as meaningless, we can interpret “Jack” by

using “whoever Joe was trying to talk about;” i.e., we can project this description as a presupposition and interpret “Jack” as anaphoric on the presupposition. Of course, for Joe’s audience the latter is the most attractive option at least when the audience does not know whether the demonstration failed.

Let us consider the case where the audience does not know, and hence engages in repair. Since Joe was talking about nothing, most subsequent sentences containing “Jack” will be false (insofar as the name is *used*), but may still in principle survive and even give rise to something resembling a referentially ungrounded name-using tradition. “Whoever Joe was talking about” may serve as an introducing description in subsequent contexts and hence as conveying a shared presupposition. Since none of the participants in the subsequent context knows who Jack is (how could they?) the description must be interpreted as a device to restrict possible satisfiers for subsequent sentences “Jack is *F*” rather than as character or part of a character (alternatively, in the case where later interlocutors knows that the introduction of “Jack” involved a failed demonstrative, subsequent discourse may be understood as being semantically flawed and the following sentences as requiring repair; see also Chapter 7). Most empty names are nevertheless probably introduced explicitly by descriptions.

We develop these issues in more detail in the present chapter. It is worth emphasizing names are introduced into a discourse by descriptions, and the difference between anaphoric and deictic names is grounded in how descriptions can play different roles: as metalinguistic facilitators for demonstration (a referentially used description or complex demonstrative) – or reviving the referent of some previous demonstration – or as projecting antecedents.

6.1.3. Names Vs. Pronouns

Another important matter for the present chapter is the differences and similarities between names and ordinary pronouns such as “he,” “she” or “it.” Even if names are categorized as special-duty pronoun, the fact that names can establish name-using traditions suggests significant differences between proper names and ordinary pronouns, which do not. The main reason is that “he” is much more flexible – *any* salient (male) individual can be the value of “he.” Importantly, the character of “he” ensures that it always picks out the most salient dref. Names, on the other hand, are specially tailored to *avoid* picking up the most salient dref in favor of a different, less salient one. If Jack is the most salient male, the correct term to use is “he;” “Jack” is felicitous only when Jack is not the most salient dref. This is, of course, the reason why names are so useful. But this difference is also the grounds for other differences. The difference in activation conditions (the fact that “he” picks out the most salient dref and “Jack” does not, and that names can, accordingly, pick up drefs much farther back in a discourse) suggests that names favor a referential interpretation if one is available. For the same reason, names do not naturally occur in explicitly anaphoric position. Our example from the Chapter 4, “if a couple called their child ‘Bambi,’ Disney would sue Bambi’s parents” is a rare exception, and even in this case “his” may be more natural.

Though I do not discuss the relationship between semantics and *concepts* in detail, it is worth noting that names likely give raise to different conceptual representations than pronouns. To employ the “file”-metaphor, understanding names requires drawing on the file an agent has associated with the name, and – conversely –

interpreting sentences involving names give rise to conceptual files labeled with the name. It may be argued that ordinary pronouns do not seem to generally give rise to anything but briefly temporary files associated with a particular dref. Now, even if this observation is correct it does not undermine the close association between names and pronouns I defend. There may be a difference in robustness of files associated referents compared to anaphoric drefs, but the *type* of file associated with the expressions would be the same. The role of a name or pronoun is to track individuals; even when the term is introduced by a description the file must be labeled as a file on an individual, not the role expressed by the introduction. The interpretation of a later appearance of a dref associated with that file has to be understood as such; the antecedent description is not the semantic or conceptual content of the expression but a constraint on the correct individuation.

In 6.2 I discuss some properties of anaphoric names and how they differ from deictic names, and in 6.3 what the difference between proper names and “ordinary” pronouns actually amounts to. In 6.4 I discuss how to understand the relatively rare phenomenon of a sustained name-using tradition associated with a descriptive name.

6.2 The Ubiquity of Descriptive Names

In this section I first discuss an objection from Jeshion to the effect that descriptive names do not exist, then point out that descriptive names are not only legitimate names, but common, and compare the roles of names to the roles of pronouns in discourse.

6.2.1 Descriptive Names as Speech Act Failures

Concerns about the stipulative nature of the introduction of descriptive names like “Julius” have led some to the quixotic position that names cannot be introduced by description at all.⁹⁰ Jeshion (2004) argues that introducing a name into the language is a particular speech act that requires that certain felicity conditions are satisfied. Any name introduction must i) be performed with the right intention (the right illocutionary act); ii) the subject performing the act must be in an appropriate position to do so; iii) the object to be named should – as far as the agent introducing the name knows – not already have a name, or the agent should be, for some principled reason, unable to find out what the name is (other potential conditions are discussed in Jeshion, 2009). Accordingly, Evans failed to introduce the name “Julius” since he had no intention of using “Julius” to actually speak about that inventor.

While Evans’s famous example may initially seem to combat the force of the semantic argument, it does so only because of its artificiality. No name has been introduced into the language [footnote omitted].

One possible move here would be to accept these points [i.e. the conditions] as applicable to ostensive names, yet deny them for descriptive names [...], maintaining that descriptive names are a special variety of names that are simply not subject to the same conditions as ordinary names. But this rejoinder sounds somewhat desperate at this stage. It is like saying that in addition to promises, there is a special class of promises, promises* – a class known only to philosophers – and for which the promiser* need not have any intention of fulfilling the act she promises* to do. Few will take this suggestion about promises seriously. I think we should take the same negative attitude about the free introduction of descriptive names [footnote omitted]. [Jeshion, 2004, p.602]

But this is puzzling. The analogy to promises suggests not only that “Julius” fails to be a real name but that “Julius” is rendered *meaningless* since the speech-act misfires; i.e. that

⁹⁰ Of course, one may argue that by definition whatever a stipulation of “Julius” introduces it is strictly speaking not a *name*, but something looking like and sharing relevant semantic properties with names; this would be a terminological point, however, insofar as the resulting terms behave, syntactically and semantically, like names.

Evans failed to accomplish anything at all. This would at least require an argument insofar as it is definitely possible to understand what discourse participants are saying when using “Julius.” Perhaps the idea might be that since “Julius” is not a name, “Julius invented the zip” does not express a proposition at all, and occurrences of “Julius” must be analyzed pragmatically. This position might appear defensible if we did not systematically utilize particular linguistic resources to introduce such names, or did not introduce them for very specific purposes. But we do. We use descriptive names in situations where we exactly want to avoid referential commitments, and where no other expression, such as a regular pronoun, would do the trick.

A similar move is suggested by Kaplan to avoid the result that the introduction of descriptive names generates singular thoughts *en masse* (and hence in order to avoid making the strange *a priori* contingent propositions such as *Julius invented the zip* too common):

Normally one would not introduce a proper name or a dthat-term to correspond to each definite description one uses. But we have the means to do so if we wish. Should we do so, we are enabled to apprehend singular propositions concerning remote individuals (those formerly known only by description). Recognizing this, we refrain. What purpose [...] is served by direct reference to whosoever may be the next president of Brazil? The introduction of a new proper name by means of a dubbing in terms of description and the active contemplation of characters involving dthat-terms [...] constitute a form of cognitive restructuring; they broaden our range of thought. To take such a step is an action normally not performed at all, and rarely if ever done capriciously. The fact that we have the means – without special experience, knowledge or whatever – to refer directly to the myriad individuals we can describe does not imply that we will do so. [Kaplan 1989, p.560 ft. 76]

Although Kaplan may be right that we generally do not introduce names this way, the relevant problem is that we *can* do it. We should not be able to generate singular thoughts

at the stroke of a pen or get contingently *a priori* propositions of this kind, but that is of course not an argument against the existence of descriptive names.

Of course, Kaplan's *worry* is well taken, for if singular thoughts could be created by the stroke of a pen – just by using the *dthat*-operator – we may come close to a *reductio* of Kaplan's *dthat*-interpretation of demonstratives. This *reductio* can be avoided, however, though at a cost. If, as Kaplan sometimes suggest, the *dthat*-operator turns an expression *dthat*(the *F*) into an *indexical* by incorporating contextual information into the character of the expression, then the possibility of creating singular thoughts by the stroke of a pen is devastating and unavoidable. But if *dthat*(the *F*) is a *demonstrative*, there is no threat. In that case, the context and a salient demonstratum are needed to individuate an object that can serve as the referent of the expression, and in the absence of a demonstratum, *dthat*(the *F*) is undefined and hence not the source of any singular thought at all. But that still leaves unexplained the fact that we do introduce names descriptively. Since “Julius invented the zip” is not meaningless, “Julius” cannot be a *dthat*-construction. But there is little reason to think that it should be; we did not try and fail to employ a demonstrative to introduce “Julius.”

6.2.2 The Utility of Descriptive Names

Situations in which it is useful to introduce a name to stand for an unknown or even arbitrary satisfier of a description arise for instance when we need to assess complex states of affairs to which we have no demonstrative access, or keep track of movements of various objects or individuals satisfying certain roles without committing ourselves to any particular individuals fulfilling those roles. An example could be crime investigators

attempting to distinguish the doings of a large number of unknown individuals at a crime scene. One could of course keep track of them by using definite descriptions (“the person who entered first;” “the person who entered through the back,” and so on), but it is often easier to keep track of them using dummy names – especially since investigators may want to be able to make counterfactual claims about these people and thus have to have a rigid expression at hand.

Claims about future states of affairs often involve introducing names for which the actual truth-maker is irrelevant to the truth conditions. For instance, meteorologists already have a pre-determined stock of names to use for storms and hurricanes, and may use these names in predictions about the future, even though the satisfiers are individuated only by description. Similarly, if I were to make a work drawing of a house I intend to build, I may give names to certain elements of the construction (e.g. single letter names “a,” “b,” “c,” and so on). Using names makes it easier to consider counterfactual scenarios – and I introduce the names precisely because i) deploying tags makes it easier to track the elements (*de re*) in counterfactual scenarios compared to definite descriptions; and ii) deploying tags enable me to keep track of more elements than using pronouns would do. Given that the house is not already built – and may perhaps never be built – “a,” “b” and “c” are not referential terms; but they are perfectly good names (I could have used “Jack,” “George,” and “Alexius” instead). In different worlds or different futures different individuals serve as designata of these names, but the names are nevertheless rigid. Keep in mind that the names cannot just refer to the *drawing*. If “b” designates a wall, the sentence “b will be able to carry the weight of the roof” may be true, which it is not if “b” merely referred to the drawing.

The names in question are context-bound; outside of the context of introduction I cannot use the names “a,” “b,” and “c” and expect my audience to know what I refer to, even if that audience was present at the original introduction. Now, most of the cases where descriptive names are used are characterized by the fact that the names do not acquire a life beyond that context (“Jack the Ripper” and “Julius” are rare exceptions). It is therefore plausible that Jeshion’s conditions are necessary for establishing a name-using *tradition*. Few cases of name introduction, even with regard to deictic names, satisfy Jeshion’s conditions, so few names survive beyond the context in which they were originally introduced. Jeshion, however, overlooks the fact that what goes on inside these contexts still requires an analysis.

6.3 Names As Pronouns

Names may be categorized with pronouns. Both names and pronouns

- i) are definite expressions and associated with familiarity constraints,
- ii) are singular terms taking individuals as values (that is, they do not in general track roles or properties, but the individuals who fill them),
- iii) have no descriptive content,
- iv) are (as opposed to general terms) associated with a character/content distinction,
- v) sustain a distinction between accommodation and binding in discourse.
- vi) are rigid.

vii) are complete noun phrases with no semantic distribution among syntactic parts.

In practice, however, proper names and pronouns also exhibit significant differences, partially as a result of the “specialized nature” of names:

i) Given that there are rarely more than one potential dref for “Alexius” in a context or discourse, a name is able to activate drefs from far back in a discourse, whereas pronouns pick up the dref that is at all times most salient (at least relative to the reach associated with their very limited lexical meanings).

ii) Pronouns such as “she” relate to antecedents partially in virtue of their linguistic meaning or character (“the most salient female”). Since drefs are generally more salient than items in an extra-linguistic context that has not yet been mentioned, the characters of pronouns suggest a preference for binding over accommodation. Names do not suggest this preference, since its character does not similarly motivate a search for an antecedent binder through a salience requirement embedded in the character.

iii) Pronouns automatically create a link to its antecedent; names must be explicitly introduced through a baptism or stipulation. The name-using traditions of names are presumably something of a side-effect.

iv) Name occurrences can project presuppositions; ordinary pronouns cannot.

In this section I discuss these similarities and differences, and their consequences. The main hypothesis is that the differences i)–iv) between names and ordinary pronouns are explained by the fact that ordinary pronouns are designed to pick up the most salient dref, antecedent or extra-linguistic referent, whereas a name is not. This difference also accounts for the usefulness of names, since it means that names can be used to circumscribe activated drefs and keep track of several members of a predicate extension simultaneously. The difference in reach is subsequently explained by the character of the pronoun (“the most salient male”) and the name (“the person called *n*” or perhaps even “the most salient person called *n*”).

6.3.1 Names and Pronouns: Tracking Individuals

Deictic name is at the time of introduction co-referential with a demonstrative pronoun. Since both are non-descriptive, an occurrence of a demonstrative and a name introduced to refer to the object it picks out, are semantically equivalent, and the two expressions are related by co-reference (rather than binding). Introducing a name to be co-referential with a demonstrative adds expressive power. A name is cognitively “lighter” than a demonstrative, and identifying objects by repeated demonstration is more cognitively demanding than labeling them with tags and picking them up by those labels. The existence of a large stock of different names (the expression), and our ability to introduce new ones when needed, facilitate our ability to retrieve referents (by referring) even when they can no longer be accessed by demonstratives since they enable us to sustain co-reference beyond the context in which demonstration is possible. “The person called *n*” will generally have far fewer possible satisfiers in a context than “the most salient male.”

The name is not as cognitively light as a pronoun of laziness, but the use of lazy pronouns, like anaphoric pronouns, is limited to picking up the most salient individual. Thus, a name is often required to keep track of different individuals when several are in play simultaneously to avoid ambiguity (compare the ambiguous “George and Jack fought. He won,” and the non-ambiguous “George and Jack fought. George won”).

A descriptive or anaphoric name serves a similar but nonetheless semantically importantly different purpose. A descriptive name is introduced to go proxy for a pronoun anaphoric on a quantified or existential antecedent, and is semantically equivalent to this pronoun (*not* to the antecedent quantifier). A name is cognitively “heavier” than an anaphoric pronoun, but lighter than a definite description, so to keep track of several descriptively introduced drefs names are helpful and less cumbersome than descriptions, in particular if we need to distinguish several individuals in a given predicate extension, for instance when I need to make claims concerning, say, the relationship between several tables in a room without knowing anything else about them. Names also have the advantage of not exhibiting scope ambiguities, and they are therefore rigid and better suited for describing counterfactual scenarios unambiguously. Anaphoric pronouns are even lighter, but if there are several individuals in play pronouns such as “it” or “he” may give rise to ambiguity.

Deictic names are useful in ways that anaphoric ones are not. Names (the expression) are specialized tags, and in virtue of being specialized they are able to survive the context of introduction and be used to pick out individuals that are no longer demonstratively available. Special facilitating conditions must presumably be in place for a name to gain traction in a linguistic community, and many deictic names are introduced

only for temporary purposes, such as nicknames or aliases. Given proper circumstances, however, a deictic name may initiate an extensive name-using tradition. Anaphoric names are less useful in this respect. Such names have a reach in a discourse that is, in principle, as long as deictic names, but rarely survive the context of introduction. The reason is that the description that serves as the name's antecedent is *also* able to track drefs across contexts. In the case of deictic names, the demonstrative used to introduce it is of course *not* able to track drefs into a context in which the demonstrated object is no longer available. Thus, when a descriptive name initiates a name-using tradition we get over-determination. For an anaphoric name to be preserved across contexts, the antecedent description must (relative to conditions to be described) be preserved as well since the description must be invoked to provide truth conditions for sentences containing the name. Since there is little reason to preserve the name when the associated antecedent description alone, which must be retained in any case, would serve the same purpose, anaphoric names predictably survive the context of introduction only rarely.

Within a given context, however, anaphoric names are useful for their ability to pick up drefs from far back in a discourse. Descriptions cannot. Descriptions pick up properties or roles; they denote whoever fill these roles, not the individuals, and contribute the roles to the truth conditions of sentences. Thus, definite descriptions play different roles in natural language than names, even anaphoric names, and anaphoric name occurrences cannot in general be substituted by descriptions.

6.3.2 Names and Pronouns: Binding, Accommodation and Presupposition Projection

Names have longer reach than pronouns. There is a restrictive limit on how much of discourse can have elapsed between a felicitous use of an anaphoric pronoun and its antecedent, but few and liberal restrictions on the felicitous use of proper names. Indeed, proper names are often taken to introduce or re-introduce referents *across* contexts, something a pronoun cannot since the free occurrence of a pronoun (without a demonstration) severely underdetermine its possible antecedents. The meaning of an occurrence of a name “John” is strictly speaking not entirely stable either, but although the name “John von Neumann” is context sensitive and may, in principle, designate different individuals on different occasions the range of possible interpretations is very limited (often to a single one) – which is, of course, the feature that allows us to maintain stable name-using traditions associated with names (the expression). Since names do not pick up the most salient referent or antecedent, however, they will often have to be explicitly introduced for the audience to be able to locate the correct referent or antecedent. These features are shared by all names, and serve to underpin a distinction between names and ordinary pronouns with respect to activation properties. This distinction is not semantic (or at least not truth-conditional⁹¹).

Another pragmatic difference can be illustrated by the occurrence of pronouns in “she will abuse it.” The pronoun “she” can in this case in principle function either as bound or as a demonstrative. Even the sentence “if the queen has unlimited power, she will abuse it” is strictly speaking ambiguous, depending on whether “she” is understood as anaphoric on “the queen” or interpreted as a demonstrative picking out some other

⁹¹ It may apply to context-update potentials, however, which is a more fine-grained semantic criterion.

salient female in the context. An audience hearing the “open” sentence “she will abuse it” will have the option of looking for a salient female who fits the predicate, or to interpret the occurrence as being anaphoric on an antecedent introducing some female or other to the common ground. If there is no demonstratively available female the anaphoric interpretation is mandatory, but even if there is, the audience will usually (barring overriding pragmatic factors) be hesitant to assign the demonstratively available female to the pronoun. Something more must be provided to encourage the demonstrative interpretation and the accommodation of a new referent. The anaphoric interpretation will hence generally trump accommodation of a referent.

Names, on the other hand, do not exhibit a preference for binding. The activation constraints associated with names enable names, as opposed to ordinary pronouns, to retrieve drefs other than the most salient one. The fact that they are *able* to do so might make speakers even expect that the purpose of a given occurrence is to circumvent the most salient dref to retrieve a different one (otherwise a less cognitively demanding pronoun is preferable). This is particularly strikingly exhibited by the fact that names, as opposed to pronouns, do not allow *de se* readings (Chierchia, 1989). Compare “Jack thinks he is a fine fellow” and “Jack thinks Jack is a fine fellow.” There is a general (pragmatic) push toward a “lighter” expression, one tailored to pick up the most salient dref, when possible, and using a “heavier” expression, such as a name (or even more markedly, a definite description or a demonstrative) only when we wish to emphasize a lack of connection between co-referential expression. Thus “Jack thinks that Jack is a fine fellow,” suggests that we are talking about two people called “Jack” (or that there is something unusual about Jack’s self-locating beliefs, as in Perry, 1979). The issues

surrounding *de se* readings are complex, but the availability of *de se* readings when using pronouns and not names, is not obviously an objection to categorizing names with pronouns. That names cannot have *de se* readings are explained by pragmatic factors pertaining to activation potential. Names are associated with the *absence* of an already activated dref. The fact that names do not pick out activated drefs, but accommodate new (or deactivated) drefs, predicts that names will not sustain *de se* readings (as well as why names must be explicitly introduced). The real question is of course not why names lack *de se* readings, but why ordinary pronouns have them (though only in subordinate position). This is a difficult one, but it is worth noticing that *de se* readings are generally available only when the pronouns are *syntactically* bound.

The most obvious explanation for the difference in activation conditions between names and ordinary pronouns is linguistic meaning. A pronoun “he” is, regardless of whether it is a demonstrative or occurs anaphorically, associated with a character along the lines of “most salient male,” and the pronoun will thus be connected to an appropriate dref (if any is available) partially in virtue of its linguistic meaning. The linguistic meaning of a name (“person called *n*”) does not immediately connect it to an antecedent in this manner. This is an important reason for the special, long-distance reactivation potential associated with names. Thus, if a name occurrence is to be bound by an antecedent description it has to be explicitly associated with this description (for instance by a stipulation) – its linguistic meaning does not suggest a one. There is hence little wonder that there are few examples of syntactically bound names.

The difference in character also explains another important difference between names and ordinary pronouns. Even in the case of descriptive names the character is *not*

the binder, but helps locate the correct binder. In a discourse the binder is usually the description used to introduce the name (as in “Vulcan, the planet between Mercury and the Sun”). But the character of a name can also be used to *project* a descriptive presupposition – in the most general case “the thing called *n*” – in cases where participants are familiar with the name and share the presuppositions, which can then be used to bind subsequent occurrences of the name (i.e. by restricting the domain it ranges over). The character of a pronoun cannot project presuppositions – “salient male,” for instance, is not an apposite presupposition.⁹²

6.4. Anaphoric Names Across Contexts

Treating “Julius” as a term the interpretation of which is determined by a previously occurring quantifier phrase requires that this phrase plays a role in the semantic analysis – a sentence “Julius is *F*” requires a context to receive a determinate truth value. If no referent is available (even through deference to a name-using tradition) we need an explicit antecedent background. In DPL the relation is captured by semantic binding represented as operations on the assignment function; in the semantics suggested in chapter 5 the antecedent issues a restriction on the domain of possible satisfiers. In 6.4.1 I describe some intuitive tests for whether a name is anaphoric or deictic and why the

⁹² The issue of presupposition projection is a thorny one, but the *possibility* of presupposition projection is accepted by virtually anyone. Stalnaker (1978) took presupposition to be a pragmatic mechanism based on a notion of “local context” (see also Heim (1983); for a recent defense of a Stalnaker-inspired account, see Schlenker, 2010). An alternative that might be better suited to the purpose of the resolution of anaphoric names would be a version of a Binding Theory of Presuppositions (e.g. van der Sandt, 1992; Geurts, 1996) or the intermediate account of Moltmann (2006a). I do not attempt to develop a comprehensive account here since it would require taking a stance on a range of questions that go relatively far beyond the scope of this dissertation.

difference does not entail any clear differences in semantic competence (or even knowledge) for the user. In 6.4.2 I discuss the potential context-escaping property of names and in 6.4.3 how to understand this phenomenon. In 6.4.4 I briefly relate my view to different theories of reference, and in 6.4.5 I discuss cases where a descriptive name evolves into a deictic name.

6.4.1 Tests for (Semantic) Binding?

Although anaphoric names resist Kripke's semantic argument the often transient nature of such names makes it difficult to apply this argument as a test for whether a name occurrence is deictic or anaphoric. The upshot of the semantic argument is that speakers with incomplete or even faulty information may nevertheless succeed in referring. Since anaphoric names generally occur only within the context of the discourse in which they are introduced, speakers will rarely possess incomplete or faulty information, and it is generally not possible to wonder whether the designatum satisfies the introducing description. This may arguably also be the case also in contexts of introduction for deictic names, however. A possible case that would illustrate the difference could be the following. Suppose the parents of a child name their son "Jack" by a demonstrative baptism. In that case it is of course possible for them to wonder whether Jack is, in fact, their son. If they name their son by stipulation "we name our son Jack," it is less clear whether this worry would make sense. Instead, it would be possible to worry that the child they are seeing is, in fact, Jack. I am wary of putting too much weight on such scenarios, however; the intuitions at work seem to be somewhat complex.

There may still be tests for whether a name is anaphoric or deictic. Compare:

(1) Julius is the inventor of the zip

(2) Judson is the inventor of the zip

Consider the question *does it matter to determining the precise circumstances under which (1) and (2) are true, that you can determine who invented the zip?* In the case of (1) the answer is “no.” In the case of (2) the answer is “yes.” Furthermore, suppose investigation led us to believe that McTaggart instead of Judson invented the zip. In that case it would intuitively be reasonable to give up (2) but not (1).

Second, (3) seems plausible, whereas (4) sounds nonsensical:

(3) If we discovered that McTaggart invented the zip, McTaggart would be Julius.

(4) If we discovered that McTaggart invented the zip, McTaggart would be Judson.

Both tests rely on epistemic distinctions in addition to semantic ones. The third test does not, but has other limitations. Let “the F ” be the description associated with a name n in a discourse, and consider:

(5) Could n be someone other than the F ?

It could not be the case that Julius is someone other than the inventor of the zip; that is, that we could be mistaken about the identity of Julius (not whether, in a counterfactual

situation, Julius did not invent of the zip – presumably Julius only contingently invented the zip). Our discourse is, to an extent, *immune to reinterpretation*. Suppose there is a stretch of discourse concerning Judson’s inventorship of the zip. If someone pointed out that McTaggart invented the zip, we could reinterpret previous discourse to be about McTaggart, not Judson, and in the process presumably admit that Judson is someone else and that the claims we made about him are strictly speaking false. A similar discourse about Julius, however, does not require reinterpretation; rather, we would interpret it as having been about McTaggart all along.

If this is correct, it suggests that the designatum of “Julius” is determined by the truth conditions for sentences involving the name only relative to a circumstance of evaluation, and is itself no part of the truth conditions or proposition, whereas in the case of “Judson” the designatum is *presupposed* for deriving these truth-conditions and hence part of the truth-conditional content. Suppose we, by considering just what it would take to invent the zip, infer that Julius must have been clever. Suppose there is a subsequent discussion about whether Judson or McTaggart invented it. Our judgment about the truth of “Julius must have been clever” would not depend on the outcome of this discussion. Since the truth of “Julius must have been clever” can be agreed upon independently of the result of the discussion the semantic contribution of “Julius” cannot depend on this outcome. Suppose that we, just by considering what it would take to invent the zip, infer that Judson must have been clever. If we subsequently came to question whether Judson or McTaggart invented the zip, we would also suspend judgment about the truth of the conclusion (that Judson is clever); at least the conclusion would no longer be licensed. Thus, the contribution of “Judson” to the discourse is sensitive to how things are with

Judson. Not only does the truth-value of “Judson must have been clever” depend on the value assigned to “Judson” – the truth-*conditions* depend on the assignment as well.

6.4.2 Escaping the Context of Introduction

In the case of anaphoric names that do not escape their context of introduction the cognitive resources required to understand them are in practice not much different from the resources brought to bear on deictic names, even though they are semantically treated as anaphoric pronouns. On Millian views, “Julius” is directly referential, yet given that Julius is stipulated to be whoever invented the zip there is no real possibility that “Julius” would be associated with incomplete information or turn out to refer to anyone who is not the inventor of the zip even if a Millian interpretation were assumed (though that observation is a *problem* for the Millian, since it gives rise to unacceptable *a priori* contingent truths). Some anaphoric names escape the context of introduction, however. In these cases there is a distinct possibility that some of the associated information should be lost between contexts of use. Since such names are supposed to be bound by their introducing description it might appear that understanding their truth-conditional contribution requires knowing the introducing description. This would be a worry, since it would suggest that anaphoric and deictic names draw on significantly different interpretive resources. As it turns out the worry is a serious one for alternative views of names, but disappears once we adopt the anaphoric view.

Fregean views encounter a similar worry about names if senses are understood as descriptions. According to Fregean views a name is associated with a sense, and in order to understand an occurrence of a name one may need to grasp the sense. It is, however,

implausible that speakers are explicitly aware of what description earlier users associated with the name, so on this view the meaning of a descriptive name (the expression) will change over time. It is sometimes argued (e.g. Sainsbury, 2002, 2005) that in order to understand an occurrence of *n* it is sufficient to associate with *n* only *some* sense or other that determines the correct reference, and that sameness of *reference* is sufficient to ground a name-using link even if the descriptive sense is not transmitted. The problem, however, does not quite go away. One of the putative strengths of a Fregean view is the promise of a solution to the problem of *empty* names, such as “Vulcan,” which is supposed to denote the planet located between Mercury and the Sun. Sainsbury’s criterion is insufficient for grounding a name-using tradition for an empty name since there is no referent for different senses to determine. Thus, even Sainsbury’s Fregean account must assume that descriptions associated with a name are retained in a linguistic community in order to have a notion of sameness of sense that is robust enough to ground a name-using tradition across contexts.

On Millian accounts referring to Aristotle by “Aristotle” requires merely that one is latches onto the correct name-using tradition, which again supplies a reference in a context, and some (perhaps causal) connection to Aristotle himself serves to keep the semantic contribution stable, regardless of descriptive information associated with him. In the case of “Julius” the name-using tradition does not extend to Judson, but to an original descriptive stipulation. There may be some worries concerning how we guarantee that uses of “Julius” in later contexts would still refer to Judson – there may be some indeterminacy with regard to whether such uses would be making false claims

about Judson or true claims about a different denotation⁹³ – but it is empty names that make the worry most obvious even on a Millian view.

For instance, to interpret occurrences “Vulcan” one must associate some description with the name. Even if “Vulcan” referred to an abstract or Meinongian, non-existent object, there would have to be some (meta-linguistic) way of ensuring that “Vulcan” is interpreted as referring to Vulcan rather than to Santa; that is, that we are using the name with the same interpretation as before and – crucially – with a different interpretation than the one we assign to “Santa.” For deictic names causal mechanisms can ensure this. For empty names there is no obvious mechanism. At some level of explanation the difference between different empty names is a difference in associated descriptions. Even if the contribution to the proposition is merely a gap (Taylor, 2000; Braun, 1993, 2003), in which case the content of “Vulcan” and “Santa” is the same, one must still admit that the ability to *convey* information with the name requires associating some description with it. So again, preserving meaningfulness (even if only on a level of pragmatics) across contexts requires associating the name with some description. Thus,

⁹³ One issue that has engendered much confusion, and that is worth mentioning here, is the possibility of reference change. Evans (1973) pointed out that “Madagascar” originally referred to a portion of the mainland of Africa, but that the reference of the name gradually changed to the island. Strictly speaking there is no room for reference change on direct reference views, however. Now, how to individuate *names* is controversial, but the two primary options are i) to take names to have a single reference each; if several people are called “Eric,” for instance, they do in fact bear phonographically identical yet different names (Kaplan, 1990); or ii) to take names to be shared between bearers. On the first view, names are partially individuated by their bearers. In that case, however, since Madagascar the portion of mainland Africa and Madagascar the island are different referents, they do not share a single name. Rather, there are two different names “Madagascar,” one that refers to the mainland, and one that refers to the island. Thus, there never was any reference change; when “Madagascar” became the name of the island, it was really a new name that came into existence (and the old name “Madagascar” for the mainland gradually faded out of use). On the second view, the referent of a name is individuated in a particular context; that is, there is only one name “Madagascar,” but it can be used to refer (directly) to different entities on different occasions, just as the name “Eric” (directly) refers to different persons in different contexts. But in that case we do not have an example of a reference change either. The name itself never had a referent; rather, the *use* of a name in a context referred. The fact that the object that conventionally serves as the referent in a given context changes is hence not strictly speaking an example of reference change.

both Fregean and Millian views require that descriptions associated with descriptive (and at least empty) names are somehow maintained by the name-using tradition. Since this is not required in the case of ordinary names, descriptive and empty names are different from ordinary deictic names – if not semantically, then at least with respect to what it takes to ground a name-using tradition and preserve meaning across contexts.

If “Julius” is an anaphoric name the situation is simpler. What does it take for two occurrences of “Julius” to have the same meaning? One misleading answer is to take two occurrences of “Julius” to mean the same if they have the same binder. This amounts to confusing the meaning of a descriptive name with the antecedent. The introducing description and the descriptive name are not synonymous. The content of an occurrence of a descriptive name is that of a bound variable, and “Julius is F ” is interpreted as an open sentence “ x is F .” The truth-conditional contribution of “Julius” relies on the antecedent, but the contribution “Julius” makes will be exactly the same if the semantic binder is “the inventor of the zip,” “an inventor of the zip,” or “the individual Evans called ‘Julius’” (though the discourse dynamics may be different). It is the same since, given that the world (and assignment) stays the same, any sentence bound by these descriptions is satisfied by the same sequence of objects.

Indeed, to keep the interpretation of “Julius” constant over occurrences and contexts we only require that the presuppositions are *extensionally* equivalent. This is the case even if “Julius” occurs under the scope of a modal operator. “Julius could have failed to invent the zip” should be “ $\diamond \neg \text{Invent-the-zip}(\text{Julius})$.” Since this sentence is, on its most natural reading, true, the antecedent binder is interpreted as taking wide scope over the modal operator. As long as we keep the circumstance of evaluation fixed relative

to the antecedent, whoever satisfies “Julius” (for the sentence to be true) will be fixed as well, even with respect counterfactual scenarios. Empty names have the same meaning in different contexts insofar as the binder ensures that the occurrences of the name are true in the same situations. Of course, this requirement may in principle be satisfied even in the face of rather drastic changes to the content of the binder – “the planet between Mercury and the Sun” and “the city called ‘Atlantis’” are extensionally equivalent. But as opposed to the alternative views mentioned the anaphoric interpretation is not committed to taking into account an anchor or referent that these different expressions determine. Radical changes in meaning is not itself a problem for the anaphoric account; such changes will of course lead to misunderstandings, but the anaphoric approach can nevertheless provide determinate truth conditions in each context of use. (We return to empty names in the next chapter.)

Let us go through the role of antecedents across contexts in a little more detail. The interpretation of a name in a context relies on presuppositions invoking name-using traditions. When names are deictic, associated information supplies a referent if it (by some mechanism) individuates the anchor the tradition, regardless of descriptions speakers associate with the name. Now, since a name does not have a standing meaning (apart from “thing called *n*,” for an anaphoric name to survive beyond context of introduction the antecedent description must (in some sense to be made precise) be retained as well to provide the interpretation of sentences in later contexts. Anaphoric names have no semantic anchor in an extra-linguistic referent, however. On the face of it, the procedure might appear to be essentially the same as the procedure for deictic names. The character, “the thing called *n*,” which already embodies a certain deference to

previous usage and hence a referential chain, and context help individuate the anaphoric chain. Of course, in this case these factors are insufficient to assign a referent, since the anaphoric chain does not supply one. At least on views that focus on speaker reference, the notion of using the name deferentially, or “with the same interpretation,” might nevertheless seem to facilitate binding the name with the original introducing description. So, should we say, then, that in the case of a descriptive name, an occurrence is bound, across contexts, by the introducing description?

I suggest, admittedly tentatively, that the answer is “no.” In the case of demonstrative names, its referent is its meaning (if anything is), and the role of associated information is to facilitate co-reference. In the case of anaphoric names the name-using tradition and associated information does not play an entirely analogous role. The anaphoric binding relation is encoded in the dynamic layer representing the continuously updated context of a discourse, and this layer is lost once that discourse is closed. Now, all uses of names are presumably, at least to an extent, *deferential*; even the character of a name “thing called *n*” embodies a degree of deference to previous usages. To use the name “with the same interpretation” helps locate the correct name using tradition, which – if everything goes well – supplies a referent. But when using an *anaphoric name* “with the same interpretation” things are different. In the referential case the content is the referent, so invoking the same content means invoking that referent. For an anaphoric name the antecedent is *not* part of the content, so an intention to use the name “with the same interpretation” is not enough to locate the binder from previous context – the content of the name is at most a “free” identity function. To interpret it, one must supply an antecedent, but that binder is local to the context in which the name is used. (It is

hence not incorrect to say that anaphoric names *never* escape contexts of introduction.

The proffered content of an anaphoric name is that of an open variable. Instead of saying that we “retain” the variable “with the same interpretation” across contexts, occurrences of syntactically identical variables are bound by binders available in the context of use, and the relationship to previous contexts is looser than for deictic names.)

Uses of “Julius” require a presupposition. If the audience is unfamiliar with the name, these presuppositions must be made explicit. If the audience is familiar with the name, one can rely on shared presupposition, but the assumptions must match to a greater or lesser degree (extensionally equivalent presuppositions) for talk about Julius to be felicitous. The presuppositions used to bind the variable may of course be the same as those individuals in the earlier context, and deferential intentions provide resources for establishing connections across contexts; even for descriptive names one defers to previous users by intending to use the name in the manner of one’s linguistic surroundings. As mentioned, deference cannot provide the semantic content in this case. Thus, when using “the thing everyone else is talking about using *n*” as a presupposition, the anaphoric name is *bound* by *this* description, not by the description binding it in a previous context, so the binder occurs in or is presupposed locally to the discourse in which the name occurs.

Deference can facilitate (extensional) equivalence across contexts, however. “the thing everyone is talking about with *n*” will by default be extensionally equivalent to the antecedent description in previous contexts. If *n* was bound by “the *F*” in C_1 , then using “whatever the others talked about with *n*” (rather than “the *F*,” which by hypothesis is unavailable) to bind *n* in C_2 will guarantee (at least barring indexicals in the description)

that the utterance “*n* is *G*” in C_1 is true iff “*n* is *G*” in C_2 is true. To this extent sameness of interpretation is achieved across contexts even though binding is local. Sameness of interpretation in this strong sense does not always occur, however. Since the anaphoric name is a variable, not directly referential, we may be a bit flexible – not all occurrences of “Julius” *must* strictly speaking make the same contribution, any more than all occurrences of “she” has to make the same semantic contribution.

The local binding requirement has some novel consequences. Since information associated with a name may change, the contribution of a descriptive name (the expression) is in principle unstable. Since the content of an anaphoric *occurrence* is the content of a variable, talking about “retaining” the name across contexts is misleading – the *expression* is retained in a linguistic community, as is usually some associated information (if the presuppositions associated with a name have little in common across contexts, then the name occurrences have little in common as well), but the interpretation is determined in the context of use. So for instance, suppose a linguistic community began to associate “Vulcan” with “the planet between Mars and Jupiter.” This community uses the name in a different way than we do, but it is hard to say that they use it *wrongly*, as long as they are not also assuming that they use it the same way we do. On the other hand, if two discourse participants associated somewhat different information with “Santa,” we cannot rule out the possibility that they, to some extent, are talking past each other. Since there is no object of which they could make false or true ascriptions, it is difficult to see that this sort of indeterminacy is something we should worry about.

Note that one can interpret utterances made in a different context C as long as one can retrieve a binding description that is extensionally equivalent to one used in C . Thus,

interpreting a text using “Julius” written in a different context under different assumptions or in a different language, is still eminently possible. That the binder is not immutable is not a barrier to communication or at least partial understanding across contexts, though misunderstanding is a risk.

As mentioned, the most straightforward binder that would achieve cross-contextual interpretation and guarantee extensional equivalence is “the individual called ‘Julius’ by the speakers in this context.” Since knowing a description like this is also what one needs to know to apply a deictic name – though in both cases, however, the description serves as a criterion for co-reference – the cognitive requirements anaphoric names and deictic names place on the speaker are pretty much equivalent. Of course, there is a vast semantic difference. In the case of “Julius” the description must be explicitly given in a representation of the semantic content – it must project a presupposition – whereas in the case of “Aristotle” the description serves, at most, to (metalinguistically) individuate the correct tradition.

6.4.3. The Notion of Context, and Theories of Reference

We assume that contexts are objective. This is the conception of context found e.g. in Lewis (1979) or Kaplan (1989). Anaphoric relations are hence encoded in an objective common ground. An utterance of a deictic name *n* is assigned a referent in virtue of properties of the context, and context keeps track of anaphoric relations. In the case of anaphoric names used with no explicit introducing antecedent, we must invoke binding presuppositions to interpret sentences in which they occur, but these presuppositions are generally fairly obvious. After all, a name is felicitously used without an accompanying,

explaining description only when all participants can be assumed to be relatively thoroughly familiar with the name. We may hence allow a relatively modest form of presupposition projection (Heim, 1992). This presupposed binder may be unhelpful to participants unfamiliar with the name, but for such participants using the name without explicit introduction is infelicitous in any case.

Although context “stores” anaphoric relations, it is less clear how we determine whether a name occurrence is deictic or anaphoric. Evidence for the *distinction* is obtained (Chapter 4) independently of any particular theory of reference (i.e. that answers the question of in virtue of what a name occurrence comes to refer to whatever it refers to), but it is in the end difficult to give a precise answer that covers all cases without committing ourselves to a particular such theory. In the case of “Aristotle,” for instance, a causal theory (e.g. Kripke, 1971; Devitt, 1974, 2012) makes the referent available through a causal link to Aristotle. If there is such a link, “Aristotle” may be used referentially. On more epistemically oriented views, it is the existence of an informational link (Evans, 1982) or the availability of *de re* knowledge of the referent (McDowell, 1984) that ensures the availability of a referent. In neither case, however, will there plausibly be a link to Whitcomb Judson in the case of “Julius” strong enough to support *reference*. Deflationary views of reference (Horwich, 1998) has sometimes been seen as an alternative, but it is worth pointing out that although deflationism is a viable perspective on the issue of what the relation of reference *is*, a metaphysical question, the question at hand requires more than a deflationary disquotational schema. After all, for an occurrence of “Aristotle” to refer to Aristotle the philosopher, rather than the shipping magnate, some element or other of the context will need to intervene, and to determine

what the use refers to we need a mechanism that Horwich cannot obviously supply. It is worth pointing out that even accepting a Kaplan-style view of names, according to which Aristotle the shipping magnate and Aristotle the philosopher bear different though syntactically identical names, would not help avoid the question. In this case the question just gets pushed one step back to a question of what determines which *name* I am using. Citing the speaker's intentions would not obviously help. One would still need a theory of reference to ensure that a speaker's intentions are determinate in these cases.

As opposed to ordinary pronouns names exhibit a preference for accommodation of a referent rather than binding. So if context can supply a referent (by whatever mechanism), a name occurrence will generally be referential. There may, however, be exceptions. If I use "Telemann, the most productive composer in history" to make a general point about being a productive composer (rather than about Telemann himself), or "John, our newest coworker" to make a point about recent hires, it might arguably be plausible to interpret the use as a case of binding, even if context could in fact supply a referent. How to interpret such cases may, however, also depend on how we draw the distinction between semantic reference (the reference an expression "has") and speaker reference (the object the speaker intends to pick out). In terms of semantic reference, the use of "Telemann" is arguably referential, whereas relative to my intentions to speak about productive composers in general, my use of the name is plausibly interpreted as bound. In fact, it may not be entirely implausible to suspect – especially if the goal of semantic analysis is to interpret the *speaker* – that the majority of name occurrences, perhaps with the exception of what Evans calls "producer names" (Evans, 1982), are bound. I leave this option as an alternative for future development.

6.4.4 Transitions from Anaphoric to Deictic Names

A name (the expression) that started life as an anaphoric name can come to be used deictically if an object satisfying the presuppositions becomes available for demonstration. “Neptune,” for instance, was introduced to designate the (then unobserved) planet that purportedly caused the perturbations of the orbit of Uranus, and was used for this purpose for a time before the planet was first observed. “Neptune” is *at present* not bound by (or even generally associated with) this description. The discovery of Neptune made a semantically relevant difference. Since “Neptune,” prior to discovery, was stipulated to be whatever planet was responsible for the perturbations, we could not have discovered that it was not (even if we knew it was only contingently responsible for the perturbations). If the name had become associated with a non-equivalent antecedent, sentences containing it would have different truth-conditions. After anchoring the name-using tradition in a demonstrated referent, however, it would in principle be possible to discover that Neptune was not responsible for those perturbations by discovering that the new anchor for the tradition lacked this property.

Given that a name functions as a pronoun and receives a referent only in the context of a particular discourse, the transition from anaphoric to deictic name is not a matter of the name changing meaning, but a matter of a single name (expression) being used differently in different circumstances. After demonstrating the object, we were left with two sources of interpretation for “Neptune;” one supplied by description, one by demonstration. The meaning did not change – the lexical meaning is still “thing called ‘Neptune,’” and the proffered content is that of a variable. Rather, the context of

discovery facilitated the felicitous use of a demonstrative “*d*,” and allowed us to conclude that Neptune = *d*. Hence, what we have is not a semantic change constituted by a change in epistemic situation. Instead, our improved epistemic situation allowed us to use demonstratives felicitously, and the *semantic properties of demonstratives* made it possible to use “Neptune” as a deictic rather than anaphoric name.

The important change would then be the subsequent (implicit) decision to let the anchor take precedence over the original stipulation. Since names prefer accommodation of referents over binding, it is unsurprising that the change occurs smoothly, although one may imagine a period of indeterminacy insofar as uses of the name would be associated with an (indeterminate) intention to use it the way it has traditionally been used. This suggests that if one discovered right after the observation that the planet observed was not responsible for the perturbations of the orbit of Uranus after all, one may be inclined to deny that one had discovered Neptune rather than accept that Neptune was discovered but did not have the properties previously associated with it. But this seems correct.

6.4.5 Transitions in the Other Direction: Indeterminacy?

Transitions from a referential into an anaphoric name might be less plausible. It is certainly possible that “Thales” has at present become a descriptive name as well. No one associates many properties with “Thales” except that he was the philosopher who thought everything was water (he was probably not even called “Thales”), and it is unclear whether “Thales” even satisfies the semantic argument anymore. Given the paucity of descriptions associated with Thales, it may in fact be impossible in principle to discover that he does not satisfy them – what would count as making this discovery? If such

discovery is impossible, the name is, at least in practice, immune to the semantic argument. Names such as “Homer,” and perhaps “Moses” and “King Arthur,” are other potential cases for which similar questions may arise. Such cases are controversial, however. It may also be difficult to draw a line between names introduced by description and by demonstration because it is somewhat difficult to determine the limits of what counts as *being demonstratively available*. The name “55 cancri b” may be a case in point – 55 cancri b is an exoplanet the existence of which is inferred by variations in its star’s radial velocity rather than by direct observation and demonstration.

Finally, there may also be descriptive uses of names that would, in other contexts, be used referentially (with the same denotation). Suppose someone says that “there is an attorney in York, called ‘Jack’ I think, who specializes in this area. Jack might help you.” Would the occurrence of “Jack” in the second sentence count as bound by the description “attorney in York called ‘Jack?’” Even if it is not, we may be inclined to treat it as bound to make sense of my subsequent beliefs about this attorney. Indeed, even when names are not bound, we may perhaps assume that audiences otherwise unfamiliar with a referent may interpret a name denoting it as bound by information given about the denotation. If this is the case, we may wish to treat referential names as bound for instance in belief ascriptions. We return to these issues in Chapter 7 when we discuss Frege’s puzzle.

6.5. Conclusion

We have argued that names and pronouns are closely enough related that they should be assigned to a single grammatical category. In this chapter we traced the differences

between proper names and ordinary pronouns such as “he,” “she,” or “it” to differences in character, and spelled out the consequences of this difference. We also accounted for the differences between, and different uses of, anaphoric and deictic names. The resulting pluralism about names gives us novel resources to deal with several familiar problems concerning proper names, and we will explore these resources in the next chapter.

CHAPTER 7: SOLUTIONS

Direct reference has commonly been associated with two major problems. First, if the semantic contribution of a name is its referent, then what should we say about *prima facie* meaningful names that do not refer to anything? “Vulcan” does not refer to anything, yet seems to be meaningful, and to mean something different than “Santa,” which does not refer to anything either. Furthermore, whereas “Vulcan exists” is false, “Vulcan does not exist” is intuitively true. Second, “Hesperus” and “Phosphorus” refer to the same object, yet seem to make different semantic contributions, insofar as it is possible to deny that “Hesperus is bright” yet assent to “Phosphorus is bright.” Now, given that Frege’s puzzle primarily arises in contexts of belief ascriptions, it is less obvious – if we assume a Russellian view of names – that the problem even ought to have a semantic solution at all; Russellian views explicitly disassociate semantic content from cognitive significance. Empty names are problematic, however, and a third substantial problem is the problem of the contingent *a priori*, discussed in chapter 5, for which pragmatic solutions seem less promising. This worry, however, is often neglected, partially perhaps because descriptive names have often but mistakenly been thought to be rare.

It is worth emphasizing the difference between Frege’s puzzle and the problem of the contingent *a priori*. One may get away with deeming “Jack believes that Hesperus = Phosphorus” true in virtue of Jack’s belief that Hesperus = Hesperus, even if Jack vigorously denies that Hesperus is, indeed, Phosphorus. The ascription does, after all, have a true reading if names are directly referential, and if we do not base semantics on cognitive significance this might simply be its semantic content. One cannot get away with deeming

“Jack knows that Judson invented the zip” to be true just because Jack knows that zips exist and are artifacts. That ascription has *no* true reading in this case. It does not matter whether semantics is based on cognitive significance or not.

If descriptivism were correct, we avoid the problems. Yet there is overwhelming evidence that names are rigid designators. Some philosophers have tried to combine descriptivism with rigidity by arguing that names are equivalent to rigidified descriptions. In a framework of two-dimensional semantics such versions of descriptivism have also promised elegant solutions to the problems (Chalmers 1996; 2006; 2012; Jackson 1998a; 1998b; Garcia-Carpintero, 2006). We have argued against treating names as rigidified descriptions (see also Soames, 2002, 2005). Furthermore two-dimensional descriptivism often involves a specially designed technical apparatus. In this chapter I first describe how our pluralism licenses a straightforward way of incorporating the motivations behind two-dimensionalism without relying on multiple intensions or propositions associated with single sentences, without making names equivalent to descriptions, and without invoking specially designed modal operators or new notions of necessity. Nevertheless, the two-dimensional behavior of singular terms will come out as a semantic phenomenon (as opposed to the metasemantic two-dimensionalism of Stalnaker, 1978, 2001, 2006, or the “metasyntactic” view of Bochner, 2013). The idea is really very simple. If we treat names as variables semantically bound by (generally) wide-scope quantifier expressions we can capture their apparently two-dimensional behavior by scope interactions between various operators. The result is a much-simplified system that captures many benefits of two-dimensional semantics in a one-dimensional framework using only resources we need in any case, such as semantic variable binding. In addition, the result is at least consistent with (though does not require)

treating propositions as Russellian, structured entities, and with a neo-Russellian rather than neo-Fregean approach to reference.

In 7.1 I briefly describe a standard version of two-dimensionalism and how it solves the problem of apparently contingently *a priori* sentences. In the following sections I show how these results can be captured by mechanisms we need in any case to account for the modal behavior of anaphoric pronouns. Descriptive names should be treated in the same way, as singular terms represented by (semantically if not syntactically) bound variables. In the following sections I apply the framework to provide novel solutions to the problems of the contingent *a priori*, empty names and, more tentatively, Frege's puzzle.

7.1. Two-Dimensional Semantics

Let us first motivate a two-dimensional treatment of descriptive names. Assume, following Evans (1979, 1982) that the name "Julius" is introduced by

- (1) Let us use "Julius" for whoever invented the zip (if anyone did).

"Julius" is sensitive to Kripke's modal argument and therefore rigid. We do not use descriptive names in ways that make the following true (Evans 1979, p.192):

- (2) If Julius had never invented the zip, he would not have been Julius.

That (2) is false shows that “Julius” is not equivalent to the (non-rigid) description “the inventor of the zip.” On the other hand, “Julius” is not directly referential either. Consider:

(2*) If anyone invented the zip, then Julius invented the zip. [Evans, 1979, p.193]

If “Julius” were referential, then (2*) would be contingent, even though we know that it is true *a priori* since (1) is a stipulation.

In the actual world Judson invented the zip, so “Judson invented the zip” and “Judson is Julius” are both true. But relative to a world in which, say, McTaggart invented the zip, “McTaggart invented the zip” and “McTaggart is Julius” should be true. So regardless of what “Julius” in fact denotes (2*) expresses a truth. “Julius” seems to be insensitive to Kripke’s *semantic argument* against descriptivism.⁹⁴ Since “Julius” is stipulated to designate whoever invented the zip, (3) has a true reading:

(3) If we were to discover that McTaggart, rather than Whitcomb Judson, invented the zip, then “Julius” *would* refer to McTaggart.

Since “Julius = Judson” is necessarily true, if true, what kind of modal operator captures the consequent of a true reading of (3)? It is tempting to suggest that it invokes an *epistemic* operator and thus a set of possibilities along a different dimension than the set of *metaphysical* possibilities associated with the operator in (2). However, even though we

⁹⁴This is not to deny that the name “Julius” could come to be more tightly associated with Judson when we learn that Judson is, in fact, the inventor of the zip, and hence come to resist assignment changes just as the semantic argument predicts. This, however, is a different matter (see the previous chapter); all I need is that the use of “Julius” at the time of introduction *t* resists the semantic argument at *t*.

naturally employ a stipulation like (1) only if we lack epistemic access to Judson (and perhaps need a rigid expression to describe counterfactual situations involving the inventor of the zip without ambiguity), it still seems to be the case that “Julius” avoids the semantic argument in virtue of being semantically stipulated to do so, not in virtue of our epistemic situation with regard to Judson. As argued in chapter 4, if there are lessons for epistemology here they appear to stem from a semantic distinction rather than vice versa.

7.1.1. The Two-Dimensional Framework

In a Fregean, two-dimensional framework “Julius” is interpreted as equivalent to an *actually*-rigidified description.⁹⁵ “If anyone invented the zip, Julius did” may be analyzed as (4) (“Julius” is the underlined part):

$$(4) \exists x(x \text{ invented the zip}) \rightarrow [\exists x: \underline{@(x \text{ invented the zip})}] (x \text{ invented the zip}).$$

(4) is contingent but can be recognized as true by anyone who understands the trivial

$$(5) \exists x(x \text{ invented the zip}) \rightarrow [\exists x: x \text{ invented the zip}] (x \text{ invented the zip}).$$

Hence, (4) appears to be an example of the contingent *a priori*. The idea behind two-dimensional semantics is that we can associate “if anyone invented the zip, Julius did” with two intensions, one also expressed by the *a priori* but necessary truth (5), the other as the

⁹⁵ I will assume that definite descriptions are restricted quantifiers. Indeed, for simplicity I represent definite descriptions simply as existential quantifiers. Nothing substantial should hinge on this.

contingent (but *a posteriori*) truth concerning the referent of “Julius” that *he* invented the zip if anyone did. We can then illustrate the behavior of “Julius” with respect to the following matrix (adapted from Stalnaker, 1978):

	<i>i</i>	<i>j</i>	<i>k</i>
<i>i</i>	T	F	F
<i>j</i>	F	T	F
<i>k</i>	F	F	T

Let *i*, *j* and *k* be index points (worlds), and let the horizontal axis represent worlds of evaluation. Relative to the matrix we define the following operators (Davies & Humberstone, 1980; Davies, 2006):

- i) $\Box\phi$ is true iff ϕ is true at each index point along the horizontal axis, or more precisely, $\Box\phi$ is true w.r.t. world *i* iff ϕ is true with respect to $\langle i, w \rangle$ for every *w*.
- ii) $@\phi$ is true with respect to a world *w* iff ϕ is true with respect to the designated or actual world w^* of the model. If $@\phi$ is w.r.t. to *any* world, then it is true with respect to *every* world, so $@\phi \rightarrow \Box@\phi$ is a theorem.

We can also capture the intuition that there is a sense in which $@\phi$ is contingent by noting that it is a contingent matter which world is denoted by $@\phi$, and we capture this by varying

not only world of evaluation but also which world plays the role of the actual world. Let the vertical axis represent worlds “considered as actual:”

iii) $\mathcal{F}\phi$ (“fixedly ϕ ”) is true iff ϕ is true at each index point along the vertical axis.

More precisely, $\mathcal{F}\phi$ is true w.r.t. world i iff ϕ is true w.r.t. $\langle w, i \rangle$ for every w .

If the embedded sentence ϕ contains no occurrence of $@$, then $\mathcal{F}\phi$ is equivalent to ϕ .

Combining the rule for $@$ with the rule for \mathcal{F} , we get:⁹⁶

iv) $\mathcal{F}@ \phi$ is true iff for every world w , ϕ is true with respect to w with w also playing the role of the actual world. More precisely, $\mathcal{F}@ \phi$ is true iff for every w_n , ϕ is true with respect to $\langle w_n, w_n \rangle$.

In other words, $\mathcal{F}@ \phi$ is true iff for every world w , ϕ is true with respect to w with w also playing the role of the actual world. $\mathcal{F}@$ is not in general equivalent to the necessity operator \Box , for whereas $\Box(@\phi)$ is true, $\mathcal{F}@(@\phi)$ is equivalent to $\mathcal{F}@ \phi$, and so to $\Box \phi$, which is false if ϕ is contingent.

With \mathcal{F} alongside \Box and $@$ we can express two notions of necessity. $\mathcal{F}@ \phi$ requires that ϕ is true with respect to every cell on the *diagonal*, or in other words: $\mathcal{F}@ \phi$ is true if ϕ is

⁹⁶ Another way of achieving the same results could employ subscripted modal operators that turn the evaluation of a formula to a world introduced by a preceding possibility or necessity operator with the same subscript. See Stephanou (2001) for details.

true with respect to any pair $\langle w, w^* \rangle$ if $w = w^*$, that is, if the world that plays the role of the world of the context is the same as the world of evaluation. That a sentence is true with respect to any world w considered as actual, meaning that the diagonal intension yields *true* for any pair $\langle w, w \rangle$, does not entail that the sentence is necessarily true, that is, that the horizontal intension yields *true* for every world.⁹⁷

Let i be a world in which Judson invented the zip, and j a world in which McTaggart invented the zip and Judson did not. (4) is “horizontally” contingent since it is false in the j -column at the i -row. If we assess (4) on the j -row, it is false at the i -column but true at the j -column; (4) is still contingent when we change the world considered as actual, but it is true at different worlds on the horizontal depending on which row we consider. (5), on the other hand, is true with respect to all cells. It turns out that (4) has the same $\mathcal{F}@$ -intension as (5) since (4) is true at every cell on the diagonal; no matter which world w is considered as actual the contingent (4) is true at w . Our *a priori* knowledge that (4) is true can then be captured by its $\mathcal{F}@$ -intension. The fact that (4) is contingent, regardless of assignment, is captured by its horizontal intension, which yields *true* at some worlds, *false* at others.

Furthermore, the necessity of identity entails that “Julius = Judson” is horizontally necessary, or \Box -true, but not that it is $\mathcal{F}@$ -true. If McTaggart had actually invented the zip (i.e. move to the j -world on the j -row) then the non-modal sentence “McTaggart invented the zip” would be true and so would “McTaggart is Julius.” It does not follow that the modal sentence “if McTaggart had invented the zip then McTaggart would have been Julius” is

⁹⁷ With respect to the system, however, it may be natural to associate with every sentence three rather than two intensions, one corresponding to the vertical axis, one to the horizontal axis (closely related to Chalmers’s *secondary intension* or Jackson’s *C-intension*), and one to the diagonal (closely related to Chalmers’s *primary intension* or Jackson’s *A-intension*).

true, for in this case we are asked to consider world j on the i -row, the row in which Julius did not invent the zip at j , but in which the value of “Julius” at j is nevertheless Judson since Judson is the value assigned at the $\langle i, i \rangle$ cell.⁹⁸

7.1.2 Worries Concerning the Two-Dimensional Framework

Despite this attractive result there are several reasons to reject the treatment. First, the \mathcal{F} operator does not correspond to any intuitive notion of necessity and seems, essentially, to be a context-shifting operator of the kind that Kaplan labeled a “monster.” Now, the \mathcal{F} operator would only behave as a monster if @ is an indexical. Davies & Humberstone (1980) deny that the actual world should be understood as an aspect of context, and treat @ as an operator rather than an indexical. It is less clear that this entirely avoids the gist of the worry; since $\mathcal{F}@ \phi$ is understood as saying that whichever world had been actual, ϕ would have been true in the actual world, the behavior of @ under the scope of \mathcal{F} still seems to require that \mathcal{F} is understood as a context-shifting operator. Now, the non-existence of monsters or context-

⁹⁸ It is unclear how well the system captures what Evans (1979) had in mind. In Evans (1982) the axiom for descriptive names (chapter 4) the description is not supposed to be a synonym for the name but to provide the assertoric or cognitively significant content. But if the name is not synonymous with the description, it should not be synonymous with the *rigidified* description either. Evans called sentences such as (2*) *superficially contingent* but *deeply necessary*. The notions correspond roughly to necessity and *aprioricity*, though for Evans’s the former distinction belongs to the level of truth conditions and the latter to the level of sense. Evans never characterized deep necessity in terms of modal operators; rather, the deep contingency or necessity of a sentence depends on what makes it true. Of course, what makes ϕ true is the same state of affairs that makes @ ϕ true, and one may suggest distinguishing truth *with respect to* a world from truth *in* a world, where the latter is given as “if w were to obtain, or were to be actual, then [ϕ] would be true” (Evans, 1979, p. 207); thus, a sentence is deeply necessary iff it is true *in every possible world*. For Evans, superficial necessity and contingency are hence matters of *properties (truth-values) of modal sentences*, whereas deep contingency and necessity, notions that according to Evans are not purely internal to semantic theory, are *modal properties of (non-modal) sentences* (Davies, 2006, p.150). Two-dimensionalism still entails the rejection of Evans’s claim that deep necessity is not explained in terms of a theory-internal notion. In a two-dimensional framework, it is; deep necessity, truth in a world, corresponds to the notion of truth at a world considered as actual.

shifting operators in natural language is debatable,⁹⁹ and they are easily defined formally, but shifting the already artificial *world* parameter seems particularly troublesome – worlds (as opposed to agents, places or times) are not normally encountered, variable index points in ordinary discourse.

Davies (2006) also suggests introducing \mathcal{D} as a primitive modal operator for truth on the diagonal. $\mathcal{D}\phi$ would thus be true with respect to w_n , with w_m playing the role of the actual world, just in case for every world w the embedded sentence ϕ is true with respect to w with w also playing the role of the actual world. After all, not only does the \mathcal{F} -operator not correspond to any natural language notion of necessity; it does not obviously do any real work in the modal analysis of natural language sentences at all unless combined with the $@$ -operator. So why not rather have a primitive operator \mathcal{D} that does the job of $\mathcal{F}@$ on its own? One obvious worry is that \mathcal{D} does not really correspond to any natural language expression either. The operator is supposed to reflect epistemic modalities. If the operator in (3) corresponds to a genuine kind of modal operator, however, then \mathcal{D} , being primitive, seems to imply that counterfactuals are inherently semantically ambiguous between epistemic and metaphysical operations.

Furthermore, the interaction with $@$ is lost. On the more complex account it is the fact that “Julius” a $@$ -rigidified description that makes it behave the way it should in the matrix. If it is the \mathcal{D} -operator that does the trick (without reference to $\mathcal{F}@$), and without being a context-changing operator, it is not clear exactly what the reasons would be for it to

⁹⁹ See for instance Schlenker (2003). Schlenker’s examples concern hyperintensional contexts in particular, and he treats attitude verbs as context quantifiers.

produce the results it should – why is McTaggart rather than Judson assigned to “Julius” in $\langle j, j \rangle$ for instance? As far as I can see there is no intuitive way to account for this unless the \mathcal{D} -operator is a context-shifting operator if the \mathcal{F} -operator is.

Second, it seems unlikely that descriptive names are abbreviations of @-rigidified descriptions, especially if ordinary names are not. Nothing about the introduction of “Julius” suggests that the apparently non-modal sentence “Julius invented the zip” should have any special modal property or be associated with more than one intension. Indeed, the stipulation that “Julius” refers to the denotation of “the inventor of the zip” did not deploy the @-operator at all, and understanding “Julius” does not require understanding that operator – it is at least conceivable that a linguistic society could have descriptive names without having that operator (Davies 2006).¹⁰⁰

Furthermore, and more generally, analyzing names as @-rigidified description incurs an explanatory commitment that is not easily cashed out, as we argued in Chapter 1. If names are directly referential singular terms, then their rigidity is explained – without invoking modality – by the absence of any semantic element in the interpretation that could be used to individuate different referents at different worlds. Hence, direct reference explains the rigidity of a name in virtue of the requirement that only the referent is relevant to determining the truth-conditions of sentences in which it occurs. The name’s modal behavior is thus predicted by its non-modal properties. If names were @-rigidified descriptions, however, modal properties would be built into the semantics of names themselves. The

¹⁰⁰Soames (2005) provides a thorough critique of two-dimensional semantics as well. Many of his arguments are based on contexts that mix modal and attitudinal operators, however, and they appear to rely on treating @ as an indexical rather than as an operator, contrary to what Davies & Humberstone presented assumes. See also Pickel (2012) for a discussion.

behavior of the name in modal contexts is correctly predicted, of course, but if we ask *why* names behave the way they do in modal contexts we encounter what seems to be an uncomfortably tight circle. The obvious answer, “because of the @-operator,” refers to the name’s modal behavior in order to explain that behavior. There may be ways around this circle, or reasons not to worry, but at the very least direct reference theories have an advantage in providing a non-modal, non-circular explanation of the behavior of names in modal contexts.

7.2. Variables

We can capture the two-dimensional behavior of descriptive names without relying on the mechanisms discussed above, if we instead treat singular terms as variables.

7.2.1 Pronouns in Modal Contexts

Consider an example we have already discussed:

(6) One man wrote the score for *City Lights*. That he was multitalented is a contingent truth.

The second sentence of (6) is true since it is a contingent truth that Chaplin was multi-talented. Likewise (6*) relies on how things are with Chaplin:

(6*) One man wrote the score for City Lights. If he had been lazier he would have left the job to someone else.

(6*) is true if there is a world in which Chaplin was lazier and did not compose the score but left the job to someone else. As (6) and (6*) show, anaphoric pronouns are rigid. If “he” in (6*) were given a standard D-type analysis, as a proxy for the non-rigid “the man who wrote the score,” then “he would have left the job to someone else” would be a contradiction. On the other hand, if the whole fragment (6) is evaluated at a different world in which Copland (and not Chaplin) wrote the score, the second sentence of (6) should be true if it is a contingent fact that *Copland* (not Chaplin) was multitalented. Hence “he” is not directly referential and is not affected by Kripke’s semantic argument.

If we treated “he” as [the x : $\mathcal{F}@ (x \text{ composed the score for City Lights})$] we would get the right results (merely using @ would just give us Chaplin). A more obvious alternative is to treat “he” as a variable bound by “one man.” Now, “he” is hardly *syntactically* bound by the antecedent “[o]ne man wrote the score for City Lights,” but as we have seen there are several ways to make “he” *semantically* bound, for instance in DPL, or by treating the antecedent as a domain restrictor for the variable interpreted as an identity function (as in UIF). Regardless of the exact mechanism the correct interpretation of the second sentence of (6) is one on which the value of “he” (the value that satisfies the sentence) is the same as the value that satisfies $\text{man}(x)$ in the evaluation of the first sentence. In standard Kripke semantics assignment functions are defined relative to worlds. For (6), and (6*) the world of evaluation for the antecedent determines an assignment of values to variables, and the value of the variable “he” (that is, the value that satisfies the whole fragment) is determined by the

antecedent. Of course, the bound occurrence of “he” is not a rigid *designator*; it does not designate anyone in particular but *ranges* over a domain. It nevertheless behaves rigidly in the sense that its value under an assignment is insensitive to any modal operator occurring in the second sentence itself (as captured e.g. by UIF). Thus the consequent of (7) will create an illusion of being contingent *a priori* if evaluated in isolation:

(7) If one man composed the score for City Lights, then he composed the score for City Lights.

Since any interpretation that satisfies (7) must assign the same value to “he” as it does to “ x ” in “man(x),” (7) is necessarily true. Nevertheless, relative to any assignment the consequent itself, “he composed the score for City Lights,” is contingent.

We obtain a pseudo-two-dimensional interpretation of “he” in a standard one-dimensional model by varying the world of evaluation for the whole fragment. Evaluating (6), for instance, relative to different indices and their associated assignment functions can be illustrated as moving down the vertical in our matrix above. At every vertical index w the assignment function associated with w assigns a value and *the second sentences* in our examples are satisfied if there is an appropriate assignment. At every index along any *horizontal* line in our matrix the second sentence is satisfied relative to the assignment associated with the world of the leftmost cell. The consequent of (7) viewed in isolation is thus contingent, but since it is satisfied by any assignment that satisfies the antecedent it is diagonally true.

7.2.2. “Julius” As a Variable

I submit that “Julius” is anaphoric on the (perhaps presupposed) antecedent definite description “the inventor of the zip” and hence bound by it, just as “he” is bound by “one man” in (6). Treating “Julius” as a variable semantically bound by its antecedent yields the modal behavior we want in a pseudo-two-dimensional (variable) system that avoids the worry that “Julius is the inventor of the zip” expresses an *a priori* contingent singular proposition. “Julius invented the zip” is contingent when evaluated on its own since the assignment that satisfies “Julius,” determined at the world of evaluation for the antecedent description “the inventor of the zip (if any),” only contingently invented the zip. “Julius” behaves rigidly with respect to this assignment. “ $\diamond\neg(\text{Julius invented the zip})$ ” is thus true, since the evaluation lets “Julius” be (semantically) bound by an antecedent that takes wide scope over the modal operator. The correct first-order interpretation of the sentence is:

$$(8) [\text{the } x_{\text{julius}}: x_{\text{julius}} \text{ invented the zip}] \diamond\neg(x_{\text{julius}} \text{ invented the zip}).$$

“Julius invented the zip” comes across as *a priori* since the full interpretation, relativized to the antecedent (assuming that an inventor exists), is the trivial:

$$(9) ([\text{the } x_{\text{julius}}: x_{\text{julius}} \text{ invented the zip}] (x_{\text{julius}} \text{ invented the zip})).$$

After all, “Julius invented the zip” seems *a priori* because “Julius” was introduced by the stipulation (2), and once that stipulation is made semantically explicit, as in (9), the result is

not only *a priori* but trivial. “Julius invented the zip” looks contingent when viewed in isolation – not because the binding description is inert, but because (10) is false:

$$(10) ([\text{the } x_{\text{julius}}: x_{\text{julius}} \text{ invented the zip}] \Box (x_{\text{julius}} \text{ invented the zip})).$$

On the other hand, “Julius = Judson” will be necessarily true if the assignment-providing antecedent is evaluated relative to the actual world and necessarily false if the antecedent is evaluated with respect to a world in which Judson did not invent the zip. Notice also that “if someone invented the zip, then Julius did” will at least have a reading on which it is contingent as well; (11) is false:

$$(11) \Box ([\exists(z): z \text{ invented the zip}] \rightarrow [x_{\text{Julius}} \text{ invented the zip}])$$

In (11) the modal operator takes wide scope over the (presupposed) antecedent that binds x_{Julius} . Given that the assignment is fixed by the antecedent (relative to world i in the matrix above), the world j , where McTaggart and not Judson invented the zip, provides a counterexample to (11).

Of course, “Julius invented the zip” will generally be uttered with the assignment-determining antecedent occurring earlier in a discourse, or even as a tacit presupposition. Thus, the antecedent will (generally) take maximal scope. But there is nothing preventing us from evaluating, say, the whole of (8) with respect to a world in which McTaggart rather than Julius invented the zip (such worlds would e.g. be part of the context set for the discourse). This is the interpretation that allows us to obtain a true reading of (3) and the

reason “Julius” is, at least on one reading, not affected by the semantic argument. The matrix above captures this behavior. The horizontal axis provides the modal profiles of “Julius invented the zip” relative to assignments. The vertical axis gives us the results of evaluating the whole fragment at different assignment-providing worlds. Since “Julius invented the zip” will be true at any world providing a satisfier for (9), “Julius invented the zip” is diagonally true. Thus, it will be natural to conclude that “Julius invented the zip” is *a priori*. Of course, the behavior is still a result of the semantic behavior of “Julius,” not epistemic concerns, but the semantics of “Julius” matters for understanding (and reflects certain facts about the context of introduction for) the name. “Julius is the inventor of the zip” is *a priori* since “the inventor of the zip invented the zip” is, but the former under “ \square ” is equivalent to (10), not to (9) under a wide-scope modal operator.

But the system I am suggesting is not two-dimensional, at least not in the ordinary sense. There is no double set of intensions involved, and we do not need a variety of modal operators since we have no @-operator “pulling us back” to the actual world that we need to manipulate. Rather, the two-dimensional matrix above is used to illustrate the behavior of embedded modal operators and why it is that certain embeddings – or interactions between ordinary modal operators – create the illusion of *a priori* contingent sentences. It really has nothing to do with epistemology at all, nor with indexicals, context-shifts, or non-standard modal operators. Instead, the behavior of descriptive names is a result of the relatively familiar phenomenon of scope interactions between modal operators and variable binders.

“Julius” is a singular term. It is not equivalent to a description, not even a description that is required to take wide scope in modal contexts (so our approach differs fundamentally from Dummett, 1993), but it can be *bound* by an antecedent description. The rigidity of

“Julius,” however, is simply a result of the fact that modal operators occurring between the bound occurrence of a variable and its binder cannot affect which individual satisfies a sentence. Relative to the actual world “Judson is Julius” is true since the satisfier is Judson, and necessarily so since “Julius” is rigid. Given a different world of evaluation “Julius” could nevertheless be someone else, since the necessity of identity is a matter of truth on the horizontal; so “Julius” does not satisfy the semantic argument. As its surface form suggests, however, “Julius invented the zip” contains no modal operator at all; the apparent two-dimensional interpretation falls out of a natural interpretation of the non-modal sentence relative to different worlds. Furthermore, we need to recognize (semantically) bound singular terms whose contributions must be understood in light of antecedent expressions – and the two-dimensional behavior that is the result – in any case if we are to obtain any satisfactory account of the modal behavior of anaphoric pronouns, independently of descriptive names. That we get a promising account of such names out of it, one that emulates standard two-dimensional treatments, is a bonus.

7.3 Empty Names

Treating “Julius” as a term whose interpretation depends on a previously occurring description requires that a presupposition affects the semantic analysis (see Chapter 6), though one may pick up the correct interpretation by assuming that Julius is governed by “whoever the other participants in the discourse is calling ‘Julius.’” The associated description is not the meaning of the name; the range of “Julius” is (if anything). As discussed in Chapter 6, the antecedent does not need to be identical across contexts or

between users for “Julius is *F*” (interpreted as an open sentence) to make the same truth-conditional contribution. Insofar as one may call anything a “name-using” tradition for “Julius” (and there will of course in practice be information that remains associated with the name across contexts), requiring intensional equivalency of antecedents across contexts is probably a too strict requirement. Insofar as descriptive names do not satisfy the semantic argument, potential change across contexts is not a problem, but what we want. In 7.3 I discuss some alternative proposed solutions to the problem of empty names, and some conditions an adequate theory must meet (Meinongian and gappy propositions views are discussed in 7.4.3). In 7.4 I apply our pseudo-two-dimensional system (combined with considerations on discourse interpretation from Chapter 5) to solve the problem(s) of empty names and negative existentials.

7.3.1 Descriptive and Empty Names

Although anaphoric names are generally introduced for context-specific purposes, empty names – names purporting to name objects that happen not to exist – provide examples of descriptive names that may survive the context of introduction. All empty names are anaphoric names. A name that works as a constant does so because it inherits its semantic content from a demonstration used to introduce the name. There may in some cases be a legitimate debate about what a demonstrative demonstrates, or exactly when it fails, but reference failure is a sufficient condition for demonstration failure. A failed demonstrative has no semantic content. Any term supposed to inherit its semantic value from a failed demonstrative, inherits nothing.

That does not mean that when demonstratives misfire *nothing* is introduced. If Jack attempt to demonstrate a planet and introduce it as “Blob,” and there is no object that he demonstrates, we need to engage in repair to interpret subsequent sentences containing “Blob.” We could, of course, treat these sentences as literally meaningless, but often – at least if we are uncertain whether Jack saw anything or not – we treat “Blob” as “whatever planet Jack talked about.” This (non-denoting) description then serves as a binder for the empty, descriptive name. All empty names (and traditions associated with empty names) originate in either a descriptive stipulation, or in some sort of repair effort generated by the failure of a demonstration. Such repair may be unintentional. For instance, suppose Jack falls into a crevice and yells “get me a line.” Someone could take him to say “get me Elaine,” and – without knowing who Elaine may be – initiate a name-using tradition by deferring reference to Jack. In this case, however, the name-using tradition is initiated by the description “whoever Jack was talking about” – which was, of course, no one.¹⁰¹

Empty names are descriptive names, and since empty descriptive names are relatively common (think of various religious or mythological beliefs) there ought to be a significant number of parallel cases where the introducing description is non-empty, insofar as it is not the fact that the name turns out to be empty that makes it descriptive, but the semantic properties of its introduction. Such cases are probably harder to (uncontroversially) identify, however. Moreover, such names quickly turn into referential ones if an appropriate referent becomes available for demonstration (Chapter 6), or if the use of the name merges with a referential tradition. There may also be cases that, since the names are used with deferential

¹⁰¹ Thanks to Hilpinen for the example.

intentions (“whoever the experts are talking about”) at least ordinary users would be unaware that they are using a descriptive name.

7.3.2. Empty Names and Fiction

We may wish to distinguish empty names from names of fictional characters. When writers of fiction introduce fictional characters, they are not attempting to refer to or pick out persons bearing the name given to them in the stories. As Schiffer points out, when James Joyce introduced the name “Buck Mulligan”

Joyce was not trying to refer to a man named ‘Buck Mulligan’ and failing miserably; he was, in the way characteristic of fiction, *making as if* to refer to a man with that name and to tell us something about him [Schiffer, 2003, p.50].

It is not entirely implausible that Joyce, by engaging in a pretending use of the name “Buck Mulligan,” created something whose name is “Buck Mulligan” and which is the referent for occurrences of that name. As opposed to “Vulcan,” “Buck Mulligan” may not be entirely empty. Although Mulligan is not a person, Mulligan is a fictional character, which (arguably) exists *as such* and can truthfully be ascribed properties such as being well known, and falsely ascribed others such as being a person (it can presumably be truthfully ascribed these properties in a pretense or fictional context, be fictionally ascribed these properties, or be ascribed fictional properties that can be mapped onto real properties such as *brilliance* and *personhood*). According to Schiffer (2000, 2003) Mulligan is a *pleonastic* entity, a minimal abstract entity created by a *something-from-nothing transformation*; that is, by the very fact that Joyce decided to use the name “Buck Mulligan” in a pretending manner. Other participants in a name-using tradition can know of Mulligan’s existence (as a fictional

character) and refer simply by engaging in a particular language game (*ibid*, p.52).

Metaphysically, fictional characters are “mere shadows of the pretending use of their names” (*ibid*, p. 59); there is nothing more to their nature than what is determined by the hypostatizing language game that recognizes them in our ontology. (A similar view is defended by Thomasson, 1999; 2003.)

I remain agnostic about whether Schiffer’s account is correct about fiction. I am sympathetic to Sainsbury’s (2007) arguments against abstract artifact theories of fiction. On such accounts, including Schiffer’s, a fictional story is a representation of fictional states of affairs, but this picture seems to get the cart before the horse; fictional accounts are representational, but characterized precisely by being representations of nothing. But according to abstract object theorists (and Meinongians), when I entertain the idea of Sherlock Holmes playing the violin, I am representing an abstract (or non-existent) object, Holmes, as *playing the violin*. Though false, the thought is true-in-fiction since Holmes *encodes* but does not *exemplify* this property. It remains obscure, however, what it means for an abstract object to encode a property that my thought then represents as being fictionally exemplified; it seems preferable to say that it is the *representation* (e.g. in the novel) of Holmes that encodes the property and does not exemplify it any more than a text about Caesar’s rise to power exemplifies *rising to power* (though it encodes this property by being a text about Caesar’s rise to power). In other words, abstract artifact theories appear to attempt to push features that properly belong at the level of representation down to the level of things represented. The approach to anaphoric names described in this section appears to

be eminently compatible with a representationalist (or irrealist) account of fiction, but developing this view in detail goes beyond its scope.¹⁰²

Even if Schiffer's view is correct it would force us to distinguish fictional and empty names. Whereas fictional names are names for fictional characters, introduced in a context of pretense, empty names are names introduced to pick up existing objects. They are defective since there is no object satisfying the descriptions used to individuate the putative referent. As opposed to Joyce in Schiffer's example, Leverrier *was* trying to refer to a planet called "Vulcan" and failing miserably. He was not engaging in pretense.¹⁰³ So *if* Schiffer's account is successful for fictional names, we must recognize a distinction between fictional and empty names, and the account cannot explain the latter.

Now, if we did treat fictional characters as anaphoric names claims about Buck Mulligan would be meaningful, and perhaps even get the correct truth-values. Given that "Buck Mulligan" was introduced by a description containing something suggesting that the referent is a character of fiction, then "the fictional character that is the first to appear in *Ulysses*" could very well serve as a binder and "Buck Mulligan" be a bound variable. Mulligan would of course not be a medical student, but could presumably fictionally be a medical student, or *encode* this property. As long as the introducing description is explicit

¹⁰² For other critiques of abstract artifact theories, see Yagisawa (2001) and Brock (2010).

¹⁰³ Adopting something like Schiffer's account has the benefit of letting us distinguish between names for fictional characters, to whom various properties are ascribed under pretense in a story, and names that are empty even within the fiction, such as "George Kaplan." The latter may then be treated as an anaphoric name, even within the scope of a pretense operator.

about the fictional nature of the character, we can make both claims that are true *simpliciter* and claims true in fiction under its semantic scope.¹⁰⁴

Schiffer's account does at least not work for empty names. Admittedly the distinction between names for fictional characters and empty names may be difficult to draw in practice. "Vulcan," as used by language-users *at present* with knowledge of the status of Vulcan, may be a fictional rather than empty use of the name, which would make the semantic properties of "Vulcan" in discourse at present different from what it had for Leverrier in a manner parallel to the semantic properties of "Neptune" before and after it was observed, respectively. In general, if empty names can turn into names for fictional characters it will be somewhat hard to come up with an unequivocal example of an *empty* name, since a name may end up being empty (rather than a name for a fictional character) only as long as language users do not know that it is empty – there are presumably plenty of examples even in contemporary linguistic societies, but the very nature of emptiness ensures that it is difficult to locate clear examples.

Even granting that there is a sense in which Mulligan exists we should not grant that there is a similar sense in which Vulcan exists (at least if we imagine ourselves in the context of the scientists speculating about the existence of Vulcan). For instance, given that it is already true that Joyce wrote *Ulysses* using the name "Buck Mulligan," adding the pleonastic entity Mulligan to our ontology will amount to a conservative ontological extension. We already agree that it is *true* that Buck Mulligan exists according to the fiction; so going from *the fictional existence of Mulligan* to the existence of *the fictional Mulligan* does not commit

¹⁰⁴ Currie (1990) suggests treating fictional characters as something akin to bound variables, but the proposal is very different from the one defended here. Since fiction is not my primary topic I will not discuss the idea in detail.

us to any truth maker we are not already committed to. It is hard to see adding Vulcan as a conservative extension. One would presumably have to say that Leverrier *unintentionally* introduced a fictional name referring to a pleonastic entity. But this is implausible. As opposed to Mulligan, who satisfied the descriptions Joyce associated with it (“a fictional character appearing in the fiction *Ulysses*”), Vulcan is *not* the denotation of the description Leverrier used to introduce the name, since this description did not invoke *fiction*. So whereas Joyce was able to individuate the referent, Leverrier must have failed wildly. Indeed, whereas the existence of Mulligan supervenes on our practices with the name involving the pretense of fiction, Vulcan does not, since Vulcan systematically fails to satisfy the descriptions we associate with it.

Pleonastic transformations are supposed to be relatively well-understood affairs that do not induce new ontological commitments. For instance, given that “Carthage was destroyed” is true, we can introduce a name for the event, *the destruction of Carthage*, without thereby making ontological commitments over and above those embodied in the first sentence. But a *false* sentence, “Chicago was destroyed,” should not introduce an event, *the destruction of Chicago* – adding it to our ontology even as a fictional object would at best make subsequent discussions about whether Chicago was (really, not fictionally) destroyed rather hard to keep track of. Similarly, introducing fictional characters is arguably innocuous once there is a fiction, a set of true sentence of the kind “according to the story, there is a medical student called ‘Buck Mulligan’ ...,” that forms a base for the pleonastic transformation, since what makes those sentences true is exactly what makes claims about Buck Mulligan true – no new truth-makers are needed. There is no similar straightforward ontological reduction for Vulcan. Furthermore, if Leverrier unintentionally lapsed into

pretense discourse it must have happened at a surprising point; the conclusion “there is a planet between Mercury and the Sun” is not part of the pretense (but straightforwardly false); it is not until Leverrier started *using* the new name that we need to appeal to pretense to explain what is going on.¹⁰⁵

7.3.3 Empty Names and Name-Using Practices

If “Vulcan” is meaningful in the sense of enabling sentences in which it occurs to be evaluated for truth or falsity in a non-defective manner we need a story not only about how it could be meaningful, but also a metalinguistic story about why it has the meaning it has and how it retains this meaning. I will argue that the main challenges to existing theories concern this metalinguistic aspect of the story. On the view defended in this dissertation its contribution is partially determined by antecedent presuppositions, and empty names may or may not make the same contribution to different discourses, and the metalinguistic constraints are accordingly more flexible.

To illustrate the metalinguistic challenge I will look at Sainsbury’s solution to the problem of empty names (Sainsbury, 2005). Sainsbury defends a truth-theoretic version of what he calls “minimal Fregeanism,”¹⁰⁶ RWR (“Reference Without Referents”). Given that a

¹⁰⁵ Not the least it seems correct to say that Joyce created Buck Mulligan, and wrong to say that Leverrier created Vulcan.

¹⁰⁶ Sainsbury’s “minimal Fregeanism” goes roughly as follows: For Frege an expression’s sense is public yet graspable by individual speakers, and determines reference. The sense is what an expression contributes to a thought about its extension. Frege saw, however, the problem of trying to identify *the* public sense of a given proper name, conceding that “we must really stipulate that for every proper name there shall be just one associated manner of presentation of the object so designated” (Frege, 1918, p. 359). According to Sainsbury, the knowledge required to understand a name is the knowledge possessed by “parties to a single name-using practice,” and the role of *sense* is “to make sameness of referent manifest; likewise the main work that must be done by difference of sense is to fail to suggest sameness of reference” (Sainsbury, 2005, p.14). Successful interpretation requires recognizing this guarantee, but as Kripke’s Paderewski cases show this manifestness is not guaranteed to elicit the right cognitive response from an audience. Sainsbury thus

theory of meaning should be “empirically adequate” and that the goal of a truth-theoretic view is to make sense of speakers’ beliefs, the fact that “Santa” and “Vulcan” are associated with different practices and beliefs means that they must have distinct meanings and require different referential axioms. Sainsbury suggests that names are associated with *reference-conditional* axioms on the form:

$$(12) \forall x(\text{“}\alpha\text{” refers to } x \text{ iff } x = \alpha) \text{ [Sainsbury, 2005, p.93]}$$

The function of a referring expression is to introduce an object, but instances of functional kinds may fail to have a function definitive of its kind (a car might stop working though still be classified as a car by its function), and a referring expression can fail in the function of introducing an object and still be a referring expression. Sentences containing empty names are treated as false in Sainsbury’s associated negative free logic.

Sainsbury presents a putative dilemma for anyone who agrees that empty names are meaningful but denies that they are referential expressions (*ibid*, p.102): either the proposed description serving to fix a name’s meaning will be inconstant across the population using the putative name, and so the publicity of meaning is lost; or we must require that some

refuses to *identify* the sense of a name with what Dummett would call “full-blooded” condition, a condition the grasp of which would be sufficient to grasp the meaning of a name for someone who did not already possess the relevant singular concept. The view is hence closer to McDowell’s (1977) attempt to spell out a name’s sense in terms of “austere,” homophonic reference axioms, axioms that are interpretive – yield truth theories that *display* the sense of a name – but contain no particular, specifiable descriptive information (and is hence what Dummett calls a “modest” theory of meaning). The reason Sainsbury avoids *propositions* is related. Propositional views, at least Russellian ones, disassociate the semantic content of (an utterance of) a sentence from the cognitive content of one who understands it. But if an adequate “theory of meaning is a theory of understanding” (Dummett, 1973, p.9), such a level of meaning is otiose – the goal of a semantic theory is not primarily to analyze the semantic properties of individual terms, but to explain the (linguistic) actions of *speakers* in terms of their beliefs and intentions. To know the meaning of an utterance of a sentence is to know its truth conditions – no more fine-grained reflection of content is necessary or possible – and to know the meaning of a term is to know the contribution it makes to the truth-conditions (which on Sainsbury’s view amounts to “immersion in a [name-using] practice”).

constant meta-linguistic description is in play, a condition that is too sophisticated to attribute to ordinary speakers (attempts to analyze “Vulcan” as “the thing called ‘Vulcan’” putatively fail since distinct empty names – distinct in the sense that they are associated with different practices – may have the same form).

But Sainsbury’s own view falls prey to his own dilemma. Take the problem of syntactically identical but semantically distinct empty names first. On Sainsbury’s own view referential axioms of type (12) would themselves not distinguish between two names “Vulcan” introduced into language by different descriptions. Is the task thus supposed to be assigned to the name-using traditions themselves? That would just push the problem one step back, for unless descriptive associations determine the meaning of the name, how do we individuate the correct name-using traditions? If the name is empty, it must be the descriptions or descriptive intentions of the producers that fix its meaning, and these must be robust enough to carry the identity conditions for the ensuing name-using tradition, since the reference-conditional axioms themselves are insufficient to distinguish two empty names. So Sainsbury’s dilemma applies to his own view. Nor, for that matter, can the axioms distinguish syntactically identical but semantically distinct *descriptive* but non-empty names. Compare ordinary names. In (12) the name on the right hand side is *used*. Using a name involves engaging with a name using practice, and this practice sustains the link to the referent itself. Syntactically identical referential names can be distinguished in virtue of a name-using tradition extending to the referent. But when the tradition terminates in a description, or descriptive intention, then surely this description – not the object itself – is responsible for the name having the meaning it has (and assigning a different meaning to different names such as “Julius” or “Vulcan”).

The main problem is to account for the mechanism by which subsequent speakers (“consumers”) are initiated into a name-using practice terminating in nothing. Sainsbury relies on the speaker who is initiated into a practice having intentions such as

(13) $\forall x$ (if the uses of n I am encountering refer to x , then I will use n to refer to x)

[*ibid.*, p.113]

But (13) does not extend to empty names. What kind of intention would the intention described in the consequent have to be, given that the antecedent is false for any x if n is empty? What, for instance, prevents a speaker who uses “Vulcan” interchangeably with “Pegasus” from satisfying (13)? (13) fails to determine which, if any, name-using tradition an occurrence of “Vulcan” will be part of. (13) does not apply in the case of empty or descriptive names, where the descriptive intention or character must determine name-using tradition. That is, for ordinary names, Sainsbury claims that all that is needed for successful initiation is the disposition on the part of the speaker to use the name with a speaker referent that aligns with the semantic referent (*ibid.* p. 116), but this will obviously not suffice when there is no referent.

Sainsbury also says, surprisingly, that if there is no referent for a name, then there is no issue of failing to be in accord (p. 122).¹⁰⁷ This admission, however, is tantamount to

¹⁰⁷ “The account of baptism and transmission does not depend upon a name-using practice having a referent. To the extent that the concept of a referent entered into the necessary intentions of users, what was required was no divergence of referent. This condition is typically met when someone is initiated successfully into an empty name-using practice: there is no divergence between the neophyte’s intended referent and the semantic referent of the uses by old hands” [p. 122]; “To say that my practice is distinct because it involves different information is to make information figure in individuating empty but not non-empty practices, introducing a multiformity antithetical to RWR” [p. 123]; and, perhaps most strikingly:

settling for the first horn of his own dilemma. If the meaning of an empty name is to be constant between speakers in a linguistic community, which is what he explicitly sets out to achieve, then the description has to be available as well. In that case, however, the motivation for his own account is in danger of being undermined. The obvious move would be to include the descriptive content as part of the metalinguistic referential axiom for empty names. The upshot would be a view reminiscent of our variable view, for if descriptive and empty names are governed by descriptive axioms, the unified theory Sainsbury strives for must be given up. But in that case there seems to be little motivation for adopting reference-conditional axioms – Sainsbury’s main invention – for ordinary names either, instead of unconditional axioms of the kind “*n* refers to *n*.”

If we treat names as bound variables then Sainsbury’s dilemma loses its bite. The worry that a description required to master the application conditions for a descriptive or empty name is too sophisticated to attribute to speakers is unfounded. To know the meaning of “Vulcan” the description “the planet between Mercury and the Sun” would do the trick, but so would the extensionally equivalent “some planet Leverrier was talking about” or “the thing called ‘Vulcan’ by people in my community.” Insofar as we don’t treat empty or descriptive names as *synonymous* with the descriptive clause used to introduce it, meaning constancy across a population of language users does not require that the descriptions the users have to know to master the name have to be the same. And if “Vulcan” ended up being associated with a wholly different (i.e. non-equivalent) description, then there would be little

“Anthropologists worry about the individuation of myths, legends, and fables. How come we find Atlantis-legends on both sides of the Atlantic? It might be a coincidence; that is, there might be two distinct name-using practices. Or it might be that there is a single legend which migrated. Even in empty cases, standard considerations of propagation in terms of initiation and continued participation are what matter” [p.123-24]. One would like to see how Sainsbury could account for the two competing hypotheses about the identity of the Atlantis myth could be explicated considering that no descriptive information can determine the referential conditions for “Atlantis” (especially if the two occurrences themselves are syntactically distinct).

point in requiring that it should do the same truth-conditional work. On our view such variation is not a problem, and it is intuitively plausible. Suppose Jack associates “Vulcan” with “the planet between Jupiter and Saturn.” Jack’s use of “Vulcan” would be objectionable because he is using the name in a non-standard manner. But his manner of using the name is non-standard exactly because he is not associating it with the description “Vulcan” is standardly associated with, not because he is talking about the wrong thing or referent – there is none (compare what we would say if he used “Augustus” to refer to Tiberius; then Jack would in a clear sense be talking about the wrong person; semantically, at least, Jack would, in virtue of being part of a particular name-using tradition, refer to Augustus).

7.4. Application I: A Variable-Name Solution to Empty Names

The main feature of our solution is straightforward. Empty names are introduced by descriptions, or at least presuppose such descriptions. If “Vulcan” is a variable bound by “the planet between Mercury and the Sun,” we will not only obtain the correct truth-values for sentences containing the name but there would be little mystery concerning their truth conditions. Since there is no assignment satisfying “Vulcan is a planet” the sentence is false. It is nevertheless eminently meaningful, since there are worlds of evaluation in which the variable “Vulcan” is satisfied. There are worlds in which the antecedent “planet between Mercury and the Sun” is satisfied.

“Vulcan could have existed” has a false reading (though as we see in 7.4.3 it has a true reading as well), since there is no satisfier of “planet between Mercury and the Sun” that is a possibly existing individual. On the other hand, “if there had been a planet between

Mercury and the Sun it would have been Vulcan” seems to have a (not only vacuously) true reading. We obtain this reading by evaluating “Vulcan is a planet” including its binding phrase, i.e. “the x_{Vulcan} : planet between Mercury and the Sun (x_{Vulcan}): x_{Vulcan} is a planet,” at a different world; that is, by moving “Vulcan is a planet” along the vertical axis in the matrix. Since there is a world relative to which both the antecedent and consequent is true, the conditional is true as well. In 7.4.1 I embed our pseudo-two-dimensionalist interpretation in discourse semantics – although sentences such as “Vulcan does not exist” and “Vulcan is not a planet” comes out as true (as long as “Vulcan” is semantically rather than syntactically bound), we need to account for why such sentences might appear problematic – while semantically straightforward, they have some striking pragmatic effects. In 7.4.3 I distinguish some important semantically relevant differences between discourse involving “Vulcan” today and such discourse at the time of Leverrier.

7.4.1 Negative Existentials

To get a better grasp of how we can properly interpret negative existentials, we may use a framework for discourse interpretation such as UIF (Chapter 5). If a speaker denies the existence of Vulcan, she is denying the binding presupposition while relying on that presupposition for the meaning of “Vulcan,” and something along the lines of UIF is required to avoid existence denials being contradictory. Mapping our pseudo-two-dimensionalism into a discourse semantics is fairly straightforward. “Evaluating the sentence involving ‘Vulcan’, including the binder, at a different world” amounts, simply, to shifting to a different world in the context set. If “a planet between Mercury and the Sun” is a presupposition for or explicitly introduced antecedent in a discourse, it is interpreted as true

in all worlds in the context set (the worlds of evaluation are different from the worlds of the context set, and we may allow the actual world to be no element of the context set), but since “Vulcan” is (semantically) *bound* by this presupposition, different worlds in the context set yield different satisfiers. “Vulcan is difficult to discover” will then be contingently true relative to context. In UIF, saying that “Vulcan” is “bound” amounts to saying that it is a member of the *Topic* – the class of planets between Mercury and the Sun – in each remaining world. “Vulcan” is *rigid* relative to these worlds but not referential, since what satisfies “Vulcan” varies from world to world.

If we take the actual world to be the *world of evaluation* for any sentence in the discourse, “there is a planet between Mercury and the Sun” is false, and the resulting *Topic* – or set of assignments satisfying the sentence in question (Chapter 5) – is the empty set. A subsequent sentence “Vulcan is a planet” will hence be false as well. “Vulcan is not a planet” is true, however, and is the negation of “Vulcan is a planet” (we need not complicate matters with a distinction between choice negation and exclusion negation). The sentence is easily interpretable, however. Since there are worlds relative to which the antecedent is true, there are worlds relative to which *Topic* is *not* the empty set, and in which “Vulcan is a planet” is true (indeed, it is true in all worlds in which the antecedent is true, though not *necessarily* true relative to those worlds).

Similarly for “exists.” We need to treat existence as a predicate, e.g. $\lambda x[\exists x(x = x)]$ – otherwise any denials of existence would be contradictory – and $\lambda x[\exists x(x = x)](\text{Vulcan})$ is true at a world w iff Vulcan is in the domain associated with w . “Vulcan exists” is false, and “Vulcan does not exist” is correspondingly true. The discourse fragment, including the antecedent, is eminently meaningful, however, since there are worlds in which the antecedent

is true and hence relative to which “Vulcan exists” is true as well. Of course, to interpret negative existentials, we must take a step back from the simplified representation of sentences containing descriptive names used thus far in this chapter. It is crucial that “Vulcan” is *semantically*, not syntactically, bound by its antecedent – “Vulcan does not exist” should not be represented as the contradictory “([the x_{vulcan} : x_{vulcan} is a planet ...] $\neg(x_{vulcan}$ is a planet),” but as “it is not the case that x_{vulcan} is a planet,” where the interpretation of the syntactically free variable “ x_{vulcan} ” is determined relative to the set of individuals satisfying the antecedent. Someone who utters “Vulcan does not exist” is not asserting that it *is* a planet that doesn’t exist, but denying that the presupposition is true while using the name in accordance with that presupposition. This effect is easily achieved as long as “Vulcan” is not *syntactically* bound by the presupposition.

In UIF (and other dynamic systems) anaphoric chains are embedded in the discourse structure, and are necessary to interpret “Vulcan,” which is hence semantically but not syntactically bound by the antecedent quantifier phrase “the planet between Mercury and the Sun.” There is no contradiction involved in uttering “Vulcan is not a planet.” From the point of view of discourse participants, who assume that Vulcan exists, the utterance may appear false. Relative to the actual world, however, since Vulcan is, in fact, not a planet an utterance of “Vulcan is not a planet” will be true, even though it denies its antecedent, and even though the antecedent is required to interpret “Vulcan is not a planet.” (Note also that if “Vulcan” is treated as an anaphoric name in the sense of UIF, one that may arguably project its own binder in virtue of its character, e.g. “thing called ‘Vulcan.’”) Insofar as negation is a sentence operator, sentences containing “Vulcan” are false, not truth-valueless.

Truth-conditionally, there is thus nothing problematic about either existence-claims or planethood-attributions involving Vulcan. We distinguish the context set from the worlds of evaluation. The context set of a discourse in which participants believe Vulcan exists does not contain the actual world. Nevertheless, each world of the context set may be mapped to a function from world-time pairs to truth-values (where the circumstances of evaluation include but are not restricted to the worlds of the context set). So for instance, let $w_{a1} \dots, w_{an}$ be the worlds relative to which there is a planet, $\alpha_1, \dots, \alpha_n$ between Mercury and the Sun. For discourse participants “Vulcan is a planet” is understood relative to $w_{a1} \dots, w_{an}$ where $\alpha_1, \dots, \alpha_n$ are the satisfiers (“Vulcan” does not *refer* to any of $\alpha_1, \dots, \alpha_n$, however, for which item satisfies “Vulcan” varies over the worlds of the context set). The interpretation of “Vulcan is a planet” is a mapping of $\langle w_{a1}, \alpha_1 \rangle, \dots, \langle w_{an}, \alpha_n \rangle$ onto the world-domain pairs of the circumstances of evaluation. Since $@w \notin \{w_{a1} \dots, w_{an}\}$, “Vulcan is a planet” is false (at the actual world), though not *necessarily* false. “Vulcan is not a planet,” on the other hand, is false at $\langle w_{a1}, \alpha_1 \rangle, \dots, \langle w_{an}, \alpha_n \rangle$ (i.e. at all worlds in the context set), but since $@w \notin \{w_{a1} \dots, w_{an}\}$, it is true at the actual world. If $@w$ is treated as the *context*, “there is a planet between Mercury and the Sun” is straightforwardly false.

Though not *necessarily* false, note that since any sentence anaphoric on the antecedent, such as “Vulcan is a planet” is interpreted as $(Vulcan = \|\lambda x_T. x\|) \in I(Planet)_{@}$, where T is the set of planets between Mercury and the Sun, “ $\langle P(Vulcan) \rangle$ ” will in fact be false as well – since Vulcan does not exist, the denotation of “Vulcan,” nothing, could not have been a planet. However, relative to some worlds in the context set worlds in which Vulcan could have been a planet are indeed accessible. I am tempted to say that this captures

a sense in which the planethood of Vulcan is an epistemic but not metaphysical possibility – relative to the worlds in which the presupposition that provides the means for interpreting claims involving Vulcan is satisfied, Vulcan’s planethood is possible, but relative to the actual world Vulcan could not have been a planet. I leave it for a later time to discuss whether a more general theory of the distinction between epistemic and metaphysical possibility could be developed along these lines.

The results accord well with an interpretation in light of the matrix above. “Vulcan is a planet” expresses a complete semantic unit, with complete truth conditions. Changing the world of evaluation for the antecedent, however, will change the truth conditions of the sentence. Given the actual world as world of evaluation, “Vulcan is not a planet” has no true reading, even relative to the presupposition. The reason the sentence may seem infelicitous – or at least hard to interpret properly – in a context in which discourse participants assume that there is a planet between Mercury and the Sun, is that their context set or common ground is inconsistent with the sentence. From the point of view of the discourse participants, who assume that Vulcan exists, the utterance of “Vulcan does not exist” will in fact come across as *a priori* false relative to the presupposition that Vulcan is a planet between Mercury and the Sun. (Since the antecedent is itself contingent, however, “Vulcan is not a planet” would nevertheless be only contingently false from this perspective.) Since the context set contain only the worlds in which there is a planet between Mercury and the Sun any contribution of “Vulcan is not a planet” or “Vulcan does not exist” will return the null set as the context update; i.e. it eliminates the complete context set. That does not mean that such sentences are false, or contradictory, or semantically problematic. In a discourse, such utterances conflict profoundly with the assumptions at work. The *semantics* of negative

existentials are straightforward, though the potentially catastrophic effects of uttering such sentences require substantial pragmatic repair. How we should proceed in cases like these (of catastrophic presupposition failure) is moot, and I will not attempt to solve the question here – any theory of discourse would need to account for the phenomenon, and the view defended here seems to be compatible with a variety of approaches.¹⁰⁸

7.4.2. Vulcan the Hypothetical Planet

Although the context just described might reflect discussions when the existence of Vulcan was still an open question, when “Vulcan” is used *at present*, the presupposed antecedent is presumably not “there is a planet between Mercury and the Sun.” We are generally aware that this is false, and uttering “Vulcan does not exist” should not engender any deep revisions of our presuppositions or require pragmatic backtracking. The presupposition cannot be one that is explicitly denying existence either. “There is no planet between Mercury and the Sun called ‘Vulcan’” does not license uses of “Vulcan” in discourse, any more than it would license discourse anaphora. *Hypotheticals* license discourse anaphora, however. “Jack allegedly had a donkey, and Harry allegedly beat it,” seems fine even if we know that Jack, in fact, did not have a donkey. Hence empty names, even “Vulcan” at present, should not be governed by an explicitly negative presupposition, but, perhaps, by “the alleged planet between Mercury and the Sun” (or perhaps a hyperintensional operator, though such cases

¹⁰⁸ Diagonalization, as suggested by Stalnaker to deal with defective contexts (1978, 1998), is not itself sufficient. There seems to be no suitable uncertainty about the proposition is expressed by “Vulcan does not exist” in this case since the speaker, *ex hypothesi*, uses “Vulcan” entirely in line with the presupposition governing the discourse – at least all speakers agree on the presupposition “thing called ‘Vulcan,’” and that agreement seems to be sufficient to block diagonalization in the manner suggested by Stalnaker. The problem would not arise if we were allowed e.g. to entertain a strictly larger context set than the set consisting of worlds in which the presupposition is true, or if participants were allowed to entertain different context sets simultaneously (as in Beaver, 1992).

involve complications that we discuss below). Given the antecedent “the x_{Vulcan} : alleged (x_{Vulcan} is a planet and x_{Vulcan} is between Mercury and the Sun” the sentence “Vulcan is not a planet” would come out as true on its most natural reading.

Suppose the antecedent for “Vulcan” is “the alleged planet between Mercury and the Sun” or – for simplicity – “a planet that could have been located between Mercury and the Sun.” The antecedent is satisfied with respect to the actual world if there is a world w accessible from the actual world in which there is a planet between Mercury and the Sun. “Vulcan” is then interpreted as the x such that for any world w , if x is a planet between Mercury and the Sun in w , then $x = Vulcan$ in w (“Vulcan” is nevertheless rigid, since with respect to a world w^* accessible from w , for any y such that $Vulcan \neq y$ in w , $Vulcan \neq y$ in w^* even though y is the planet between Mercury and the Sun in w^* and Vulcan is not).

There are some notable constraints on uses of anaphora in hypothetical contexts (Roberts, 2005). Such contexts require repetition of the subjunctive operator in the anaphoric sentence. Operator scope is confined at most to the sentence in which it is uttered, and information introduced within the scope of an operator is only accessible within that scope. “It is possible that Jack had a donkey. Harry beat it,” does not have the reading “it is possible that Jack had a donkey, and it is possible that Harry beat it.” Insofar as the antecedent for “Vulcan” is hypothetical, subsequent sentences containing “Vulcan” will usually have to be *explicitly* subjunctive (“Vulcan would have been F ”) as well.

Anaphoric links can nevertheless be sustained. The domain restrictions needed to interpret anaphora occur in possible worlds of evaluation, not worlds of the context set. The antecedent makes a restriction relative to the domain of counterfactual worlds accessible from the context set, and subsequent claims such as “Vulcan would have been F ” are

evaluated relative to these counterfactual domain restrictions. They would nevertheless be true, if true in a possible world accessible from the worlds of the context set, i.e. at a different line on the vertical axis of the matrix above. An explanation for why the intensional operator must be repeated is presumably that the effects on discourse by a sentence utterance must be effects on the context set itself, not merely on counterfactual worlds accessible from the context set. “Vulcan could have been *F*” rules out worlds in the context set from which “Vulcan is *F*”-worlds are inaccessible or sufficiently distant, so discourse contributions must always be relativized to the context set – any effects in the counterfactual worlds, then, are caused indirectly, through changes in the context set. (Note that the counterfactual worlds may themselves not be in the context set.)

Furthermore, and importantly, under the presupposition that Vulcan is the *alleged* planet between Mercury and the Sun, “Vulcan does not exist” is still felicitous, and must – since it involves no repeat of the intensional operator signaling evaluation relative to the hypothetical context – be interpreted as a claim about the actual world (to the effect that the actual world is not among the worlds in which there is a planet between Mercury and the Sun). The *Topic* relative to which “Vulcan” is interpreted is hypothetical, since the antecedent involves a hypothetical context. Hence, the domain restriction relative to which “Vulcan” is interpreted is the domain of a counterfactual world. “Vulcan does not exist” is thus true relative to the actual world iff the individual that is a planet between Mercury and the Sun in the counterfactual world does not exist in the actual world.¹⁰⁹

¹⁰⁹ In fact, the situation is slightly more complicated and gives rise to a somewhat surprising puzzle. Let “Vulcan does not exist” be $\neg @Ev$ (that is, “Vulcan does not *actually* exist” – a similar effect can be obtained by other means; the exact formulation does not matter here). Since “*v*,” however, is an identity function from *Topic*, the set of things that are (the) planet between Mercury and the Sun in the non-actual world *w*, “*v*” is a rigid designator relative to its assignment in *w*. $\neg @Ev$ is true iff the object *y* such that *y* is

As with other descriptive names, the semantic contribution of an utterance of an empty name depends on context and available presuppositions. Speakers are talking about the same thing (in any relevant sense) – their contributions to a discourse are properly coordinated – iff they associate the name with (extensionally) equivalent binding phrases. If empty names are descriptive names, however, there is no obvious reason why the interpretation of sentences containing “Vulcan” should be the same across contexts unless the associated descriptions were preserved. As opposed to ordinary names, empty names have no anchor in referents to keep their meaning stable. Since we lift this requirement we avoid the difficult metalinguistic problem of ensuring cross-contextual constancy of interpretation. Using “Vulcan” in contexts governed by different (presupposed or explicit) antecedents may, at least in one sense, amount to talking about slightly different things.¹¹⁰ Each antecedent would produce a meaningful discourse with well-defined truth conditions, however. If this means that people using “Vulcan” are sometimes talking a little past each other it is hard to see this as a substantial objection. Interpreting claims about Vulcan expressed by Leverrier, for instance, may require retrieving some parameters from the context of utterance, which is in any case a plausible condition for interpretation.

the planet between Mercury and the Sun in w does not exist in the actual world. In that case, however, it is in fact not immediately clear whether $\neg @Ev$ is true or false. To be true it would require that y does not exist in $@w$, but although there is no planet between Mercury and the Sun in $@w$, insofar as *being a planet* is not a necessary property, it might in principle be the case that $\neg @Ev$ is false. The individual that is a planet between Mercury and the Sun could, in fact, in principle exist in a different guise in the actual world (it could of course not be a planet between Mercury and the Sun). If this consequence is unpalatable, it may be possible to avoid it for instance through appeals to a closeness relation for counterfactuals (everything else being equal, a world in which some actual object served the role of Vulcan might be more distant than a world in which a new object served the role).

¹¹⁰ Though there is a sense in which the discussion is about the same thing as well; as long as an antecedent is empty, an occurrence of “Vulcan is F ” will receive the correct truth-conditions regardless of whether the antecedents in use are (intensionally) distinct.

7.4.3. Empty Names, Gappy Propositions, and the Role of Descriptions

We can hence solve the problem of empty names without invoking metaphysical considerations about (hierarchies of) non-existent or special abstract objects (Parsons, 1980; Salmon, 1987, 1998; Zalta, 1985, 1988, 2003). There is a sense in which our view bears some resemblance to views that interpret sentences containing empty names as expressing gappy propositions (Taylor, 2000, Adams & Dietrich, 2004, Wyatt, 2007). On the view we have defended sentences containing empty names are open and not truth-evaluable in isolation from context, and there is nothing in particular to which such names refer. But instead of semantically expressing a gap (and rather *conveying* relevant information to discourse participants in terms of associated descriptions), we provide semantic contents for empty names in virtue of available presuppositions.

Gappy propositions views are problematic for several reasons. First, speakers are often unaware of whether a name is empty, so speakers using “*n* is *F*” will be unable to determine whether they express a proposition or merely conveys one. Competent and rational speakers are generally able to recognize initial confusions over the semantics–pragmatics distinction, however, at least after sufficient training (Braun, 1993, p.460; Reimer, 2001, p.25), and should thus be able to recover the semantic content, the gappy proposition, upon educated reflection. But recovering this proposition is impossible if one does not know that the name in question is empty. Furthermore, since conversational implicatures should be cancellable, if the illusion of truth-evaluability of sentences containing empty names is due to a semantics–pragmatics confusion, the illusion should be removed by cancelling the descriptive propositions that are pragmatically implicated; but it is not (Green, 2007; Piccinini & Scott, 2010). A different line is taken by Taylor (2000), who explains the

contribution of empty names in virtue of a special type of pragmatic *intrusion* (not implicature) drawing on associated descriptions that comes into effect when context is unable to supply a referent. As Everett (2003) points out, the approach seems *ad hoc*. If speakers themselves are unaware of whether the name they use is empty or not, why is the intrusion not in effect in uses of *non-empty* names?

Our variable-interpretation is not susceptible to these worries. Names can work in different ways in different contexts, but the difference is of course not due to pragmatic effects that come into play because the name is empty. Instead, the difference is a consequence of the semantic role of the name in that particular context. I admit that speakers may be unaware the semantic differences between the different name occurrences, but there are good reasons for this (see Chapter 6). Regardless of whether a name occurrence is anaphoric or deictic it is a singular term and not, for instance, a covert description. The difference concerns rather the role of contextually available associated presuppositions. In principle it should be possible to determine, for the speaker, whether a given use of a name is anaphoric or deictic, but in practice it will often be impossible. The reason is an important effect of character and how character connects to name-using traditions. The character “thing called *n*” invokes a tradition, and implies that what Evans would call “consumer users” of names use names *deferentially* – that is, defer to previous users to supply the meaning.

With regard to a sentence containing “Caesar” I defer the meaning of the name to earlier users including, through chains of deference, the original antecedent of the name, which would be either a demonstrative or a description. If the original antecedent is a demonstrative, then deference is a matter of co-referring with that antecedent. Whether it is or not is of course outside of the control of a present (consumer) name-user. Thus, there is a

certain degree of uncertainty about meaning associated with name uses. But this is not a particularly controversial claim. If the meaning of a name is the referent, then since we cannot know for certain whether the name is empty, Millians would also be committed to such uncertainty. Now, if the name was originally introduced by a description, we *could* say that the present use is anaphoric on this description. However, in Chapter 6 I argued that the constitutive reliance on the original description would prevent a name from initiating a proper name-using tradition. In these cases, then, the governing antecedent for a particular name occurrence is the description presupposed in that context of use (whatever it may be). In either case there is no particular reason to think that speakers should be aware of whether names are descriptive or not, and – since names are singular terms – it would not matter to how we use it. Of course, if speakers had immediate access to whether their sentences expressed *singular* or *general propositions* things would be difficult, but the fact that we defer meaning in the case of proper names means precisely that we do not have transparent access to the proposition, and as we argued at length in Chapters 1–3 we cannot read off the grammatical category to which a name belongs whether a proposition is singular or not. And as opposed to gappy propositions view, the contribution of a name is not a matter of mechanisms that come into play *because* the proposition expressed is defective, and is not a matter of invoking pragmatic mechanisms or implicatures that ought to be cancellable.

On Braun's (1993, 2003) version of a gappy propositions view atomic gappy propositions are not meaningless but false, whereas complex propositions, including belief ascriptions and negations, may be true. Instead of invoking implicatures to explain the information speakers associate with sentences containing empty names, Braun appeals to *ways of grasping* gappy propositions, psychological states under which gappy propositions

can be believed. “Vulcan exists,” “Santa exists,” and “Father Christmas exists” all have the same *semantic* content, however. But intuitively, the latter two are more intimately connected. To account for this intuition, Braun invokes propositional guises under which these propositions may be believed. But that gives us a dilemma parallel to the metalinguistic worry associated with e.g. Sainsbury’s view: I) On the one hand, if such guises – psychological entities – are so fine-grained and private that they may vary from context to context or utterance to utterance, it is difficult to see how they can be invoked to account for intuitions concerning sameness of meaning between utterances. II) If they *can* account for these intuitions in a systematic manner, they appear to be much more tightly connected to the language in which these sentences are uttered (and semantics) than to psychology, and would thus look much more like Fregean senses (Mousavian, 2008), which is exactly what Braun wants to avoid.

On our variable-view sameness or difference between utterances is a matter of semantics, not psychology. Though presuppositions associated with a name may in principle vary across contexts (though not with psychology), it is no surprise that they will generally be the same. Presuppositions are nevertheless not Fregean senses, and do not constitute the meaning of the name. In short, gappy propositions views need to explain how the contributions of “Santa” can remain the same across uses, and how the contributions of “Santa” and “Vulcan” differ, even if the explanation is relegated to the level of pragmatic factors or implicatures. To account for this one must invoke associated descriptions and some form of systematicity in virtue of which the contributions can remain (more or less) constant. Such views are hence no more, and likely less, parsimonious than the variable view, and have additional counterintuitive results.

Meinongian views encounter metalinguistic obstacles as well. The context of utterance a name determines the reference of that occurrence in virtue of its character, the associated name-using tradition, and other contextually salient information. In the case of “Aristotle,” the invoked name-using tradition supplies the anchoring individual, and it does not matter what other information speakers may associate with Aristotle; Aristotle himself ensures a constant interpretation across contexts. But if “Vulcan” refers to an abstract or Meinongian object, why expect such constancy of meaning? On both Meinongian views and abstract-object views the individuation of a correct referent is problematic. The correct non-existent or abstract referent for an occurrence of “Vulcan” is an object that encodes the property of being a planet between Mercury and the Sun, but there is a potentially infinite range of non-existent objects that encode this property.

The problem of *individuation* for Meinongian theories (which non-existent planet between Mercury and the Sun does “Vulcan” refer to) is well-known (see Parsons, 1980; Zalta, 2003; Sainsbury, 2007 for discussions), and is certainly not made easier by the fact that the properties Vulcan is supposed to encode include relational properties to concrete existing objects (how does a non-existent or abstract object encode *being between Mercury and the Sun*). Even if individuation conditions for Vulcan could be defined, however, it is unclear how we guarantee *sameness of reference* across contexts. After all, a referent is determined in part by the context. Whereas ordinary referential names determine referents partially in virtue of links to the introducing demonstrative, or by causal links to the referent, neither mechanism is available in the case of “Vulcan.” So even if we could individuate a non-existent reference of “Vulcan” in a context we seem to get an additional problem: How do we ensure that the same referent is assigned in every context in which “Vulcan” is used.

Indeed, it is unclear that anything could even in principle ensure that different occurrences of “Vulcan” received the same referent – in the case of ordinary referential names the range of potential candidates is relatively limited, but the range of suitable non-existent objects is not.

What the Meinongian should *not* say is that it is associated descriptions that are supposed to individuate a referent in a context. Such descriptions change over time (think for instance of the changes in descriptions associated with Santa), and in the case of a slight shift in associated descriptions at the meta-linguistic level – maybe so fine-grained that there is no intensional difference between them – there is no obvious way of preventing different occurrences of “Vulcan” from being assigned different Meinongian referents. Views according to which empty names pick out abstract objects face the same kind of worry. Such objects are easy – too easy – to create, so given that different information is associated with Vulcan in different contexts, why think that the different occurrences refer to the same rather than different abstract objects? However, as long as we assume that the propositions expressed are singular – which is the motivation between Meinongian and abstract-object approaches – it is required that the referent is the same for conversation to work.

Thus, letting empty names refer to non-existent or abstract objects does not make it easier to ensure that different occurrences of “Vulcan” receive the same interpretation than on a variable view. Insofar as “Vulcan is *F*” expresses a singular proposition and the descriptions do not occur in the semantic content of sentences, interpretation across contexts may in fact become impossible. Of course, the worry is avoided by requiring the associated descriptions to remain constant across contexts to ensure that the same referent is individuated, but this requirement rather defeats the purpose of Russellian accounts of such names and undermines the unification of the view in terms of cognitive resources required to

understand empty and non-empty names. Meinongian and abstract object views achieve a certain semantically unified view of name occurrences, but as we have argued at length this is not a particularly significant advantage over our pluralism, which provides a unified categorization of singular terms, which Russellian accounts do not.

On our variable view the associated descriptions are part of semantics; they range over the (concrete, at least in the case of purported planets) objects existing at the world of evaluation, a far more restricted domain than Meinongian and abstract objects, and interpreting a sentence containing “Vulcan” requires no more than relying on a presupposition roughly equivalent to the presupposition used in the sentence’s context of utterance. One main benefit over Meinongian views, then, is that we do not require descriptions to be constant across contexts in order to ensure cross-contextual interpretation of sentences – extensional equivalence is sufficient. Capturing a whole discourse, including presuppositions, from the point of view of a different discourse might furthermore be viewed as a matter of degree – whereas reference isomorphism is not a matter of degree, intensional isomorphism may be. Though I leave the idea of degrees of intensional isomorphism somewhat open, the idea is that whereas the contents of the names *n* and *m* are either identical or different, the content of “bearded, jolly fellow who brings presents on Christmas *Day*” is *closer in meaning* to “bearded, jolly fellow who brings presents on Christmas *Eve*” than to “planet between Mercury and the Sun.”¹¹¹ In short, Meinongians must at the very least ensure that the descriptions associated with “Vulcan” remain fully stable, and this implausible requirement may not even be enough.

¹¹¹ A full account of degrees of intensional isomorphism would require a detailed account of the derivational background of the structured propositions or intensions.

Since empty names are not referential, our view requires nothing similar – descriptive constancy across contexts is non-essential.

7.5 Application II: Frege’s Puzzle

Applying our framework to solve Frege’s puzzle is perhaps less intuitively obvious but is worth considering.¹¹² The puzzle of why “Hesperus = Phosphorus” is informative, whereas “Hesperus = Hesperus” is not,¹¹³ and how Hesperus = Phosphorus could be an empirical discovery, concerns the intersection between semantics and epistemology or psychology, and at least Russellian theories of propositions may in principle deny that the puzzle should strictly speaking receive a semantic solution at all (though pragmatic factors can be used to explain why statements involving “Hesperus” and “Phosphorus” may have different effects on a conversation; see Salmon, 1986; Stalnaker, 1998; Soames, 2002).¹¹⁴ On Fregean views,

¹¹² Cumming (2008) develops a sophisticated view according to which names in belief contexts can be treated as (open) variables that are assignment sensitive and can give rise to *de dicto* readings of names in certain situations, so an idea along these lines would not be entirely new.

¹¹³ Note, as Salmon (1986) points out, that “Hesperus = Hesperus” expresses a proposition which is arguably different from the one expressed by “Hesperus is self-identical,” which can be represented by $\lambda(x)[x = \textit{Hesperus}](\textit{Hesperus})$ and $\lambda(x)[x = x](\textit{Hesperus})$, respectively. The propositions are equivalent, but the properties involved are very different.

¹¹⁴ The pragmatic two-dimensionalist solution of Stalnaker is rather different than the one I will suggest here. On Stalnaker’s externalist picture, since names are referential and referents themselves (not representations of such referents) are parts of the worlds of the context set, discourse participants may exhibit confusion or error over how to update the context as discourse evolves. Since names are directly referential and first-order identities are necessarily true, there will be certain pieces of informative information (such as “Hesperus = Phosphorus”) that appear not to issue such updates (which for Stalnaker can only be a matter of deleting worlds). Stalnaker accounts for such confusion by pragmatic factors – it is interpreted as a confusion about what proposition is, in fact, expressed. For Stalnaker, an utterance *u* updates a context by intersecting it with two effects: the proposition that *u* has been made (“the Commonplace Effect”) and the proposition *p* expressed by *u* (“the Essential Effect”). If there is uncertainty in the context about the meaning of *u*, then the Essential Effect cannot take place, which discharges the pragmatic mechanism of *diagonalization* as follows: first, a propositional concept is generated from *u* (the propositional function is a function from worlds to propositions, which, since propositions are (functions from) worlds (to truth-values) in Stalnaker’s system, amounts to an accessibility relation on worlds). This propositional concept maps each world of the context set onto *p* (what *u* in fact expresses); but rather than

on which cognitive significance and semantic content are more intimately related, co-referential names are not interchangeable in attitudinal contexts – a name denotes its customary sense rather than its reference in such contexts. However, it must arguably be possible to disregard this effect; (14) has a true reading as well.

(14) The ancient Babylonians believed that Venus was a star

Kripke's puzzle (Kripke, 1979) concerning someone who believes that *Paderewski has musical talent* and that *Paderewski does not have musical talent* by failing to recognize that Paderewski the prime minister and Paderewski the pianist is the same person, is a challenge even to Fregean solutions. Since the sense of an expression is its contribution to a proposition, and since communication proceeds by way of grasping such senses, there is a limit to how fine-grained or sensitive to idiosyncratic cognitive states senses can be. There are also Russellian approaches (Richard, 1990; Crimmins, 1992) that adopt something resembling Frege's technique by letting the complement clauses of "that"-phrases be interpreted differently from ordinary occurrences by treating "that" itself as a context-shifter, as demonstrating the proposition-under-a-guise, or perhaps by invoking various contextually determined features.

updating the context set with p , diagonalization means updating the context set with information from the propositional concept. We do this by determining for each world w in the context set, whether the proposition that the propositional concept projects at w , includes w , and cut away the worlds in which it does not. What remains is then the worlds w relative to which the propositional concept projects w , and this set of worlds is the diagonal proposition. Uncertainty or confusion about for instance first-order identities, then, can be accounted for by diagonalization, and it is a pragmatic rather than semantic form of two-dimensionality. In the take on Frege's puzzle I provide below, I do not attempt to model the *discourse* updates engendered by discourse-relative confusions that result in Stalnaker's effects. On the other hand, Stalnaker's account of such confusion does not extend entirely smoothly to belief ascriptions.

Given that Frege's puzzle appears to be a puzzle only if we invoke concerns that do not obviously belong to the domain of semantics, it is not entirely clear that our account even *ought* to solve it. The considerations made in 7.5 are thus somewhat tentative; the proposed solution may also exhibit a whiff of *ad hoc*ness. As opposed to "Vulcan" or "Julius," "Hesperus" is not particularly plausibly interpreted as a descriptive name at all. To obtain a solution we must hence grant that there is a sense in which hyperintensional operators (or "that") can work as context-shifters or context-updaters that take the agent's belief context into consideration. In 7.5.1 I explain the basic operations, in 7.5.2 some technical resources and alternative accounts; in 7.5.3 I discuss Kripke's puzzle, and in 7.5.4 Geach sentences.

7.5.1. A Solution to Frege's Puzzle: The Basic Moves

There may, however, be good reasons to treat hyperintensional operators as context shifters. Although (14) has a true reading since the Babylonians believed, of the object Venus, that this object was a star, it may also have a false reading. If it does, it must be because the embedded clause is somehow sensitive to the presuppositions and context of the ancient Babylonians and whether they would or would not for instance assent to the sentence "Venus is a star." Thus, the false reading of (14) emerges by embedding parameters of the ancient Babylonians' context, or at least a suspension of our context in favor of a (possibly generalized) interpretation of their context, into *our* context of utterance for (14). (Similarly, to interpret occurrences of "Vulcan" uttered by Leverrier, we may have to recover his presuppositions, at least to the extent that the presuppositions we use are extensionally equivalent to at least some of Leverrier's presuppositions.)

Such shifts in contextual parameters for expressions within hyperintensional contexts might also license a special treatment of names. For grammatical reasons “Hesperus” is a singular term even in the context of belief ascriptions, and semantically it can contribute nothing but an individual to the propositions expressed by “Jack believes that Hesperus is not bright.” But remember that a name only gets an interpretation in context – it is *occurrences* of names that receive interpretations based on available contextual information. Whereas “Hesperus” is usually used as a constant, there is no barrier, in principle, to treating occurrences of “Hesperus” differently when it occurs embedded in an attitudinal context. We *could* of course continue to treat it as a constant even in this case, but that might prevent us from interpreting Jack in a manner that makes his beliefs intelligible. The best way to make “Jack believes that Hesperus is not bright” intelligible would thus be to interpret “Hesperus” as sensitive to *his* presuppositions, and to do this we must interpret it as a variable. So if there is a sense in which “Jack believes that Phosphorus is bright” and “Jack believes that Hesperus is *not* bright” could be true together – without attributing a contradiction – it must be because what Jack associates with the names is in some respect relevant to the truth conditions of the ascriptions. Thus, if we want to interpret Jack’s beliefs the interpretation must be sensitive to *Jack’s* presuppositions. Note that we do not treat “Hesperus” even in Jack’s belief context as *equivalent* to descriptions; rather, we attempt to make the contributions of sentences containing “Hesperus” sensitive to Jack’s other assumptions about Hesperus (within the context of Jack’s belief) in a manner that allows us to predict what he would or would not assent to. “Jack believes that” is hence a context shifting operator.

“Hesperus is not bright” could be “the x_{Hesperus} : $F(x_{\text{Hesperus}})$: x_{Hesperus} is not bright,” where “ F ” represents some part of Jack’s beliefs about Hesperus. When “ F ” is evaluated at

the actual world, then insofar as “*F*” ascribes some predicate true of Hesperus (such as “thing others call ‘Hesperus’”) Jack’s belief that Hesperus is not bright is false. But the belief would not be *necessarily* false or contradictory, even in combination with his (true) Phosphorus-belief, since whatever actually satisfies *F* (Venus) may fail to be bright with respect to some worlds. “Hesperus” would nevertheless behave rigidly, and this enables us to interpret Jack’s counterfactual beliefs about the individual value of “Hesperus.”

It would furthermore be the case that “Hesperus = Phosphorus” is necessarily true, if true, even in belief contexts. To see this, let “*F*” and “*G*” denote some presupposition associated with Jack’s Hesperus and Phosphorus beliefs, respectively. So Jack believes that “the $x_{Hesperus}$: $Fx_{Hesperus}$: $x_{Hesperus}$ is not bright” and “the $x_{Phosphorus}$: $Gx_{Phosphorus}$: $x_{Phosphorus}$ is bright.” Relative to the actual world Venus is the satisfier of “*F*” and “*G*,” so $x_{Hesperus} = x_{Phosphorus}$. Given the necessity of identity “ $\Box (x_{Hesperus} = x_{Phosphorus})$ ” is true. In the context of Jack’s beliefs, however, “ $\Box (x_{Hesperus} = x_{Phosphorus})$ ” is (15):

$$(15) [\text{the } x_{Hesperus}: Fx_{Hesperus}] [\text{the } x_{Phosphorus}: Gx_{Phosphorus}] \Box (x_{Hesperus} = x_{Phosphorus})$$

(15) is true, but not *necessarily* true, since it is false relative to worlds in which *F* and *G* have different satisfiers. Hence, we may attribute to Jack coherent doubts about the identity of Hesperus and Phosphorus, even though his beliefs are (if you wish) *de re* and he recognizes the rigidity of names and the necessity of identity, in virtue of doubting that “*F*” and “*G*” have the same satisfier.

That is the gist of it. But let us be very careful about what we are trying to achieve. The proposal is *not* a form of descriptivism (instead, it is related to the treatment of belief

contexts in Kaplan, 1968). It is emphatically not the case that the meaning of “Hesperus,” or its contribution to the truth-conditions of belief ascriptions to Jack, is captured by, equivalent to, or determined by the description “the star that shines in the evening,” or even a cluster of descriptions that Jack may associate with Hesperus. Since “Hesperus” is interpreted as a variable its meaning is, if anything, any potential element of the domain over which it ranges. The binding presuppositions do not give the meaning or sense of *Hesperus* but the truth-conditions of “Hesperus”-involving *ascriptions*. We need not even know what Jack’s presuppositions are to interpret the ascription “Jack believes that Hesperus is not bright.” We may use a dummy description *D* to stand in for presuppositions Jack might have, and let *D* bind the occurrence of “Hesperus.” Jack’s belief is false if any part of *D* is *extensionally* equivalent to any definite description actually satisfied by Venus, including “the thing we call ‘Hesperus’” (if it is not it becomes hard to determine what Jack’s belief is about in any case), and the *ascription* is true if *D* is extensionally equivalent to any description Jack associates with Hesperus.

This is sufficient to provide the correct truth conditions for an ascription of the belief that Hesperus is not bright to Jack. We would of course not capture exactly how Jack updates his beliefs when he learns that Hesperus = Phosphorus. But the point is not to capture Jack’s cognitive and epistemic states. These states, in all their glorious details, do not matter to the meaning of “Hesperus,” a singular term, even in belief contexts. The point is rather to model belief ascriptions in a manner that allows the identity statement to be at least potentially informative and doubts about the identity claim to be coherently interpreted. Insofar as “Hesperus is not bright” is “ $x_{Hesperus}$ is not bright,” *any* presupposition that Venus actually satisfies, including “being called ‘Hesperus’” would suffice to avoid attributing

contradictions as long as we allow the presuppositions binding “Hesperus” and “Phosphorus” to come apart in some possible worlds.

7.5.2. The Role of Contexts

Though ordinary names such as “Hesperus” and “Phosphorus” are referential, and substitutable, the idea is that attitudinal contexts are ambiguous and *allow* a reading on which one interprets attitudinal operators as context shifters, where attitudes are special kinds of contexts. A fairly common way of interpreting attitude reports in dynamic systems is to treat the clause embedded under the attitude verb *not* as applying to the context representing the common ground but rather as concerning a different context, one that also represents the content of Jack’s propositional attitude content (Stalnaker, 1987). We can call this the *secondary context* of a discourse, as opposed to the *primary context* that represents the common ground. The secondary context will include possibilities not included in the common ground. In its simplest form one may let every world in the primary context be associated with an accessibility relation R_{jack} . The information state of Jack’s beliefs (relative to the primary context C_1) is the union of the worlds determined by the worlds in C_1 ; i.e., the secondary context $C_2 = \bigcup_{w \in C_1} \{w^* \mid w^* R_{jack} w\}$ (alternatively, the secondary context can be identified with the set of set of worlds standing in R_{jack} to C_1 ; see Dekker & van Rooy, 1998). Attitude reports are interpreted as

$$16. \langle C_1, C_2 \rangle + \text{Jack believes that } p = \langle C_1, C_2 + p \rangle$$

In UIF the union of contexts must be represented in the salience structure. However, even then it would be unclear how to deal with Frege's puzzle, since introducing a secondary context itself does not license a binding of the otherwise deictic name "Hesperus."

In order to achieve the desired interpretation we let "believes that" be an operator that may potentially block semantic binding and which (in virtue of being a context-shifter) requires that expressions are interpreted in isolation from the context in which we make the belief ascription. Instead of letting expressions be interpreted in light of the presuppositions that govern the context of utterance, we want the embedded clause to be sensitive to the agent's presuppositions. We may, of course, treat such contexts as a merging of the context set that represents the agent's belief states (the secondary context) with the context that represents the discourse, but it is strictly speaking not necessary to interpret belief ascriptions that the secondary context reflects the agent's beliefs in any strict sense. The agent's presuppositions are generally inaccessible to the primary context. Unless the presuppositions are made explicit (as in "Jack believes that Hesperus is a star rather than the planet Venus") we will have to use information associated with "Hesperus" in the primary context to define a binder. Fortunately there is an obvious candidate available, namely "the thing called 'Hesperus.'" We can let the characters of "Hesperus" and "Phosphorus" project local presuppositions to capture Jack's perspective.

Although we do not have access to Jack's belief states, at least these are presuppositions that we can assume are at least extensionally equivalent to some of Jack's presuppositions (otherwise we cannot ascribe any belief involving "Hesperus" and "Phosphorus" at all). They even allow us to interpret the beliefs of ancient Babylonians who did not have the names "Hesperus" and "Phosphorus" but had *de re* beliefs to the effect that

“ x is bright” (where “ x ” in fact denotes Venus) extensionally equivalent to “the thing called ‘Hesperus’ (in *our* context) is bright.” Note, again, that these presuppositions are presuppositions available in the primary context. Thus, we do not strictly speaking need a secondary context for the purposes of interpreting “Hesperus” and “Phosphorus” in belief ascriptions to Jack. “Jack believes that Hesperus \neq Phosphorus” is true, given that $T = \{x \mid x$ is called “Hesperus”} and $T^* = \{y \mid y$ is called “Phosphorus”}, iff

$$(17) \text{BEL}_{\text{Jack}} : [\langle \|\lambda x_T. x\|, \|\lambda y_{T^*}. y\| \rangle \notin I(=)]$$

The content of BEL in (17), that the thing called “Hesperus” (understood *de re*) is different from the thing called “Phosphorus” (understood *de re*) is false, as it should be. But there are worlds in which the presuppositions would have allowed the content of his belief to be true – even though “Hesperus” and “Phosphorus” are interpreted as (rigid) singular terms in belief contexts. Unless we explicitly ascribe something else to Jack, the presuppositions we use when ascribing beliefs are *true* relative to our discourse context.

The approach I am considering makes belief ascriptions minimally sensitive to belief states of the ascriber. As opposed to Fregean accounts we do not try to capture the cognitive significance of the information for the ascriber, how she represents the information or what, precisely, she believes about Hesperus. The purpose of making the semantics minimally sensitive to belief contents is to achieve a proper account of the contribution the *ascription* makes to the context of the ascriber. In order to capture this we need is to make the ascription to Jack capture the fact that there is *some* way in which Jack could coherently believe what we attribute to him, not how, precisely, Jack represents that information to himself.

Thus, our account differs significantly from Fregean treatments of belief contexts. Fregean interpretations treat the expressions of clauses embedded under attitudinal operators as denoting something other than their values outside of such contexts (i.e. the customary senses of the expressions). By contrast, we treat expressions occurring under the scope of attitudinal operators as having the same kind of values as elsewhere, and not as reflecting how agents represent them. A well-known problem for Fregean accounts is the problem of iterated belief contexts; for instance, in “Jill believes that Jack believes that p ” the clause “Jack believes that p ” must refer to a customary sense as well, so the occurrence of “ p ,” which itself occurs in a belief clause in which it refers to a sense, must refer to the customary sense of the customary sense of “ p ,” and so on in a hierarchy of senses. We do not need such iterations. In “Jack believes that Mark Twain believes that Hesperus = Phosphorus, but that Samuel Clemens believes that Hesperus \neq Phosphorus” the singular terms may be bound, but the binders at all levels are presuppositions available in the context of the ascribers, us, and we need no hierarchy of senses or contexts to make sense of the ascription.

Instead, our approach has more in common with that of Fine (2007). Fine takes “Jack believes that Hesperus is F ” and “Jack believes that Phosphorus is F ” to have the same truth conditions, but adds a layer of *relational semantics* to the interpretation. In addition to Russellian propositions, Fine defines a relation over referents in (a sequence of) propositions that holds if the referent is *represented as the same* in discourse (which it is, at least to an approximation, if “no one who understands the discourse can sensibly raise the question of whether it is the same;” *ibid.*, p.40). Fine argues that Fregean *ways of representing* an object are a myth, an attempt to change the question of whether Hesperus and Phosphorus is represented as the same to a question about the intrinsic nature of representations; that is, to

whether they are represented in *the same way*, an move that he views as a residue of psychologism. Similarly, our account does not let the intrinsic nature of the representations of Hesperus and Phosphorus play a role, and we follow Fine in letting the semantics be potentially sensitive to whether the referents are represented as the same, though as opposed to Fine we transpose the layer to the context in which the sentences occur. It is context that determines the restrictions on assignments to individual terms. (In UIF the effect takes the form of making salient restrictions available through what I have called *Topics*).¹¹⁵

7.5.3. Kripke's Puzzle

Suppose Jack believes that Paderewski had musical talent and that Paderewski did not have musical talent because Jack fails to recognize that Paderewski the statesman is identical to Paderewski the pianist (Kripke 1979). The case presents few new challenges over and above the Hesperus case. Instead of viewing “Paderewski” as a referential term we view it as a variable. In “Jack believes that Paderewski had musical talent” and “Jack believes that Paderewski did not have musical talent” the two occurrences are bound by different presuppositions. Jack would use identical variables to express his beliefs, and the contributions of the variables are the same since “Paderewski” is a singular term with no descriptive content and the occurrences range over the same domain (though the domains are not the same in all worlds). That “Paderewski is *F*” can nevertheless reflect different

¹¹⁵ Admittedly, however, the semantics of UIF, presented in Chapter 5, does not incorporate the dynamic effects of context into the grammar itself, as opposed to versions of DPL (see Muskens, 1996). The value of a singular term is (always) an identity function, but instead of letting the assignment restrictions be reflected in the grammar we encode it in what amounts to the (“dynamic”) lexicon of the term. The restrictions are hence not represented compositionally – compositionality is retained for a sentence (or at least for individual clauses), but the systematic effects of context on the interpretation of a sentence – how sentence interpretations are related – are not compositional effects. Perhaps it is possible to remedy this situation, but that goes beyond the scope of this dissertation.

thoughts, or that “Paderewski = Paderewski” can be informative, is no more mysterious than the fact that two occurrences of the anaphoric “him” can be used to express different thoughts if the antecedents are different, or that two *occurrences* of the variable “*x*” can be bound by different antecedents.

The case may appear more complicated since the characters of the two name occurrences are the same, so that letting characters project binders for the names might not be enough to provide an interpretation. But this is, notice, exactly how it should be. We cannot make sense of Jack’s belief denying Paderewski = Paderewski unless we know something more about Jack’s error. Unless we wish to attribute a blatantly contradictory belief to Jack, we need some way of representing Jack’s presuppositions concerning Paderewski as being only contingently co-extensional. We could, of course, use “Person 1 called ‘Paderewski’” and “Person 2 called ‘Paderewski’” as binders if we had no clear idea of what caused Jack’s confusions regarding Paderewski, but in general even the *ascription* is infelicitous without further explication, i.e. along the lines of “Jack thinks Paderewski the statesman is different from Paderewski the pianist,” and that ascription involves differentiating information that may be used as part of the binding presuppositions that distinguish the occurrences.

If occurrences of “Paderewski” in belief contexts changed values from referent to customary (descriptive) sense, on the other hand, Paderewski cases would indeed be mysterious given that there is a limit to how fine-grained Fregean senses could be. In our framework changes in descriptions associated with “Paderewski” do not affect the meaning of “Paderewski.” The presuppositions are not different *senses* of “Paderewski,” even though modal evaluations of sentences containing “Paderewski” are affected by whether the binder

is D or D^* . The differences between D and D^* , however, may be as fine-grained (or not) as they need to be in order to make sense of Jack's beliefs. Since we know that there must be at least a pair of associated descriptions that are to blame, we can illustrate it as a second-order quantification $\exists D \exists D^* [\exists x_{Paderewski} D_{Paderewski} \wedge \exists x_{Paderewski} D^*_{Paderewski}] (x_{Paderewski} \neq x_{Paderewski}^*)$.

7.5.4. Geach Sentences

Even if belief contexts may allow reinterpretation of names, the mechanisms do not seem to license reinterpretation of ordinary pronouns. Whereas names allow presupposition projections based, for instance, on the character of the name, the characters of ordinary pronouns do not. I have not encountered any compelling example of a case where an anaphoric pronoun suggests local binding, however.

(18) A thief broke into the house. Jill believes he stole the silver.

A proper interpretation of (18) precludes a referential reading of "he." "he" is nevertheless semantically bound by "a thief," which takes wide scope over the attitudinal operator:

(19) A burglar broke into the house. Jill believes that if we had installed an alarm he wouldn't have managed to steal it.

The value of "he" is the assignment that satisfies "a burglar." The change of world of evaluation for the second sentence does not affect the assignment. Singular terms are rigid even when they function as variables.

In other cases the *antecedent* for a semantically bound term occurs under a belief operator. The most famous example, from Geach (1967), is

(20) Hob thinks a witch has blighted Bob's mare, and Nob wonders whether she killed Cob's sow.

If all we had were syntactic binding (20) would be a serious challenge. Quantifiers should not be able to scope outside of attitudinal contexts, Nob is not wondering whether *some* witch or other killed Cob's sow but whether the witch that Hob thinks blighted Bob's mare (not the witch that *actually* killed Bob's mare, for there is no such witch), and the person *uttering* (20) should not be committed to the existence of witches. In the terminology of DRT, Hob's and Nob's thoughts should converge on the same *dref*.

The occurrence of "she" cannot project a local binding presupposition, but picks up the most salient female (a strictly speaking possible reading of (20), where "she" picks up Bob's mare, is presumably ruled out as less salient). The relevant antecedent is "a witch that blighted Bob's mare." Of course, this antecedent has (*ex hypothesi*) no satisfier in the context set of the discourse, but may have one relative to an accessible set of possible worlds representing Hob's belief context, or a set of worlds of evaluation in which we ascribe something true to Hob (of course, if the speaker is open to the existence of witches, then these worlds will also be part of the context set). "she" will then have a satisfier relative to that world and the associated domain. In other words, in any world in which a witch blighted Bob's mare, Nob's belief is about that witch. Does that mean that Hob and Nob have beliefs about the same witch? Of course not. There are no witches to have beliefs about, and Nob

and Hob may entertain wildly different psychological states. What matters is that (20) *represents them* as converging on a common witch. This might be appropriate for instance because Hob's and Nob's beliefs have a common source (no such link between Hob's and Nob's states is suggested if we replace "she" by "a witch").¹¹⁶

Note that to achieve the correct interpretation an attitudinal operator has to be repeated – it is "wonders whether" that ensures that the world(s) endorsed by Hob and Nob are represented as the same. Thus, we also obtain the correct interpretation of

(21) Hob thinks a witch has blighted Bob's mare. She doesn't exist. Nob nevertheless wonders whether she killed Cob's sow.

"She doesn't exist" involves no attitudinal operator and must be evaluated relative to the actual world (or worlds in our context set). It is true, since the class of witches that blighted Bob's mare is empty in the actual world. The last sentence is nevertheless felicitous, since "she" is rigid. "she" ranges over the – in the actual world empty – domain determined *at the world of evaluation of the antecedent*, and this may not be empty.

7.6. Conclusion

That "Hesperus" and "Paderewski" could be treated as variables may not mean that we always have to treat belief ascriptions in this manner. Our framework is flexible enough to allow "Hesperus" to be anaphoric on the ascriber's presuppositions or treated as a constant in

¹¹⁶ Though even this may not always be needed for (20) to be appropriate; see Edelberg (1992).

other contexts (which we need in order to capture the *true* reading of (14)). Now, I am nevertheless somewhat hesitant concerning this strategy for dealing with Frege's puzzle. It relies on treating "Hesperus" at least sometimes as a variable, which is less easily justified than treating "Julius" as a variable. As opposed to "Julius" ordinary occurrences of "Hesperus" appears to satisfy the semantic argument and be directly referential, so we have to assume at least that hyperintensional operators may signal some kind of special treatment. Regardless of the success or not of our treatment of Frege's puzzle, however, pluralism about proper names provides a compelling and systematic solution to the problem of the contingent *a priori* and the problem of empty names, as illustrated for instance by embedding our account in a (pseudo-)two-dimensional system. I leave it for later work to investigate further the variants of Frege's puzzle, and the extent to which the solution to empty names would also covers the case of fiction.

REFERENCES

- Abbott, B. (2001). Definiteness and Indefiniteness. In Horn, L.R. & Ward, G. (eds.). *The Handbook of Pragmatics*. Oxford: Blackwell: 122–50.
- Adams, F., & Dietrich, L.A. (2004). Whats in a(n Empty) Name? *Pacific Philosophical Quarterly* 85: 125–48.
- Adams, R.M. (1979). Primitive Thisness and Primitive Identity. *Journal of Philosophy* 76: 5–26.
- Almog, J. (1986). Naming without Necessity. *Journal of Philosophy*, vol.83: 210–42.
- Almog, J., Perry, J. & Wettstein, H. (eds.) (1989). *Themes from Kaplan*. Oxford: Oxford University Press.
- Bach, E. & Partee, B. (1984). Quantification, Pronouns, and VP Anaphora. In Groenendijk, J., Janssen, T. & Stokhof, M. (eds.): *Truth, Interpretation and Information*, Foris, Dordrecht: 99–130.
- Bach, K. (1987). *Thought and Reference*, Oxford University Press.
- Bach, K. (1999). The Semantics-Pragmatics Distinction: What it is and Why it Matters. In Turner, K. (ed.). *The Semantics-Pragmatics Interface from Different Points of View*. Oxford: Elsevier: 65–84.
- Bach, K. (2002). Giorgione was So-Called Because of His Name. *Philosophical Perspectives* 16, *Language and Mind*: 73–103.
- Balaguer, M. (2005). Indexical Propositions and De Re Belief Ascriptions. *Synthese* 146: 325–55.
- Balaguer, M. (2011). Is There a Fact of the Matter between Direct Reference Theory and (Neo-)Fregeanism? *Philosophical Studies* 154(1): 53–78.
- Barcan Marcus, R. (1961). Modalities and Intensional Languages. *Synthese* 13(4): 303–22.
- Barker, C. (2007). Direct Compositionality on Demand. In C. Barker & P. Jacobson (eds.). *Direct Compositionality*, Oxford: Oxford University Press: 102–31.

- Barker, C. & Shan, C. (2008). Explaining Crossover and Superiority as Left-to-Right Evaluation. *Linguistics and Philosophy* 29: 91–134.
- Barwise, J. & Cooper, R. (1981). Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4(2): 159–219.
- Barwise, J. & Perry, J. (1983). *Situations and Attitudes*. MIT Press.
- Baumann, P. (2010). Are Proper Names Rigid Designators? *Axiomathes* 20: 333–46.
- Beaver, D. (1992). The kinematics of presupposition. In Dekker, P. & Stockhof, M. (eds.), *Proceedings of the Eighth Amsterdam Colloquium*, University of Amsterdam: ILLC: 17–36.
- Berger, A. (2002): *Terms and Truth*, MIT Press.
- Berger, A. (2006). Precis of *Terms and Truth*: Reference Direct and Anaphoric. *Philosophy and Phenomenological Research* vol. 72(3): 640–9.
- Bezuidenhout, A. & Reimer, M. (eds.) (2004). *Descriptions and Beyond*, Oxford University Press, Oxford.
- Bochner, G. (2013). The Metasyntactic Interpretation of Two-Dimensionalism. *Philosophical Studies* 163: 611–26.
- Boolos, G. (1984). To Be is to Be the Value of a Variable (or to Be Some Values of Some Variables). *The Journal of Philosophy* 81: 430–49.
- Boolos, G. (1985). Nominalistic Platonism. *The Philosophical Review* 94.3: 327–44.
- Braun, D. (1993). Empty Names. *Noûs* 27(4): 449–69.
- Braun, D. (1998). Understanding Belief Reports. *The Philosophical Review*, vol.107, no.4: 555–95.
- Braun, D. (2002). Cognitive Significance, Attitude Ascription, and Ways of Believing a Proposition. *Philosophical Studies* 108: 65–81.
- Braun, D. (2003). Empty Names, Fictional Names, Mythical Names. *Noûs* 39(4): 596–631.
- Breckenridge, W. & Magidor, O. (2012). Arbitrary Reference. *Philosophical Studies* 158: 377–400.

- Brock, S. (2010). The Creationist Fiction: The Case against Creationism about Fictional Characters. *Philosophical Review* vol 119(3): 337–64.
- Burge, T. (1973). Reference and Proper Names. *Journal of Philosophy* 70(14): 425–39.
- Burgess, A. (forthcoming). Metalinguistic Descriptivism for Millians. *Australasian Journal of Philosophy* DOI: 10.1080/00048402.2012.711760.
- Carnap, R. (1947). *Meaning and Necessity*. Chicago: University of Chicago Press [1988].
- Chalmers, D. (1996). *The Conscious Mind*. New York: Oxford University Press.
- Chalmers, D. (2002). On Sense and Intension. *Philosophical Perspectives* 16: 135–82.
- Chalmers, D. (2006). The Foundations of Two-Dimensional Semantics. In García-Carpintero, M. & Macià, J. (eds.) (2006): 55–140.
- Chalmers, D. (2011). Propositions and Attitude Ascriptions: A Fregean Account. *Noûs* 45 (4): 595–639.
- Chastain, C. (1975). Reference in Context. In Gunderson, K. (ed.): *Language, Mind and Knowledge*, Minnesota Studies in the Philosophy of Science 7, University of Minnesota Press: 194–269.
- Chierchia, G. (1989). Anaphora and Attitudes *De Se*. In Bartsch, R., van Benthem, J. & van Emde Boas, A. (eds.): *Semantics and Contextual Expression*, Foris, Dordrecht, Netherlands: 1–32.
- Chierchia, G. (1995). *Dynamics of Meaning*. Chicago University Press, Chicago.
- Church, A. (1943). Review of Quine's 'Notes on Existence and Identity.' *Journal of Symbolic Logic*: 45–47.
- Corazza, E. (2002). Description-Names. *Journal of Philosophical Logic* 31 (4): 313–25.
- Cresswell, M.J. (1985). *Structured Meanings*. MIT Press, Cambridge, MA.
- Cresswell, M.J. & von Stechow, A. (1982). *De Re* Belief Generalized. *Linguistics and Philosophy* 5(4): 503–35.
- Crimmins, M. (1992). *Talk About Beliefs*, MA: MIT Press.
- Cumming, S. (2008). Variabilism. *Philosophical Review* vol.117 no.4: 525–54.

- Cumming, S. (Unpublished). *Proper Nouns*. Dissertation submitted to the Graduate School, New Brunswick Rutgers, available from:
http://mss3.libraries.rutgers.edu/dlr/TMP/rutgers-lib_22845-PDF-1.pdf.
- Currie, G. (1990). *The Nature of Fiction*. Cambridge: Cambridge University Press.
- Davies, M. (1981). *Meaning, Quantification, Necessity: Themes in Philosophical Logic*. Routledge & Kegan Paul.
- Davies, M. (2006). Reference, Contingency, and the Two-Dimensional Framework. In García-Carpintero M. & Macià, J. (eds.): 141–75.
- Davies, M. & Humberstone, L. (1980). Two Notions of Necessity. *Philosophical Studies* 38: 1–30.
- Dekker, P. (1993). Existential Disclosure. *Linguistics and Philosophy* 16: 561–88.
- Dekker, P. (2012). *Dynamic Semantics: Studies in Linguistics and Philosophy vol. 91*. Springer, Dordrecht.
- Dekker, P. & van Rooy, R. (1998). Intentional Identity and Information Exchange. In Cooper, R. & Gamkrelidze, T. (eds.), *Proceedings of the Second Tbilisi Symposium on Language, Logic and Computation*. Tbilisi State University, Tbilisi.
- Dever, J. (2001). Complex Demonstratives. *Linguistics and Philosophy* 24: 271–330.
- Devitt, M. (1974). Singular Terms. *Journal of Philosophy* LXXI: 183–205.
- Devitt, M. (1981). *Designation*. New York: Columbia University Press.
- Devitt, M. (2012). Should Proper Names Still Seem So Problematic? In Bianchi, A. (ed.): *New Essays on Reference*. Forthcoming.
- Donnellan, K. (1966). Reference and Definite Descriptions. *Philosophical Review*, 75: 281–304.
- Donnellan, K. (1978). Speaker Reference, Descriptions, and Anaphora. In Cole, P. (ed.), *Syntax and Semantics 9: Pragmatics*, New York: Academic Press: 47–68.
- Donnellan, K. (1979). The Contingent *A Priori* and Rigid Designators. In French, P., Uehling, T. & Wettstein, H. (eds.): *Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press: 45–60.
- Dummett, M. (1973). *Frege: Philosophy of Language*. New York: Harper and Row.

- Dummett, M. (1975). What is a Theory of Meaning? In Guttenplan, S. (ed.). *Mind and Language*. Oxford University Press.
- Dummett, M. (1981). *The Interpretation of Frege's Philosophy*. London: Duckworth, and Cambridge MA: Harvard University Press.
- Dummett, M. (1991). *The Logical Basis of Metaphysics*. London: Duckworth, and Cambridge MA: Harvard University Press.
- Edelberg, W. (1992). Intentional Identity and the Attitudes. *Linguistics and Philosophy* 15(6): 561–96.
- van Eijck, J. & Visser, A. (2010). Dynamic Semantics. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/dynamic-semantics/#EncDynTypLog> (accessed 2.24.13).
- Elbourne, P. (2001). E-type Anaphora as NP Deletion. *Natural Language Semantics* 9(3): 241–88.
- Elbourne, P. (2005). *Situations and Individuals*. MIT Press, Cambridge, MA.
- Evans, G. (1973). The Causal Theory of Names. *Proceedings of the Aristotelian Society*, Supplementary Volume 47: 187–208.
- Evans, G. (1977). Pronouns, Quantifiers and Relative clauses. *Canadian Journal of Philosophy* 7:3: 467–536.
- Evans, G. (1978). Speaker Reference, Descriptions, and Anaphora. In Cole, P. (ed.): *Syntax and Semantics 9: Pragmatics*, New York: Academic Press: 47–68.
- Evans, G. (1979). Reference and Contingency. *The Monist* 62: 161–89. Repr. In Evans (1985a): 178–213; page references are to the latter.
- Evans, G. (1980). Pronouns. *Linguistic Inquiry* 11: 337–62.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.
- Evans, G. (1985a). *Collected Papers*. Oxford: Oxford University Press.
- Everett, A. (2003). Empty Names and ‘Gappy’ Propositions. *Philosophical Studies* 116: 1–36.
- Everett, A. (2005). Recent Defenses of Descriptivism. *Mind & Language* vol.20 no.1: 103–39.

- Fine, K. (1985). *Reasoning With Arbitrary Objects*, Blackwell.
- Fine, K. (2003). The Role of Variables. *The Journal of Philosophy* vol.100 no.12: 605–31.
- Fine, K. (2007). *Semantic Relationism*. Blackwell Publishing.
- Fitch, G. W. (2004). *Saul Kripke*. Chesham: Acumen.
- Fodor, J.D. & Sag, I. (1982). Referential and Quantificational Indefinites. *Linguistics and Philosophy* 5: 355–98.
- Føllesdal, D. (1986). Essentialism and Reference. In Hahn, L.E. & Schlipp, P.A. (eds.). *The Philosophy of W.V. Quine*. LaSalle, IL: Open Court: 97–115.
- Forbes, G. (1989). *Languages of Possibility*. Oxford: Blackwell.
- Forbes, G. (2011). The Problem of Factives for Sense Theories. *Analysis* vol.71:4: 654–62.
- Frege, G. (1893/1903). *Die Grundlagen der Arithmetik: eine logisch-mathematische Untersuchung über den Begriff der Zahl*. Breslau: W. Koebner. Translated as *The Foundations of Arithmetic*, by Austin, J.L., Oxford: Basil Blackwell, 1974.
- Frege, G. (1918). Der Gedanke. Eine logische Untersuchung. *Beiträge zur deutschen Idealismus* 2, 1918-19: 58–77. In English as The thought. A logical inquiry. In Geach, P. & Black, M. (eds.): *Translations from the Philosophical Writings by Gottlob Frege*, Blackwell, Oxford 1952.
- García-Carpintero, M. (2006). Two-Dimensionalism: A Neo-Fregean Interpretation. In García-Carpintero, M. & Macià, J. (2006): 181–205.
- García-Carpintero M. & Macià, J. (eds.) (2006). *Two-Dimensional Semantics*. Oxford, Clarendon Press.
- Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic.
- Geach, P.T. (1967). Intentional Identity. *Journal of Philosophy*, vol. 74, no. 20: 627–32.
- Geurts, B. (1996). Local Satisfaction Guaranteed. A Presupposition Theory and Its Problems. *Linguistics and Philosophy* 21: 545–601.
- Geurts, B. (1997). Good News About the Description Theory of Names. *Journal of Semantics* 14: 319–48.

- Geurts, B. (1999). *Presuppositions and Pronouns*. Elsevier, Oxford.
- Green, M. (2007). Direct Reference, Empty Names and Implicature. *Canadian Journal of Philosophy* 37(3): 419–48.
- Greenberg, M. & Harman, G. (2006). Conceptual Role Semantics. In Lepore, E. & Smith, B. (eds.), *The Oxford Handbook of Philosophy of Language*, New York: Oxford University Press: 295–322.
- Groenendijk, J. & Stokhof, M. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy*, 14: 39–100.
- Groenendijk, J. & Stokhof, M. (2000). Meaning in Motion. In von Stechow, K. & Egli, U. (eds.), *Reference and Anaphoric Relations*. Kluwer.
- Gundel, J., Hedberg, N., & Zacharski, R. (1993). Cognitive Status and the Form of Referring Expressions in Discourse. *Language* 69(2): 274–307.
- Haraldsen, F. (2012). Reference, Rigidity, and Contingent Identity. *SATS: Northern European Journal of Philosophy* 13(2): 116–27.
- Haraldsen, F. (forthcoming). What Russell Couldn't Describe. *Australasian Journal of Philosophy*.
- Haukioja, J. (2012). Rigidity and Actuality-Dependence, *Philosophical Studies* 157: 399–410.
- Hawthorne, J. & Lepore, E. (2011). On Words. *The Journal of Philosophy* 108(9): 447–85.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- Heim, I. (1983). File Change Semantics and the Familiarity Theory of Definites. In Bäuerle, R., Schwarze, R., & von Stechow, A. (eds.). *Meaning, Use, and Interpretation of Language*. Berlin: deGruyter: 164–89.
- Heim, I. (1990). E-type Pronouns and Donkey Anaphora. *Linguistics and Philosophy* 13: 137–77.
- Heim, I. (1992). Presupposition Projection and the Semantics of Attitude Reports. *Journal of Semantics* 9: 183–221.

- von Heusinger, K. (2004). Choice Functions and the Anaphoric Semantics of Definite NPs. *Research on Language and Computation* 2: 309–29.
- Higginbotham, J. (1987). Indefiniteness and Predication. In Reuland, E. & ter Meulen, A. (eds.). *The Representation of (In)definiteness*. Cambridge: MIT Press: 43–70.
- Horwich, P. (1998). *Meaning*. Oxford, Oxford University Press.
- Hughes, G.E. & Cresswell, M.J. (1968). *An Introduction to Modal Logic*. Methuen, London.
- Jackson, F. (1998a). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Jackson, F. (1998b). Reference and Description Revisited. In Tomberlin, J.E. (ed.). *Philosophical Perspectives, 12: Language, Mind and Ontology*. Malden, MA: Blackwell Publishers: 201–18.
- Jacobson, P. (1999). Towards a Variable-Free Semantics. *Linguistics and Philosophy* 22: 117–84.
- Jacobson, P. (2000). Paycheck pronouns, Bach-Peters sentences, and variable-free semantics. *Natural Language Semantics* 8, 77–155.
- Jeshion, R. (2001). Donnellan on Neptune. *Philosophy and Phenomenological Research* vol.63 no.1: 111–35.
- Jeshion, R. (2004). Descriptive Descriptive Names. In Bezuidenhout, A. & Reimer, M. (eds.) (2004): 591–612.
- Jeshion, R. (2009). The Significance of Names. *Mind & Language* 24 no.4: 370–403.
- Kadmon, N. (1990). Uniqueness. *Linguistics and Philosophy* 13: 273–324.
- Kalish, D. & Montague, R. (1964). *Logic: Techniques of Formal Reasoning*. Harcourt, Brace and World, Inc, New York.
- Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J.; Janssen, T; & Stokhof, M. (eds.): *Formal Methods in the Study of Language*. Mathematical Centre Tracts 135, Amsterdam: 277–322.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.
- Kanazawa, M. (1994). Weak vs. Strong Readings of Donkey Sentences in a Dynamic Setting. *Linguistics and Philosophy* vol. 17(2): 109–58.

- Kanterian, E. (2010). Puzzles About Descriptive Names. *Linguistics and Philosophy* 32: 409–28.
- Kaplan, D. (1968). Quantifying In. *Synthese* 19: 178–214.
- Kaplan, D. (1973). Bob and Carol and Ted and Alice. In Hintikka, J. et al. (eds.) *Approaches to Natural Language*, Reidel, 1973: 490–518.
- Kaplan, D. (1977). Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals. In Almog, J., Perry, J. and Wettstein, H. (eds.) (1989): 481–563.
- Kaplan, D. (1986). Opacity. In Hahn, L.E. & Schlipp, P.A. (eds.). *The Philosophy of W.V. Quine*. LaSalle, IL: Open Court: 229–89.
- Kaplan, D. (1989). Afterthoughts. In Almog, J., Perry, J. and Wettstein, H. (eds.) (1989): 567–614.
- Kaplan, D. (1990). Words. *Aristotelian Society Supplement* 64, 93–119.
- Karttunen, L. (1976). Discourse Referents. In McCawley, J. (ed.): *Syntax and Semantics: Notes from the Linguistic Underground*, vol.7: 363–86.
- Katz, J. (1990). Has the Description Theory of Names been Refuted? In Boolos, G. (ed.). *Meaning and Method: Essays in Honor of Hilary Putnam*. Cambridge, UK: Cambridge University Press.
- Katz, J. (1994). Names without Bearers. *Philosophical Review* 103: 1–39.
- Katz, J. (2001). The End of Millianism: Multiple Bearers, Improper Names, and Compositional Meaning. *The Journal of Philosophy* 98: 137–66.
- King, J.C. (1987). Pronouns, Descriptions and the Semantics of Discourse. *Philosophical Studies* 51: 341–63.
- King, J.C. (1994). Anaphora and Operators. *Philosophical Perspectives* 8, Logic and Language: 221–50.
- King, J.C. (1995). Structured Propositions and Complex Predicates. *Noûs* 29(4): 516–35.
- King, J.C. (2004). Context Dependent Quantifiers and Donkey Anaphora, *Canadian Journal of Philosophy* 34: 97–127.
- King, J.C. (2007). *The Nature and Structure of Content*. Oxford University Press, Oxford

- Kneale, W. (1962). Modality, De Dicto and De Re. In Nagel, E., Suppes, P., & Tarski, A. (eds.) *Logic, Methodology and the Philosophy of Science: Proceedings of the 1960 International Congress*. Stanford University Press: 622–33.
- Kratzer, A. (1989). An Investigation into the Lumps of Thought. *Linguistics and Philosophy* 12: 607–53.
- Kripke, S. (1959). A Completeness Theorem for Modal Logic. *Journal of Symbolic Logic* 24: 1–14.
- Kripke, S. (1963). Semantical Considerations for Modal Logic. *Acta Philosophica Fennica* 16: 83–94.
- Kripke, S. (1972). *Naming and Necessity*. Cambridge: Harvard University Press; 2nd edition [1980]; page references are to the latter.
- Kripke, S. (1979). A Puzzle About Belief. In Margalit, A. (ed.) *Identity and Individuation*, New York: New York University Press: 135–64.
- Kroon, F. (1987). Causal Descriptivism. *Australasian Journal of Philosophy*, 65: 1–17.
- Kroon, F. (1989). Circles and Fixed Points in Description Theories of Reference. *Noûs* 23: 373–92.
- Larson, R. & Ludlow, P. (1993). Interpreted Logical Forms. *Synthese* 95: 305–55.
- Larson, R. & Segal, G. (1995). *Knowledge of Meaning*. Cambridge: MIT Press.
- Lewis, D. (1970). General Semantics. *Synthese* vol. 22, 1-2: 18-67.
- Lewis, D. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8: 339–59.
- Lewis, D. (1984). Putnam's Paradox. *Australasian Journal of Philosophy* 62: 221–36.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Loar, B. (1976). The Semantics of Singular Terms. *Philosophical Studies* 30, 353–77.
- Ludlow, P. (1994). Conditionals, Events, and Unbound Pronouns. *Lingua e Stile* 19: 3–20.
- Ludlow, P. & Segal, G. (2003). On a Unitary Semantical Analysis for Definite and Indefinite Descriptions. In Reimer, M. & Bezuidenhout, A. (eds.): 420–36.

- Maier, E. (2009). Proper Names and Indexicals Trigger Rigid Presuppositions. *Journal of Semantics* 26: 253–315.
- Martí, G. (2003). The Question of Rigidity in New Theories of Reference. *Noûs* 37: 161–79.
- Matushansky, O. (2006). Why Rose is the Rose. In Bonami, O. & Hofherr, P.C. (eds.). *Empirical Issues in Formal Syntax and Semantics* vol.6: 285–308.
- McCulloch, G. (1989). *The Game of the Name*. Oxford: Oxford University Press.
- McDowell, J. (1977). On the Sense and Reference of a Proper Name. *Mind* 86: 159–85.
- McDowell, J. (1984). *De Re* Senses. In Wright, C (ed.). *Frege: Tradition & Influence*. Oxford: Basil Blackwell: 98–109.
- McGinn, C. (1982). Rigid Designation and Semantic Value. *Philosophical Quarterly* vol. 32: 97–115.
- McKinsey, M. (1986). Mental Anaphora. *Synthese* 66: 159–75.
- Mill, J.S. (1843). *A System of Logic*. London: Longmans, 1947.
- Moltmann, F. (2006). Unbound Anaphoric Pronouns: E-type, Dynamic, and Structured Propositions Approaches. *Synthese* 153: 199–260.
- Moltmann, F. (2006a). Presupposition and Quantifier Domains. *Synthese* 149: 179–224.
- Montague, R. (1973). The Proper Treatment of Quantification in Ordinary English. In Hintikka, J., Moravcsik, J. & Suppes, P. (eds.). *Approaches to Natural Language*. Dordrecht: Reidel: 221–42.
- Moro, A. (2007). Existential Sentences and Expletive *There*. In Everaert, M. & Riemsdijk, H. (eds.): *The Blackwell Companion to Syntax II*. Wiley-Blackwell: 210–36.
- Mousavian, S.N. (2010). Neo-Meinongian Neo-Russellians. *Pacific Philosophical Quarterly* 91(2): 229–59.
- Muskens, R. (1996): Combining Montague Semantics and Discourse Representation. *Linguistics and Philosophy* 19: 143–86.
- Muskens, R. (2011). A Squib on Anaphora and Coindexing. *Linguistics and Philosophy* 34: 85–9.

- Muskens, R., van Benthem, J. & Visser, A. (1997). Dynamics. In van Benthem, J. & ter Meulen, A. (eds.) *Handbook of Logic and Language*. Elsevier/MIT Press: 587–649.
- Neale, S. (1990). *Descriptions*. Cambridge: MIT Press Books.
- Neale, S. (1993). Term Limits. *Philosophical Perspectives* 7: 89–123.
- Neale, S. (2005). Pragmatism and Binding. In Szabó, Z. (ed.) *Semantics versus Pragmatics*. Oxford University Press: 165–285.
- Neale, S. (2009). Term Limits Revisited. *Philosophical Perspectives* 22: 375–442.
- Nelson, M. (2002). Descriptivism Defended. *Noûs* 36(3): 408–35.
- Ninan, D. (2012). Propositions, Semantic Values, and Rigidity. *Philosophical Studies* 158(3): 401–13.
- Parsons, T. (1980). *Nonexistent Objects*. New Haven: Yale University Press.
- Pasniczek, J. (2002). Equating Categorially Names and Quantifiers Within First-Order Logic. *Logic and Logical Philosophy* (10): 119–29.
- Partee, B. (1989). Binding Implicit Variables in Quantified Contexts. In Wiltshire, C. et al. (eds.): *Papers from the Twenty Fifth Meeting of the Chicago Linguistic Society*. Chicago: Chicago Linguistic Society: 342–65.
- Peregrin, J. & von Heusinger, K. (2004). Dynamic Semantics with Choice Functions. In Kamp, H. & Partee, B. (eds.) *Context Dependence in the Analysis of Linguistic Meaning*. Elsevier Science, Amsterdam.
- Perry, J. (1989). The Problem of the Essential Indexical. *Noûs* 13(1): 3–21.
- Peters, S. & Westerståhl, D. (2006). *Quantifiers in Language & Logic*. Clarendon Press: Oxford.
- Piccinini, G. & Scott, S. (2010). Recovering What Is Said With Empty Names. *Canadian Journal of Philosophy* 40(2): 239–73.
- Pickel, B. (2012). Rigidification and Attitudes. *Philosophical Studies* 158: 43–58.
- Prawitz, D. (1985). *Natural Deduction: A Proof-Theoretical Study*. Almqvist & Wiksell.
- Quine, W. v. O. (1943). Notes on Existence and Necessity. *The Journal of Philosophy*, vol. 40: 113–27.

- Quine, W. v. O. (1947). The Problem of Interpreting Modal Logic. *Journal of Symbolic Logic* vol.12: 43–48.
- Quine, W. v. O. (1956). Quantifiers and Propositional Attitudes. *The Journal of Philosophy*, vol. 53: 177–87.
- Quine, W. v. O. (1977). Intensions Revisited. In *Theories and Things*. Cambridge MA: Harvard University Press: 113–23.
- Recanati, F. (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Recanati, F. (2000). *Oratio Obliqua, Oratio Recta: an Essay on Metarepresentation*. Cambridge, MA: MIT Press.
- Reimer, M. (2004). Descriptively Introduced Names. In Bezuidenhout, A. & Reimer, M. (eds.) (2004): 613–29.
- Richard, M. (1990). *Propositional Attitudes*. Cambridge: Cambridge University Press.
- Roberts, C. (1996). Information Structure: Towards an Integrated Formal Theory of Pragmatics. In Yoon, J.H. & Kathol, A. (eds.): *OSUWPL vol.49: Papers in Semantics*, The Ohio State University Department of Linguistics.
- Roberts, C. (2003). Uniqueness in Definite Noun Phrases. *Linguistics and Philosophy* 26: 287–350.
- Roberts, C. (2004). Pronouns as Definites. In Bezuidenhout, A. & Reimer, M. (eds.): Oxford: 503–43.
- Russell, B. (1905). On Denoting. *Mind* 14: 479–93.
- Russell, B. (1911). Knowledge by Acquaintance and Knowledge by Description. *Proceedings of the Aristotelian Society* (New Series), Vol.XI, (1910–11): 108–28.
- Sainsbury, R.M. (2002). *Departing From Frege: Essays in the Philosophy of Language*. London: Routledge.
- Sainsbury, R.M. (2005). *Reference Without Referents*. Oxford: Clarendon Press.
- Sainsbury, R.M. (2007). *Fiction and Fictionalism*. London: Routledge.
- Salmon, N. (1982). *Reference and Essence*. Oxford: Basil Blackwell. Repr. Prometheus Books [2005]; all page references are to the latter.
- Salmon, N. (1986). *Frege's Puzzle*. MIT Press/Bradford Books.

- Salmon, N. (1986a). Reflexivity. *Notre Dame Journal of Formal Logic* 27(3): 401–29.
- Salmon, N. (1987). Existence. In Tomberlin, J. (ed.): *Philosophical Perspectives I: Metaphysics*. Atascadero, Calif.: Ridgeview: 49–108.
- Salmon, N. (1989). How to Become a Millian Heir. *Noûs* 23: 211–20.
- Salmon, N. (1990). A Millian Heir Rejects the Wages of Sinn. In Anderson, C.A. & Owens, J.: *Propositional Attitudes; the Role of Content in Logic, Language, and Mind*. Stanford, California: Center for the Study of Language and Information, Stanford University: 215–47. Repr. in Salmon (2007): 3–31; page references are to the latter.
- Salmon, N. (1998). Nonexistence. *Noûs*, vol.32 no.3: 277–319.
- Salmon, N. (2002). Demonstrating and Necessity. *The Philosophical Review* 111 no.4: 497–537. Repr. in Salmon (2007): 67–99.
- Salmon, N. (2003). Naming, Necessity, and Beyond. *Mind*, 112: 475–92.
- Salmon, N. (2006). A Theory of Bondage. *The Philosophical Review*, vol.115 no.4: 415–48. Repr. in Salmon (2007): 113–40.
- Salmon, N. (2007). *Content, Cognition, and Communication*. Clarendon Press, Oxford.
- van der Sandt, R. (1992). Presupposition Projection as Anaphora Resolution. *Journal of Semantics* 9: 333–78.
- Schiffer, S. (2000). Pleonastic Fregeanism and Empty Names. In Everett, A. & Hofweber, T. (eds.): *Empty Names, Fiction, and the Puzzles of Non-Existence*. University of Chicago Press: 3–16.
- Schiffer, S. (2002). Meanings. In Campbell, J.K. et al. (eds.). *Meaning and Truth*. New York: Seven Bridges Press: 79–104.
- Schiffer, S. (2003). *The Things We Mean*. Oxford University Press.
- Schlenker, P. (2003). A Plea for Monsters. *Linguistics and Philosophy* 26: 29–120.
- Schlenker, P. (2010). Local Contexts and Local Meanings. *Philosophical Studies* 151: 115–42.
- Smith, A.D. (1984). Rigidity and Scope. *Mind* vol.93: 177–93.

- Smullyan, A. (1947). Review: W. V. Quine, 'The Problem of Interpreting Modal Logic.' *Journal of Symbolic Logic* vol.12(4): 139–41.
- Soames, S. (1987). Direct Reference, Propositional Attitudes and Semantic Content. *Philosophical Topics* 15: 47–87
- Soames, S (1989). Review of Gareth Evans' *Collected Papers*. *The Journal of Philosophy* 86: 141–56.
- Soames, S (1995). Revisionism about Reference: A Reply to Smith. *Synthese* vol.104, no. 2: 191–216.
- Soames, S (1998). The Modal Argument: Wide Scope and Rigidified Descriptions. *Noûs* 32: 1–22.
- Soames, S (2002). *Beyond Rigidity*. Oxford University Press.
- Soames, S (2005). *Reference and Description: The Case against Two-Dimensionalism*. Princeton University Press.
- Soames, S (2006). Kripke, the Necessary A Posteriori, and the Two-Dimensionalist Heresy. In García-Carpintero, M. & Macià, J. (2006): 272–92.
- Soames, S (2007). Descriptive Names vs. Descriptive Anaphora. *Philosophy and Phenomenological Research* vol. 72(3): 665–73.
- Soames, S (2010a). *What is Meaning?* Princeton University Press, Princeton, NJ.
- Sommer, F. (1982). *The Logic of Natural Language*. Clarendon Press, Oxford.
- Sosa, D. (2001). Rigidity in the Scope of Russell's Theory. *Noûs* 35: 1–38.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics* 9. New York Academy Press: 315–32.
- Stalnaker, R. (1987). Belief Attribution and Context. In Grimm, R. & Merrill, R. (eds.) *Contents and Thought*. University of Arizona Press, Tucson: 143–56.
- Stalnaker, R. (1998). On the Representation of Context. *Journal of Logic, Language, and Information* 7. 3–19.
- Stalnaker, R. (2001). On Considering a Possible World as Actual. *Proceedings of the Aristotelian Society, Supplementary Volume* 75: 141–56.

- Stalnaker, R. (2006). Assertion Revisited: On the Interpretation of Two-Dimensional Modal Semantics. In García-Carpintero, M. & Macià, J. (2006): 293–309.
- Stanley, J. (1997). Names and Rigid Designation. In Wright, C. & Hale, R. (eds.). *A Companion to the Philosophy of Language*. Oxford, Blackwell Press: 555–85.
- Stanley, J. (2002). Modality and What is Said. *Philosophical Perspectives* 16: 321–44.
- Stanley, J. (2002a). Nominal Restriction. In Preyer, G. & Peter, G. (eds.). *Logical Form and Language*. Oxford: Oxford University Press: 365–88.
- Stanley, J. & Szabó, Z. (2000). On Quantifier Domain Restriction. *Mind and Language* 15: 219–61.
- Steedman, M. & Stone, M. (2006). Is Semantics Computational? *Theoretical Linguistics* 32(1): 73–89.
- Stephanou, Y. (2001). Indexed Actuality. *Journal of Philosophical Logic* vol.30 no.4: 355–93.
- Sullivan, A. (2005). Rigid Designation, Direct Reference, and Modal Metaphysics. *Pacific Philosophical Quarterly* 86: 577–99.
- Sullivan, A. (2007). Rigid Designation and Semantic Structure. *Philosophers' Imprint* 7 (6): 1–22.
- Sullivan, A. (2009). Against Structured Referring Expressions. *Philosophical Studies* 146 (1): 49–74.
- Sutton, J. (2001). The Contingent *A Priori* and Implicit Knowledge. *Philosophy and Phenomenological Research* 63: 251–77.
- Szabó, Z. G. (2005). The Loss of Uniqueness. *Mind* 114: 1185–1222.
- Szabolcsi, A. (2010). *Quantification*. Cambridge University Press.
- Taylor, K. (2000). Emptiness without Compromise. In Everett, A. & Hofweber, T. (eds.). *Empty Names, Fiction and The Puzzles of Non-Existence*. CSLI Publications: 37–60.
- Thomason, R.H. & Stalnaker, R.C. (1968). Modality and Reference. *Noûs*, 2, 359–72.
- Thomasson, A. (1999). *Fiction and Metaphysics*. Cambridge: Cambridge University Press.
- Thomasson, A. (2003). Speaking of Fictional Characters. *Dialectica* 57(2): 207–26.

- Wiggins, D (1975). Identity, Designation, Essentialism and Physicalism. *Philosophia* V: 1–30.
- Winter, Y. (1997). Choice Functions and the Scopal Semantics of Indefinites. *Linguistics and Philosophy* 20: 399–467.
- Wittgenstein, L. (1953). *Philosophical Investigations*. 3rd ed. Anscombe, G.E.M. (transl.), New York: Macmillan Co.
- Wyatt, N. (2007). The Pragmatics of Empty Names. *Dialogue* XLVI: 663–81.
- Yagisawa, T. (2001). Against Creationism in Fiction. *Philosophical Perspectives* 15: 153–72.
- Zalta, E. (1983). *Abstract Objects*. Dordrecht: Dordrecht Reidel.
- Zalta, E. (1988). *Intensional Logic and the Metaphysics of Intentionality*. Cambridge: MIT Press.
- Zalta, E. (2003). Referring to Fictional Characters. *Dialectica* 57: 243–54.
- Zeevat, H. (1999). Demonstratives in Discourse. *Journal of Semantics* 16: 279–313.
- Ziff, P. (1960): *Semantic Analysis*. Ithaca: Cornell University Press.
- Zimmermann, T.E. (1998). On the Epistemic Role of Discourse Referents. In Moss, L. et al. (eds.), *Logic, Language, and Communication*. CSLI Publications, Stanford.