

PROPER NOUNS

by

SAMUEL J. CUMMING

A dissertation submitted to the
Graduate School–New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
Graduate Program in Philosophy
written under the direction of
Prof. Ernest Lepore and Prof. Jason Stanley
and approved by

New Brunswick, New Jersey

October, 2007

ABSTRACT OF THE DISSERTATION

Proper Nouns

By SAMUEL J. CUMMING

Dissertation Directors:

Ernest Lepore and Jason Stanley

This dissertation is an experiment: what happens if we treat proper names as anaphoric expressions on a par with pronouns? The first thing to notice is that a name's 'antecedent' can occur in a discourse prior to the one containing the name. An individual may be introduced and tagged with a name in one context, and then retrieved – using the name – in a later context. To allow for discourse-crossing anaphora, in addition to the usual cross-sentential anaphora, a revision of discourse semantics is in order. Essentially, we must countenance discourse referents that *span* contexts, and think of contexts, not as islands, but as nodes connected to each other by the discourse referents they share.

Discourse semantics gives rise to a new notion of *content* determined by *discourse reference* rather than pure reference. In a space of contexts structured by shared discourse referents, discourse content becomes *transmissible*. For a piece of content may be sent from one context to another whenever the discourse referents bundled up in the content are held in common by the two contexts.

The final step is to treat the cognitive state of an agent as just another kind of context, and so a potential recipient of discourse content. Discourse content is more fine-grained than traditional 'singular' content, and so is a better fit for our pre-theoretic intuitions about communication and attitude reporting. This is illustrated by applying the theory to Frege's puzzle, a puzzle of Loar's about communication, Kripke's puzzle about belief, Geach's intentional identity and new breed of 'mixed' *de re–de dicto* sentence

Acknowledgements

This will be short and sweet. First, I will appreciate my committee: Jason, who helped me bridge the gap between philosophy and linguistics, drawing on his own trailblazing experience. Ernie, for his indefatigable comments and questions, which always tested me on the points most in need of testing. John, whose timely suggestions got me out of several jams. And Matthew, who ran the long course, reading late drafts and participating in eleventh-hour meetings. I owe all of these people personal debts, as each has had a different, equally important role in my development. However these debts must remain, well, personal. Thanks also to Mark Richard for his trenchant suggestions and useful criticisms, and Jeff King for smoothly assuming John's role at my defence.

The central ideas in this dissertation were coaxed into existence during meetings with Adrian Brasoveanu in Frankfurt and New Brunswick. Around this time I also benefitted from meetings with Chung-Chieh Shan, who had just joined the faculty at Rutgers. My interest in names and attitude reports sprang from an example (of 'mixed' *de re-de dicto*) that was the topic of a Graduate Student presentation at Rutgers and a later presentation at the Harvard-MIT Graduate Student Conference. Before that, it was hotly discussed (along with untold others) with my office-mate, Adam Sennet. I cannot fail to include, on the list of unofficial committee members, Maria Bittner, from whom I learned many of my most philosophical lessons about language. Others who devoted their time to this project (and to no small effect) were Roger Schwarzschild, Mark Baker and Hans Kamp.

Thanks are due to everyone in the vibrant inter-disciplinary language community, who made Rutgers a wonderful place to reach intellectual maturity: Daniel Altshuler, Luvelt Anderson, Josh Armstrong, Veneeta Dayal, David DeVault, Carlos Fasola, Gabe Greenberg, Michael Johnson, Slavica Kochovska, Karson Kovakovich, Karen Lewis, Sarah Murray, Ángel Pinillos, Ted Sider, Will Starr and Jason Turner. Thanks to Will and Sarah, too, for the use of their apartment, in which much of this was written.

Many thanks to the questioners and commentators at presentations on various portions of this material: Joseph Almog, Louise Antony, Nicholas Asher, Chris Barker, Nuel Belnap, Selim Berker, Daniel Bonevac, Robert Brandom, Phillip Bricker, Justin Broackes, Tyler Burge, John Campbell, Herman Cappelen, John Carriero, David Chalmers, Shamik Dasgupta, Josh Dever, Cian Dorr, James Dreier, Antony Eagle, Kenny Easwaran, Andy Egan, Frankie Egan, Jerry Fodor, Thony Gillies, Hannah Ginsborg, Michael Glanzberg, Peter Godfrey-Smith, Warren Goldfarb, Anil Gupta, Ned Hall, Christopher Hill, Harold Hodes, David Kaplan, Richard Kayne, Philipp Keller, Sean Kelly, Boris Kment, Peter Koellner, Peter Ludlow, John McDowell, John MacFarlane, Paolo Mancosu, Tony Martin, Bob Matthews, Line Mikkelsen, Jessica Moss, Reinhard Muskens, Bernhard Nickel, Alva Noë, Terry Parsons, Paul Pietroski, Graham Priest, Jim Pryor, Greg Restall, Thomas Ricketts, Sherrilyn Roush, Jonathan Schaffer, Joshua Schechter, Philippe Schlenker, John Searle, Kieran Setiya, Susanna Siegel, Nicholas Silins, Hans Sluga, Barry Stroud, Eric Swanson, Anna Szabolcsi, Barry Taylor, Rich Thomason, Brian Weatherson, Mark Wilson, Dean Zimmermann and Eytan Zweig.

Thanks to my parents, who after gestating and rearing me, were selfless enough to leave me to my own devices – even when these led me abroad – and never asked for anything in return (except this acknowledgement).

Finally, thanks to Jessica Rett, for copious emotional support, for intellectual scrutiny, and for being a babe. . .

Table of Contents

Abstract	ii
Acknowledgements	iii
Introduction	1
1 Proper Names	5
1.1 Name-Spotting	5
1.2 Names in Discourse	9
1.3 Irregulars	19
1.4 Meaning on the Fly	23
2 Context	29
2.1 Semantics: Subjective vs. Objective	30
2.2 Linguistic Context and Dynamic Meaning	33
2.3 DRT, a Study in Subjectivism	37
2.4 Context as Objective Representation	38
2.5 The Nature of Discourse Referents	40
2.6 Names: Denotation	42
2.7 Names: Reference	44
3 Content	47
3.1 Referential Isomorphism	48
3.2 Loar's Counterexample	49
3.3 Denotational Isomorphism	53
3.4 Transmission of Content	55
3.5 Diagonalization	57
4 Attitudes	63
4.1 Attitudes and Truth	64
4.2 The Semantics of Attitude Reports	71

4.3	Indefinites and Attitude Reports	73
4.4	The Dynamics of Attitude Reports	74
4.5	Hob and Nob	77
4.6	Paderewski	79
4.7	<i>De Re</i> and <i>De Dicto</i>	81
4.8	Previous Accounts	86
	Appendix	92
	Curriculum Vitae	108

Introduction

Semantics is the study of the meanings of words, and how they interlock to form the meanings of larger units. When we ask what a word means or how meanings combine in natural language, our answer is often guided by analogy with a simpler language where semantics already has traction.

For instance, proper names (e.g. *Ernest*) are often classed with constant terms in mathematics (e.g. 0), while pronouns (e.g. *he*) are associated with variables (e.g. x). The argument for this is elementary. Whereas proper names and constants have only one kind of occurrence in a sentence (respectively, formula), pronouns, like variables, have two kinds of occurrence. They can appear bound by a quantifier:

- (1) a. Every physician healed himself
 b. For all x , $x + 0 = x$

And may also appear ‘free’ or unbound:

- (2) a. She healed herself
 b. $x + 0 = x$

Proper names and constants, it would seem, are not bindable:

- (3) a. Every physician healed Ernest
 b. For all x , $0 + 0 = 0$

In the elementary comparison above, we only considered units of language as large

as a sentence. Language, of course, comes in bigger parcels: paragraphs, chapters, trilogies and, in spoken form, dialogues. Parcels of sentences that are in some sense coherent¹ are called *discourses*.

As soon as we admit discourses into the comparison, it breaks down. The following mini-discourse includes an example of a ‘bound’ name:

- (4) Tampa was home to a serial killer named Bobby Joe Long. Long was known as ‘the Classified-Ad Rapist’.

In (4), the indefinite expression *a serial killer named Bobby Joe Long* (which *mentions* the name *Bobby Joe Long*) appears to bind the occurrence of *Long* in the second sentence. This is not binding in the usual syntactic sense, which would require the name to occur in the same sentence as its antecedent, but rather in the *dynamic* sense pioneered by Kamp (1981) and Heim (1982).

The indefinite *licenses* the use of the name. Without it, the name sounds odd and the discourse loses coherence (i.e. it becomes less of a discourse):

- (5) Tampa was home to serial killers. Long was known as ‘the Classified-Ad Rapist’.

Others before me (e.g. Geurts, 1997) have claimed that names are anaphoric expressions which need to be licensed. I differ from my predecessors in how I treat *discourse-initial* names. According to the anaphoric theory, a name occurring without an antecedent in the discourse is anomalous, yet examples of such names abound. For instance, if we have a friend in common called Joachim, I can open a conversation by saying: ‘So, Joachim dropped by today’. Indeed, it would be distinctly odd for me to introduce Joachim to you with an indefinite.

Previous accounts have suggested that discourse-initial names are licensed by a

¹For instance, the medley consisting of the opening sentence of *Anna Karenina*, the penultimate sentence of Pericles’ funeral oration and the first line of Rilke’s Duino Elegies is not a discourse; although quotations compiled in a meaningful way, like Walter Benjamin’s *Passagenwerk*, could be considered one.

procedure known as accommodation (Lewis, 1979b). That is to say, more-or-less, that such names are made acceptable by our behaving as though they already are – by our disposition to ignore the absence of an explicit antecedent. This explanation sounds highly artificial. For one thing, it implies that we would accept the insertion of an antecedent if one were offered,² but that just doesn't ring true in the Joachim case.

Instead, I suggest that the antecedents for discourse-initial names occur in *previous* discourses. A name will have several 'antecedents' in this sense, since the participants in the discourse are likely to have been introduced to the name under different circumstances. Properly speaking, the 'antecedent' for a name is not an indefinite expression – it is not an expression at all, but rather a *tree* of anaphorically connected utterances. The utterance of a name is felicitous in a context if and only if there is such a tree to serve as its antecedent. A tree will serve, just as an indefinite would, so long as the participants in the conversation are able to single it out as the tree to which the utterance belongs.

While the study of names in discourse is new, names have long been a topic in the philosophy of language. A number of well-known philosophical puzzles accompany the attempt to combine a referential semantics for names with a theory of content. Famously, the notion of 'singular' content has a host of intuitive drawbacks. The singular content of the sentence *Hesperus is Phosphorus* is trivial, and yet the sentence is *a posteriori*. Attitude contexts are not 'Shakespearean' (Kripke's phrase), since they can't stand (without adulteration of their logical properties) the substitution of co-referring names.

I build on the discourse semantics of names developed in Chapters 1 and 2 to develop an alternative notion of content that is sensitive to discourse relations. Since I allow anaphoric relations *between* discourses, in addition to the usual discourse-

²Geurts (1997) claims that names prefer binding to global accommodation.

internal relations, my version of discourse content is transmittable between contexts. Furthermore, discourse content is strictly more fine-grained than singular content, and as a result is a better fit for our pre-theoretic intuitions about communication and attitude-reporting (as brought out by Frege's puzzle).

Attitude reports, on my account, relate agents to discourse-sensitive content. Thus I predict that the discourse properties of the subordinate clause in the reporting sentence are mirrored in the state of affairs reported on. In particular, anaphora between reports, including attributions to different agents, must be replicated in the world (by anaphoric relations between *contexts*) for the reports to be true. This anaphoric connection between different agents is precisely Geach's notion of 'intentional identity', which he judges to be a condition on anaphorically-connected attitude sentences (Geach, 1967). The present account derives this judgment from an independently motivated notion of content.

Chapter 1

Proper Names

In this chapter, I survey the empirical landscape of proper names (and their component proper nouns), with a focus on their behaviour in discourse. I provide a brief round-up of the syntax of names, before introducing the reader to the study of discourse.

The central claim of the chapter is that names are anaphoric ('definite') expressions with attentional demands in between those of pronouns and definite descriptions. This claim, and the foundations of discourse semantics, will be refined further in the next chapter.

In the second half of the chapter, I argue that the semantic contribution of a proper name is basic rather than composite. In particular, the proper *noun* constituents of proper names are not meaningful. Rather, the *form* of the proper noun aids in the interpretation of the name.

1.1 Name-Spotting

The Stoics are credited with establishing the proper name as a part of speech. Previously, Plato had used the same word – *onoma* – to cover both proper names and common nouns. Dionysius Thrax (170–90 BCE) was the first to taxonomically dis-

tinguish proper names (*onoma kyrion*) from common nouns (*onoma prosegorikon*). This practice was copied by Roman grammarians and transmitted to modern times. According to Algeo (1973, 2), the distinction was ‘based on the semantic contrast between “individual quality” and “common quality”’. Paradigmatically, a proper name applies uniquely – to the possessor of a particular ‘individual quality’ – while a common noun applies to all who share some common quality.¹

The medievals worried about names that applied to more than one person (e.g. *Zeno*). Do such names signify a common, rather than an individual quality (jeopardizing the semantic basis for the grammatical category)? Most thought: no, such names still signify individual qualities; they just signify more than one – they are ambiguous.² Kaplan (1989) suggests that what seem to be different uses of a name are in fact uses of different names; there is a stock of names spelled *Z-e-n-o*, each of which is univocal (and individually applicative).

There is no doubt that names, in context, make a particular reference. Common nouns, while they can also be used to make a particular reference, always do so *through* a general property, or ‘role’.³ This indirect method of reference has side effects. For example, the use of the noun *thief* in (2), by evoking a role with certain propensities and motivations, efficiently explains the action described by the verb:⁴

(2) The thief smashed the passenger window.

Payne and Huddleston (2002) make a further distinction between proper *names* and proper *nouns*. A proper noun is a word-level unit of the category *noun* (ibid.,

¹Also see (Husserl, 1901), who replaces ‘individual quality’ with ‘monovalent intension’ and ‘common quality’ with ‘polyvalent intension’.

²In his commentary on the *Categories*, Paul of Venice espouses the opposite view, that (some) names signify a common quality (Ashworth, 2006).

³This distinction is not exceptionless, since we can also name roles (e.g. *Miss Rhode Island*).

⁴Or consider the exclamation (1) which, but for the consideration of the role used to pick out the referent, presents an entirely redundant proposition (since the addressee presumably knows the identity of his topic perfectly well).

(1) That’s my sister you’re talking about!

516). Proper names are noun phrases (syntagms).⁵ For instance, the proper name *Jessica Alba* consists of two proper nouns: *Jessica* and *Alba*. Proper names can consist of other parts of speech, too. For instance, *Brooklyn Bridge* contains the common noun *Bridge* in addition to the proper noun *Brooklyn*. *The Raritan River* also contains the determiner *the*. *The Bronx* combines a determiner and a proper noun. Finally, *the Golden Gate Bridge* has no proper nouns at all.

There is no limit on the forms proper names can take. For instance, personal names in Mohawk are verb phrases. But we don't need to look beyond English to find unusual forms like the brand name *I Can't Believe It's Not Butter!* and the band name *the Yeah Yeah Yeahs*. Since any string of words might in principle be a name, the internal grammar of names would appear exceedingly liberal. At first blush.⁶

$$\textit{proper name} ::= \textit{word}+$$

This grammar treats names as 'flat' strings of words. However certain names are clearly built up from more complex expressions – including other names. For example, *Brooklyn Bridge* is composed of the noun *bridge* together with the borough name *Brooklyn*. *The New York City Public Library* is built up from *public library* and *New York City* – itself built from *city* and *New York*.

Moreover, the string *the New New York Public Library* is structurally ambiguous. It might be parsed as a name for a revamped public library in New York – *the New (New York) Public Library* – or else as a name for an old library in a futuristic city – *the (New New York) Public Library*.

I will therefore suppose that names have a hierarchical grammar. Indeed, there is not one grammar of proper names, but a great number of them (Carroll, 1985). For instance, official names of persons in most Western cultures consist of (at least) first and last names. It is usually noticeable if someone has a first name for a surname or

⁵In addition to occurring as arguments, proper names can also modify other nouns, as when *Howard* is used as a modifier in *a Howard spokesperson*. For more unusual uses, see §1.3.

⁶Cf. McDonald (1996, 34)

vice versa, so I assume these come from different, overlapping allotments:

$$\textit{person-nominal} ::= \textit{first-noun} + \textit{sur-noun}$$

$$\textit{person-name} ::= (\textit{title}) \textit{person-nominal} (\textit{coda}) \mid (\textit{title}) \textit{sur-noun} \mid \textit{first-noun}$$

The names for landforms, waterways, structures and institutions have a standard blueprint. Such schemas typically include a common noun, indicating the semantic type, and a ‘proper nominal’ (either a single noun or a complex expression) indicating the particular member of that type. The blueprint may place them in either order (*Treasure Island*, *Cape Horn*) and may or may not include a definite determiner.

Here is a sample grammar for hotel names:

$$\textit{hotel-nominal} ::= \textit{proper-nominal} \textbf{Hotel}$$

$$\textit{hotel-name} ::= \textbf{the} \textit{hotel-nominal} \mid \textbf{the} \textit{proper-nominal}$$

Titles (e.g. of television shows) follow a flat paradigm:

$$\textit{title-nominal-1} ::= \textit{word} +$$

$$\textit{title-1} ::= \textit{title-nominal-1}$$

Since the common noun is semantically active in hotel names (e.g. the Tribeca Grand Hotel is a hotel), but not in titles (e.g. The Wire is not a wire), we might be tempted to conclude that the former are not, after all, names, but are in fact definite descriptions *modified by names* (like *the Howard spokesperson*). If this were true, we could maintain a flat architecture for true proper names. However, there is a semantic difference between a complex name like *the Tribeca Grand Hotel* and a name-modified noun phrase like *the Howard spokesperson* – only the latter will readily designate a role, as opposed to an individual (Kripke, 1972):

(3) Today, the Howard spokesperson was a woman.

Some names are shortenings or ‘reduced forms’ of others. For instance *Richelieu* is a reduced form of *Armand Jean du Plessis de Richelieu*. I assume the syntax includes a *reduction relation* ‘▷’ defined over names. For instance, a full personal name can reduce to any of its first nouns as well as its sur-noun (the latter with or without the title):

$$[\text{NP } (title) [\text{Nom } first-noun_1 \dots first-noun_n \text{ sur-noun }] (coda)] \triangleright [\text{NP } [\text{Nom } first-noun_i]]$$

$$[\text{NP } (title) [\text{Nom } first-nouns \text{ sur-noun }] (coda)] \triangleright [\text{NP } (title) [\text{Nom } sur-noun]]$$

In many cases (and exemplified by hotel names below), we arrive at the reduced form of a complex name by removing the common noun:

$$[\text{NP } the [\text{Nom } proper-nominal \text{ Hotel}]] \triangleright [\text{NP } the \text{ proper-nominal}]$$

1.2 Names in Discourse

Names have a further property which we may provisionally (while confined to English) call *definiteness*.⁷ Definiteness, as I am using the term, is a property of expression types that constrains their use in discourse. As a first approximation, definite expressions require identifiable (or ‘hearer-old’) referents (Lambrecht, 1994; Prince, 1979). Other definite categories are: definite descriptions, demonstratives and pronouns. Indefinite expressions include noun phrases beginning with: the indefinite article, the determiner *some*, any numerical determiner, and a special colloquial use of *this*.⁸

Among languages in possession of a definite article, the categorization of names as definites is supported by (oft-rehearsed) evidence.⁹ As we have just seen, whole

⁷The terminology is parochial because in other languages ‘definiteness’ can be conferred by ‘word order, the presence or absence of a numeral, a case-marking particle’ and other means (Lambrecht, 1994, 79). Judging from this, it is a solecism to attribute ‘definiteness’ to expression types. Nevertheless, I will persist in this familiar usage.

⁸A proto-categorization along these lines can be found in Varro’s *De Lingua Latina*, a notable exception to the classical pattern presented earlier (Algeo, 1973, 3). Varro groups proper names (*nomina*) together with ‘this’ and ‘that’ (*pronomina*), and apart from common nouns (*vocabula*) and ‘who’ and ‘what’ (*provocabula*).

⁹For instance: Geurts (1997); Van Langendonck (1998, 154–157).

classes of English names are formed by attaching a definite article (the nonpareil of definiteness markers) to a ‘proper nominal’. While personal names do not appear with a definite article in English, they do in Modern Greek, European Portuguese, Bavarian German and Pima, among others (Matushansky, 2006b).

While certain titles begin with an indefinite article – *A Midsummer Night’s Dream*, *An Inconvenient Truth* – I assume the article is a component of the (flat) proper nominal, as evidenced by such constructions as *the An Inconvenient Truth Movie Forum*.¹⁰ Other cases, like *An Engelbert Santana to see you*, will be discussed in §1.3–1.4.

To call names definite is not yet to indulge in descriptivism.¹¹ Descriptivism, the thesis that names function semantically like definite descriptions, was Russell’s answer to certain puzzles in epistemology and metaphysics. From the discourse perspective, definite descriptions denote objects that are familiar (the denotational semantics is quite different from Russell’s famous account),¹² and in this sense they are similar to names. However, they differ from names in their capacity to designate *roles*, as mentioned earlier.

Names contrast with indefinite noun phrases. Russell (1905) analysed indefinites as existential quantifiers, restricted by their nominal component. Indeed, Frege’s invention of the existential quantifier condensed complex syllogistic patterns involving premises with indefinite subjects.

Consider the following (invalid) argument:

$$(4) \quad \frac{\text{A man is big} \quad \text{A man is small}}{\text{A man is big and small}}$$

We can formalize the argument, using Russell’s analysis of the indefinite, as follows:

¹⁰www.rottentomatoes.com

¹¹Abbott (2002).

¹²Russell analysed ‘The *F* is *G*’ as ‘Exactly one thing is *F*, and it is *G*’. Note that the quantificational subject in the analysis is an indefinite NP, as it begins with a (complex) numerical determiner.

$$\frac{\exists x[big(x)] \quad \exists x[small(x)]}{\exists x[big(x) \wedge small(x)]}$$

We are then free to confirm its invalidity by standard means. We could find a countermodel, or we could construct a proof in a proof system (monadic predicate logic is decidable).

While the countermodel method relies on Russell’s semantics, the proof method does not. For it is conceivable that the indefinite article has the same introduction and elimination rules as the existential quantifier, but a different semantics:

$$\frac{F(n) \quad G(n)}{\text{An } F \text{ is } G} \text{ a-intro} \qquad \frac{\text{An } F \text{ is } G \quad 'c' \text{ is new}}{F(c) \wedge G(c)} \text{ a-elim}$$

The rule that explains the invalidity of (4) is (a-elim), which requires the instantiating term to be *new* to the derivation. While the standard model theory for the existential justifies this rule, it is not the only justification possible. In fact, we also find justification in the *discourse properties* of indefinites.

The familiarity theory of definiteness (Miklosich, 1874; Christophersen, 1939) states that an indefinite NP must refer to an individual that is new to the discourse, while a definite NP must refer to one that is already familiar. The familiarity theory informs the OED entry for *a*:

...It is especially used in first introducing an object to notice, which object, after being introduced by *a*, is kept in view by *the*; as ‘I plucked a flower; this is the flower.’

The familiarity theory can account for the invalidity of (4) – on the assumption that we treat the premises as we would a coherent discourse.¹³ From the fact that indefinites introduce a new referent, we can infer that the two occurrences of *a man* in the premises of (4) do not co-refer. The argument is therefore invalid.

An important wrinkle was introduced to the familiarity theory by Karttunen (1976). He proposed that a form of the generalization was sustained even where there was no particular reference.

¹³Chastain (1975), 207.

(5) If you find an error, make a note of it.

The speaker of (5) is not talking about a particular error – the sentence has a universal interpretation. It is nevertheless true that there is a link between the indefinite *an error* and the pronoun *it*. Karttunen coined *discourse referent* (*dref*) as a term for the ‘thing’ that, according to the familiarity theory, is introduced by the former and retrieved by the latter. It is not (in the case above) a particular error, but rather a particular variable, or *repository* of errors (Muskins, 1996). Discourse referents will be one of the topics of Chapter 2.

Definite expressions, though all require familiar referents, are not all of a piece. The space of attention is structured into a ‘givenness hierarchy’ (Gundel et al., 1993).¹⁴ The bottom rung is occupied by *uniquely identifiable* discourse referents: those that can be singled out (in context) by their role. This status is sufficient for retrieval by definite description (the nominal of the description provides the identifying role). The next rung up holds the *activated* drefs – those represented in the ‘short-term memory’ of the discourse. This status is sufficient for retrieval by a demonstrative (while mere unique identifiability is not). Certain public events, either linguistic (a previous mention) or extra-linguistic, bestow activated status on a dref. Finally, there is a stack of *in-focus* drefs: those at the current centre of attention. Pronouns select the dref at the top of this stack. Prominent mention of a dref – for instance from the subject position of a sentence – will place it on the top of the *in-focus* stack.

$$\underbrace{\text{in-focus} > \text{activated} > \text{uniquely-identifiable}}_{\text{familiar}}$$

Gundel et al. omit proper names from their hierarchy. Indeed, the attentional requirements of names are complex. They are sensitive to genre and register and, moreover, to the form of the name (full or reduced). As a first pass, the full form is

¹⁴I am simplifying their account somewhat.

used to introduce a name into the discourse, while the reduced form is adopted only after the name has been introduced.

On its first appearance in a news article, a name assumes its most dilated form. It is usually ramified with appositive modification, or else preceded by an anarthrous noun phrase (termed an ‘embellishment’ by Payne and Huddleston):¹⁵

- U.S. President George W. Bush
- Professional naked cowboy Robert John Burck
- Jessica Rett, a prominent fashion designer
- the Tribeca Grand Hotel

Another way to introduce a name is with a naming construction (Hahn, 1969; Matushansky, 2006a):

- a convicted serial killer in Florida named Bobby Joe Long

In subsequent occurrences, a reduced form of the name is used (Mani et al., 1993, 45):

- President Bush, Mr Bush, Bush
- Burck
- Rett
- the Tribeca Grand

The name-plus-indefinite-appositive combination serves both (i) to introduce a new discourse referent and (ii) to allow subsequent (reduced) occurrences of the name to retrieve that discourse referent. A name-plus-definite-appositive combo serves to

¹⁵See also: <http://itre.cis.upenn.edu/~myl/language-log/archives/001628.html>

(i) uniquely identify a familiar discourse referent and (ii) associate the name with that discourse referent.¹⁶

Full names inherit their discourse status from the NPs in their entourage. Reduced names, which have no entourage, denote discourse referents that have already been activated *by prior mention of the name* in the discourse. This makes them difficult to place on a linear hierarchy. Stretching the paradigm of Gundel et al. (1993), we might add *in parallel* a discourse status for each name α (call it α -activated), where a discourse referent that is α -activated has status sufficient to be picked out by the name α . A full-name-plus-NP combination that denotes a dref u bestows α -activated status on u , for every α that is a reduction of the full name. For instance, use of the phrase *U.S. President George W. Bush* – denoting u_{gwb} – updates the context to include *Bush-activated*(u_{gwb}) and *George-activated*(u_{gwb}).

Note that reduced names are *ceteris paribus* preferred to definite descriptions:

- U.S. President George W. Bush ... Bush/the president
- Customer Chris Roth ... Roth/??the customer
- Jessica Rett, senior economist at the World Bank ... Rett/??the economist

This suggests that α -activated $>$ *uniquely-identifiable*, for each name, α . When available, a name is usually preferred to a demonstrative. However, I do not conclude that α -activated $>$ *activated*. Once again, the situation is complex. It seems a

¹⁶From Mani et al. (1993, 45):

Firstly, we hypothesize that for certain genres of text (for example, Wall Street Journal news stories), new references are introduced by information occurring in the immediate syntactic environment of the proper name. (What the precise set of such genres is remains to be determined, but our initial set includes the most common forms of news stories and excludes literary narratives.) Many of these local contextual clues reflect felicity conventions for introducing new names. New names of people (as well as organization names, and to some extent location names) are generally accompanied by honorifics and various appositive phrases which help anchor the new name reference to mutually assumed knowledge.

demonstrative can take the place of a reduced name only if the utterer wishes to distance him- or herself from the individual referred to. This is explained by the fact that names convey a degree of empathy, a consideration orthogonal to attentional status.¹⁷

$$\text{in-focus} > \left\{ \begin{array}{l} \text{name}_1\text{-activated} \\ \text{name}_2\text{-activated} \\ \dots \\ \text{activated} \end{array} \right\} > \text{uniquely-identifiable}$$

Names behave differently in dialogue. Firstly, unfamiliar names may be introduced in reduced form (*This is George*). Second, if the name is familiar to the participants in the dialogue, then the discourse-initial situation is the same as the post-introduction situation in formal genres. The reduced form may be used without preamble – so long as there aren't other, equally salient persons who go by the same name.

For instance, if you and I know several Georges, none of whom is particularly salient, then my use of the reduced name *George* is infelicitous. You can't tell which one I'm referring to. I must batten the reduced form with extra proper nouns, or suitable descriptions, until my utterance becomes interpretable. Afterwards, I can revert to the reduced form, since I will have raised the salience of that interpretation of the name.¹⁸

Thus I must use the full form *Gareth Evans* in a dialogue situation where there are equally salient Gareths, and I must use the phrase *Gareth Evans the philosopher*,¹⁹ when there is more than one salient individual by the name of *Gareth Evans*.

According to our theory, a reduced name α can only be used to denote a familiar discourse referent with (at least) the status α -activated. The only way to get that

¹⁷Relatedly, the conventions of storytelling dictate that principal characters have names while secondary characters are referred to by a definite description identifying their role (Sanford and Lockhart, 1990).

¹⁸Note that the same situation could arise in a written document (for instance, a personal memoir about George Orwell and George Sand).

¹⁹Note that, unlike introductions, here the descriptive component is in 'close' apposition with the name (i.e. no intonation break or punctuation mark intervenes).

status is by an earlier mention with the name's full form. As it stands, then, the theory predicts that reduced discourse-initial names are not possible.

One way to resuscitate the theory is by appeal to accommodation. Lewis (1979b) writes:

There is one big difference between baseball score and conversational score. Suppose the batter walks to first base after only three balls. His behavior would be correct play if there were four balls rather than three. That's just too bad – his behavior does not at all make it the case that there *are* four balls and his behavior *is* correct. Baseball has no rule of accommodation to the effect that if a fourth ball is required to make correct the play that occurs, then the very fact suffices to change the score so that straightaway there are four balls.

Language games are different. As I hope my examples will show, conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play.

The idea would be that, so long as the participants acquiesce, the speaker may behave as though he has already suitably introduced his names. Clearly, if they are not familiar with a particular name, or if there is potentially more than one person he could mean by the name, the participants are unlikely to acquiesce, and he is left holding a presupposition violation.

As I understand him, Geurts (1997) advocates this approach. He assumes that the familiarity presupposition carried by a name will be accommodated, so long as it doesn't clash with the beliefs of the audience. Since the presupposition, for him, is simply that someone of that name exists, such a clash is unlikely. If the audience is already familiar with a person of that name, then the dref about which the speaker has expressed a commitment may be linked (by a 'counterpart relation') to this familiar dref. However, the felicity of the utterance does not seem, on Geurts' theory, to ride on this.

The accommodation theory makes an unfortunate prediction. Suppose you and I have a friend in common, called Grantham; and suppose I begin a conversation by saying to you, 'So, I saw Grantham the other day.' According to the theory, you

merely acquiesce in my use of the reduced, un-introduced name. I could have avoided accommodation by saying instead, ‘So, I saw a friend of ours called Grantham’, or ‘So I saw Grantham, a friend of ours.’ But, in spite of what the theory says, such utterances are infelicitous.

Instead of resorting to accommodation, my approach is to reject the assumption that discourses are ‘insular’. That is to say, that the anaphoric presuppositions on expressions in a discourse must either be satisfied in that discourse or else be accommodated. I want to broach the alternative view that discourses are connected by anaphora to other discourses, that anaphoric expressions in one discourse might find their antecedents in an earlier one.

Allow me to explain. Suppose the limits of a (dialogical) discourse were defined more-or-less by temporal continuity. Thus, two workmates might have a conversation over lunch one day (constituting one discourse), and then continue the conversation (while commencing a new discourse) at lunch the next day. We can even imagine an individual being introduced in the first discourse (e.g. a new love-interest) and becoming salient enough to be retrieved by a pronoun at the beginning of the second discourse (‘Did he call?’). Now, a natural way of describing this would be to say that the pronoun at the start of the second discourse is anaphoric to some indefinite expression embedded in the first. However, this explanation is impossible if discourse boundaries are impervious to anaphora.

The argument for trans-discourse anaphora does not rely on any controversial account of the size and shape of discourses. If Grantham is a mutual friend, then I can use his name around you, even though we never had a conversation in which his name was first introduced. We might, after all, have known Grantham before we knew each other, and so our respective ‘introduction’ events were different. One couldn’t, in this case, extend the boundaries of the discourse in any natural way to include an introduction event for the name. For one thing, whose introduction would we choose, mine or yours?

I nevertheless persist in calling names anaphoric, and in eschewing accommodation. As I would have it, my use of the name ‘Grantham’ is licensed, in part, by the fact that you (my audience) have previously been introduced to the name (in this discourse, or an earlier one). The name is familiar to you; it is activated in your ‘long-term’ memory. This same fact explains why it would be infelicitous for me to re-introduce the name with an indefinite, even at the start of a new discourse. The explanation is the usual one: indefinites must denote brand-new discourse referents, and cannot refer to familiar ones (including those in long-term memory).

Having a name conjure up a familiar dref in long-term memory is necessary but not sufficient for a felicitous utterance. The utterance must also confine the audience to a unique interpretation (as we saw above, *George* is infelicitous if there are multiple, equally salient Georges).

There is one further condition. Suppose I utter the name *Gareth Evans*, intending to refer to the late British philosopher. Let’s suppose further that you have never heard of the philosopher, though you do know an Australian politician by that name (and no-one else). I have produced a name for which you have a unique interpretation, but clearly something has gone wrong. Even if the discourse went smoothly after that (my preposterous claims about the ex-treasurer went unchallenged), there was, in an objective sense, something infelicitous about my utterance.

If we trace anaphoric paths back, one from the use of the name as I intended it, and another from the use of the name as you construed it, they will not meet. My path traces back to the person from whom I first heard of the philosopher, back further to the place he heard it from, and so on. Your path, running back through your introduction to the name, wends its way to a quite different origin. Starting from the other end – the baptisms of Gareth Evans and Gareth J. Evans – we can trace out distinct ‘anaphoric’ trees by following the transmission of each name through different discourses.

The third and final felicity condition on the use of names may now be put as

follows: the name in the long-term memory of the hearer must be connected by a path on an anaphoric tree to the name as it appears in the long-term memory of the speaker. In fact, we can compose a necessary and sufficient felicity condition thus: the utterance of the name must enable the audience to (uniquely) determine the anaphoric tree to which the utterance belongs.²⁰

Coordination on an anaphoric tree – as required for the successful interpretation of a name – is a generalization of the notion of coordination on an antecedent. The latter is the requirement, as it is usually stated, on the successful interpretation of a pronoun (the ‘tree’ in this case is just a linear process with its origin in the same discourse). In fact, there is a strong affinity between names and pronouns – although names, being more numerous, are specialized for ultra-long-distance anaphora, while pronouns, which have stronger attentional requirements, are consigned to the short-distance niche.

In this section, I have established that names are definite expressions with attentional requirements one notch lower than pronouns, and that a name is felicitous whenever it enables coordination on an anaphoric tree.

1.3 Irregulars

TWEEN GIRL: Duran Duran was not a Duran
or a Duran, so just cease and desist.

– *Borders near Penn Station*

CHICK ON CELL: I always end up dating people
whose names aren’t actually their names.

– *Harlem*

www.overheardinnewyork.com

In the first half of the chapter, proper names were analysed as anaphoric expressions

²⁰I am supposing that the anaphoric heritage of an utterance is determined by the utterer’s anaphoric intention (Grice, 1957).

of a certain sort. In this section, we will look at some unusual uses of proper names and proper nouns.

Burge (1973) provides examples in which proper names seem to function as common nouns. For instance, they can appear with an indefinite determiner:

- (6) a. He looks like an Alfred
 b. An Alfred Russell is here to see you, Sir

Sentence (6a) could be accounted for without revising our account of proper names, since we could argue that *Alfred* is a proper *noun* (which may combine with determiners just like other nouns). In (6b), however, we surely have a personal name, rather than a noun, combining with the indefinite article. Burge notes that in both examples, the name could be replaced by the nominal expression *person called Alfred (Russell)* while retaining much the same meaning.

Note that the metalinguistic interpretation is delicately connected with the configuration. The names in the examples above cannot simply be exchanged:

- (7) a. He looks like an *(older) Alfred Russell
 b. ?An Alfred is here to see you

Alfred North Whitehead is *literally* – Burge contends – an Alfred. But is he also ‘literally’ an Alfred (North) Whitehead? Proper names, as opposed to proper nouns, often sound strained in predicate position, unless they take on a metaphorical or metonymic interpretation:

- (8) a. That’s not an El Greco!
 b. You’re no Jack Kennedy
 c. She’s the next Grace Kelly
 d. It out-Herods Herod
 e. Please do a Napoleon for the camera

Outside of predicate position, it is acceptable to use a proper name in a metalinguistic sense. Poring over a telephone directory, I announce:

(9) There are two Alfred Whiteheads in Manhattan

Sentence (6b) appears to be a different sort of use to (9), however. It serves to introduce a new discourse referent, which it simultaneously associates with the name *Alfred Russell* (enabling later retrieval of that referent by that name). Earlier, we saw this same feat accomplished by a metalinguistic indefinite (*a serial killer named Bobby Joe Long*), but also by an argumental name in appositive construction with an indefinite NP (*Jessica Rett, a prominent fashion designer*). We need not assume that the name in (6b) contributes a metalinguistic property. Instead, we may assume it, like the argumental name, has its usual function of referring to a discourse referent, which the indefinite article marks as new to the discourse. This analysis can account for the infelicity of (7b), since the introductory occurrence of a name should take its full form (in a formal setting).

It has been popular, since Burge, to take the metalinguistic interpretation as basic. Burge, and others in league with him, have had little to say about examples (8a)-(8d), and usually set them aside for separate treatment, emphasising that they are not literal uses. Burge himself doesn't distinguish proper names from proper nouns (his article precedes Payne and Huddleston's by three decades). Unfortunately, apart from (6b), his examples all involve proper nouns. If we take (6b) out of the equation (as an argumental name marked for indefiniteness), then he doesn't marshal *any* evidence of proper names occurring as predicates.

Matushansky (2006b), who is alive to the name-noun distinction, gives a metalinguistic analysis of proper nouns. She generates metalinguistic interpretations of personal proper names – when these are combinations of proper nouns, as in (9) – by *intersecting* the metalinguistic denotations of the constituent nouns. Thus an Alfred

Whitehead is, for Matushansky, one who is both an Alfred and a Whitehead.²¹

The analysis extends to other name-types; Matushansky treats the semantic interaction between proper and common noun denotations once more as intersection. So a Raritan River is a river that is a Raritan (and thus *the* Raritan River is the unique river that is a Raritan). Nevertheless, this sounds strained. The Raritan River is not obviously a Raritan (nor is it quite correct to say that it is called Raritan - it is called *the* Raritan!). A George Washington Bridge would then be a bridge that is both a George and a Washington, but I doubt very much that such a bridge is ever denominated *George*. Worse still, *the Golden Gate Bridge*, which contains no proper nouns, receives an absurd literal interpretation: the bridge that is a golden gate. Matushansky has fallen into a trap as old as Plato's *Cratylus*.

Let's return to the Raritan River. It is called *the Raritan River* and also *the Raritan* (these are its full and reduced names). It is not, however, *a* the Raritan (River). I can speak of two Raritan Rivers, and I suppose I can say that each of those is a Raritan River (one of many).²² This seems to indicate (as described in the syntax given earlier) that the determiner couples with the proper nominal, *Raritan River*. It might be preferable, then, to treat the proper nominal as a semantic atom (with a metalinguistic interpretation), rather than attempting to distribute its meaning among its syntactic daughters. For instance, any object with the name *George Washington Bridge* would count as a George Washington Bridge. A Golden Gate Bridge would be something having the title *Golden Gate Bridge*.

Should we take the next step, and analyze the proper name *the Raritan River* compositionally – as combining the meaning of the definite article and the proper nominal meaning (a metalinguistic property)?²³ There is a problem: the referent of the name, on this analysis, is not required to be a river (it is merely required to be

²¹The fact that first-nouns and sur-nouns overlap means that this analysis is too weak as it stands, for it will follow (incorrectly) that every Alfred Russell is a Russell Alfred. Moreover, someone who goes by two different first names, say *Donald* and *Tony*, will absurdly count as a Donald Tony.

²²I can go further and say: *There's a Raritan on this map, but it's the avenue, not the river.*

²³Matushansky (2006b); Elbourne (2005); Larson and Segal (1995).

the unique object called *Raritan River*). Since *Raritan River* is a passable name for a band or a TV show, the analysis fails to capture our earlier claim that the common noun was semantically active in names of this type.

The puzzle for the compositional analysis can be posed thus: *the Yellow River* (*qua* river-name) must denote a river.²⁴ Yet one can speak of there being two Yellow Rivers, one of them a waterway and the other a concerto. Where does the restriction to rivers come from in the first case, if absent from the second?²⁵

There is a further problem. In many cases where a proper nominal occurs as a modifier, it is not interpreted metalinguistically (neither is it interpreted metaphorically). A Howard spokesperson is not a spokesperson named Howard, but a spokesperson *for* Australian Prime Minister John Howard.

1.4 Meaning on the Fly

I propose an ecumenical approach to the irregular name data. While I allow that the proper noun *Alfred* in (6a) and the proper nominal *Alfred Whitehead* in (9) contribute metalinguistic properties to the semantic calculation of those sentences, the same is not so of the name in (10):

(10) Alfred Russell is here

As I have been saying, the semantic piece contributed by an argumental name is always a (familiar) discourse referent. While the interpretation of the name surely

²⁴What I mean by this is that while an utterance of the string *the Yellow River* need not necessarily denote a river, an utterance syntactically parsed as a river-name (one can force this parsing with a conjunction: *the Yangtze and Yellow Rivers*) must be interpreted as denoting a river.

²⁵Ernie Lepore (p.c.) provides a counterexample: the Morning and Evening Stars are neither of them stars (though we can form a conjunction of their nominals). Much earlier, it was remarked that the Holy Roman Empire was none of the above. While such names are properly classed as star names and empire names, they are *misnomers*. I propose that misnomers are analogous to cases of misdescription (Donnellan, 1966; Kripke, 1977); although misnomers never have ‘semantic reference’ distinct from their ‘speaker reference’ (e.g. *the Evening Star* doesn’t semantically refer to the first *star* to appear in the evening, as opposed to the first planet), they do incur (strictly speaking) a semantic violation.

depends on its syntactic form, it does not proceed via a semantic calculation involving a metalinguistic property. That is to say, the dref denoted by *Alfred Russell* is not the result of intersecting, or otherwise combining the denotations of the proper nouns *Alfred* and *Russell*, nor is it the result of combining the denotation of the proper nominal *Alfred Russell* with a hidden definite determiner. Instead, its denotation is simple, and determined by the pragmatic operation of anaphora resolution, which takes its cue, not from the semantic values of the component proper nouns, but from their form.²⁶

I also allow the proper nominal modifier *Howard* to contribute the property of being ‘for’ Prime Minister Howard in the expression *a Howard spokesperson*, and the name *El Greco* in (8a) to contribute the property of being a painting by El Greco.

Grice had something to say against such accounts: don’t multiply meanings beyond necessity! The ecumenical account is in clear violation of this injunction. The number of possible interpretations of a humdrum name like *George* is limited only by one’s imagination (or, more precisely, by the combined imaginations of the participants in the discourse).

Before passing judgment, we should consider the motivation for Grice’s principle. Why is he so bent on reducing the semantic portfolios of words? Presumably, he is concerned about the learnability of the semantic theory. A theory of English that exceeds a certain size becomes unlearnable, leaving it open to a straightforward empirical disproof. I doubt that Grice could have put a figure, even a round one, on that size, but this doesn’t mean that, *ceteris paribus*, semantic parsimony is not a

²⁶Meaningful expressions occurring in names, such as the common noun *river* in the river name *the Raritan River*, contribute to the grounding of names (when they are first introduced) by way of their semantic properties (such words occurring in band names or titles, on the other hand, do not). This is analogous to the work of the noun in the appositive *the Raritan, a river in New Jersey*. I am less sure that the semantic value of *river* helps to retrieve the anaphora once it has been introduced, in the way *philosopher* helps in the close appositive *Gareth Evans the philosopher*. While *the Raritan River* can be used to secure unique reference in a situation where both the Raritan River and the Raritan Bridge are salient (whereas just *the Raritan* could not), it is not clear that the explanation of this is semantic, as opposed to formal. The utterance distinguishes between the two resolutions on the basis of form just as well as it does on the basis of known facts (i.e. that the Raritan River is a river, while the Raritan Bridge isn’t).

prudent strategy.

We should not forget, however, that our linguistic knowledge is topped-up with other background resources, including facts about the discourse context (Clark, 1996). It is no violation of Grice's razor that, for instance, a demonstrative contributes different denotations on different occasions of use, since these interpretations are supplied by the context, rather than an encyclopaedic lexicon.

More germane to the topic of irregular names are cases of familiar expressions occurring in an unfamiliar syntactic setting. The 'nonce' denominal verb *porched* has a clear interpretation in Clark and Clark (1979)'s example, *The newsboy porched the newspaper*, but could be interpreted differently in a different context. In such cases, instead of relying on a conventional meaning – *porched* has none – speakers and hearers collaborate to create interpretations from scratch, aided by linguistic and contextual resources.

Irregular proper names are in many ways analogous to extempore denominal verbs. The unusual syntactic configuration *do a Napoleon* (Clark and Gerrig, 1983) signals that a nonce interpretation is intended, but precisely *which* interpretation is left open by the grammar. The phrase has a different meaning in each of the following environments:

- (11) a. Please do a Napoleon for the camera
 b. Hitler did a Napoleon in Russia

Likewise, the phrase *two William Blakes* – as in: *I bought two William Blakes* – could mean two prints, two books of poetry, or two letters by Blake (Clark and Gerrig, 1983). It doesn't seem as though any one of these is *the* conventional meaning of the phrase, while the rest are non-literal or non-standard uses. Rather, it is apparent that there is *no* conventional meaning of the phrase, and that the unusual syntax signals that one is being made up on the spot (by differentiating the phrase uttered from the word that occurs in the lexicon).

This observation, while it doesn't account for such spontaneous interpretations,²⁷ at least removes the concern about learnability. These meanings, because negotiated on the spot, are not part of the semantic theory that an English speaker must learn. The use of nonce expressions taxes the hearer's creativity, not his or her rote memory.

There is another theoretical impulse, almost contrary to Grice's, that suggests we avoid, where possible, appeal to pragmatics. Semantic rules work more smoothly and are less cognitively demanding than the open-ended resources of pragmatics. In accounting for the effortless way we use language, we should (therefore) favour semantic explanations over pragmatic ones.

While I feel that this complaint sets out on the wrong foot, the exercise of following it through will prove worthwhile. Let's suppose, for the sake of argument, that among irregular uses of names the metalinguistic interpretation predominates. Following the anti-pragmatic impulse, we decide that names (which at deep syntax are nominals or common nouns) conventionally denote metalinguistic properties, and that argumental names are really nominals in combination with a silent definite determiner. The question now is: does this result in a smoother overall account of processing (across all uses of names)?

Proper names that occur as canonical arguments – by far the most common sort of occurrence – are now the coerced form, the syntactic fish out of water. A listener must add, to the metalinguistic property contributed by the lexicon, a definite determiner meaning not associated with any (visible) component of the expression, and then must compute the meaning of the whole.

In fact, the problem of incomplete definite descriptions means that even more work is required. So far, we have interpreted the name *John* as synonymous with the description *the individual called John*, which is 'incomplete', since there is more than one individual called John. Numerous solutions have been proposed for the general

²⁷I don't develop such an account here. Consult Clark and Gerrig (1983); Clark (1996); Hobbs et al. (1993).

problem (famously, Neale (1990) identifies ‘implicit’ and ‘explicit’ approaches), all of which demand extra work from the audience. For argumental names, Matushansky (2006b) proposes a (further) hidden component denoting a *naming convention*. Thus *John* is synonymous with ‘the individual called John according to convention *c*’. Since *c* is not the conventional meaning of any component of the utterance, its interpretation must be determined pragmatically, by collaboration between the speaker and the audience.

Elbourne’s means of solving the incomplete description problem results in an account that is manifestly worse off processing-wise than the ecumenical approach. Elbourne ‘completes’ the description thus: ‘the individual called John and identical to *u*’, where *u* is a variable assigned to John in the context (Elbourne, 2005). The ecumenical account says that the argumental name *John* denotes the variable (dref) *u* (by itself). On Elbourne’s account, the hearer must determine which variable is meant in addition to doing all the work mentioned in the last paragraph.

Note too that the metalinguistic property in Elbourne’s account is a spinning wheel: it could be dropped without loss, since the remainder of the restrictor uniquely specifies an individual (the value of *u*).²⁸ On my account the *syntactic form* of the name *John* helps the hearer to determine the speaker’s choice of variable/dref, and this is presumably true for Elbourne, too. The name’s metalinguistic denotation is truly ineffectual.

The only place where the semantic account does better than the ecumenical account is, naturally enough, in the metalinguistic interpretation of irregular names. Irregular names without a metalinguistic interpretation, by contrast, are difficult to treat on that account. First, there arises the question of how the conventional metalinguistic interpretation is to be excluded. On the ecumenical account, unfamiliar

²⁸Indeed, the metalinguistic clause may actually lead to faulty predictions, depending on the precise understanding of the predicate *called John*. For instance, I might use the name ‘John’ to refer to a man with a different name (i.e. a man who would not seem to satisfy the predicate *called John*) because, for instance, he is in a witness protection program and his name must be concealed.

syntax flags a non-conventional interpretation. Here, however, we have names appearing as nominals – their natural integument, by the lights of the semantic theory.

Even once we have excluded the metalinguistic interpretation, the resulting account is once again ‘strictly dominated’ by the ecumenical theory.

Take an example like:

(12) That’s not an El Greco!

The ecumenical account goes as follows. The name *El Greco*, which ordinarily denotes the dref u_{gre} , must in this case contribute a property-level denotation, in order to combine with the meaning of the determiner *an* (we prefer to coerce open-class expressions over functional ones). In this case, since one of the salient facts about El Greco is that he painted masterpieces, the leading candidate for the nonce meaning is *painting by El Greco*. To arrive at this interpretation, we found an appropriate relation and applied it to the customary denotation of *El Greco*.²⁹

On the semantic approach, the hearer must hit on the same relation, but must additionally do all the extra work (either by Elbourne’s or Matushansky’s method) to derive the argumental interpretation of *El Greco* from the metalinguistic one.

In this section, I have provided an alternative, ecumenical account of the irregular name data. This account does not transgress Grice’s principle of semantic parsimony, since the extra interpretations of names are not part of the lexicon, but are created on the fly using the resources of the discourse context. I have also shown, in a rough-and-ready sketch of linguistic processing, that the ecumenical approach would actually have less work generating the necessary readings than the metalinguistic theory of names.

²⁹Note that any property the hearer might reasonably connect to the speaker’s utterance of *El Greco*, whether derived from its semantics, its syntax (as with metalinguistic interpretations) or its phonetics, is a potential denotation of the coerced use.

Chapter 2

Context

This chapter is about context, but we must warm to our topic. In the first section, I introduce two general *perspectives* on the enterprise of semantics: the objective, which conceives of natural language semantics as an autonomous, rule-governed system, and of semantic properties as applying directly to natural language expressions; and the subjective, which admits a ‘conceptual semantics’ of mentalese, but rejects the autonomy of natural language semantics. At best, a natural language sentence (on the subjective view) has secondary semantic properties *derived* from the mental structures it gives rise to. These two perspectives on semantics interact with context: in the first place, they promote different understandings of the context of interpretation; additionally, the empirical problem of linguistic context, and its dynamic solution, are thought (by, for instance, the developers of DRT) to require a subjective, representational version of context, and so tell in favour of the subjective perspective.

In the second half of the chapter, I visit a recent series of articles by (combinations of) Stone, DeVault and Thomason, advocating a realist, yet representational account of context (and thus providing the objectivist with a solution to the problem of linguistic context). I comment on the place of discourse referents in this picture, which offers a literal interpretation of Karttunen’s inflationist terminology. Finally, I

produce the central claims of the chapter: that drefs can occur in different contexts, and that names denote such *context-spanning* drefs. These results are informed by the investigations of the previous chapter, but are only possible on an objectivist account of context (since subjective contexts cannot, by definition, share drefs).

2.1 Semantics: Subjective vs. Objective

In this section, we will debate the autonomy of semantics from psychology. By the objective perspective on semantics, I mean the position taken by David Lewis in the following quote:

I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or a population. Only confusion comes of mixing these two topics. (Lewis, 1972, 170)

Lewis discriminates semantics from psychology. The objective tradition descends from the original insights of Frege, who convincingly argued that the study of meaning was not *eo ipso* the study of psychology, and in particular that the *sense* of a sentence could be an abstract, rather than a psychological entity (Frege, 1953, 1892). Another precursor of this perspective is the semantics of formal languages, including Tarski's recursive definition of truth for predicate logic (Tarski, 1956).

Of late there has been a subjective turn in natural language semantics. The philosophical foundations for this movement were laid in (Fodor, 1975). Others, notably Chomsky (2000), have argued for the dissolution of natural language semantics altogether. Semantic judgments about natural language are accounted for, on the subjectivist version of things, by way of logically prior semantic facts about cognitive intermediaries:

Conceptual Semantics... is concerned most directly with the form of the internal mental representations that constitute conceptual structure and

with the formal relations between this level and other levels of representation. . . . Conceptual Semantics is thus a prerequisite to [natural language] semantics: the first thing one must know about an English sentence is its translation into conceptual structure. Its truth conditions should then follow from its conceptual structure plus rules of inference, which are stated as well in terms of conceptual structure. (Jackendoff, 1994, 132)

The subjective approach folds semantics into psychology. A subjectivist explanation of some ostensibly semantic fact, such as that the sentence s is true iff p , must include a psychological lemma (that utterances of s induce ‘conceptual structure’ s'). Some prefer conceptual semantics to its abstract cousin for its greater *gravitas* and scientific credentials. Others, like Fodor and Chomsky, take philosophical exception to an autonomous natural language semantics, assuming that it must be continuous with a theory of ‘everything else’.

It is not inordinately clear what the ‘translation’ of an English sentence into conceptual structure might amount to (we will look at a specific example in the next section). Will there be *rules* governing the translation of English into mentalese, or merely non-normative generalizations? Note that if there *are* such rules, then there will be a derived system of rules – the result of putting the former together with the semantic rules for conceptual structures – associating English expressions with semantic properties. These properties might be unexpected (e.g. DRT provides English with a non-compositional semantics), but they are nonetheless anathema to the semantic nihilist.

If there are no translation rules, but only generalizations, then the subjectivist theory can’t be normative. It won’t be able to distinguish, on its own terms, the intuitively different cases of producing a malapropism and, say, failing to instantiate Zipf’s generalization connecting a word’s meanings with its frequency. Failing to attach a high number of meanings to a commonly-used word is not a form of error; it is merely an outlier on a chart. For the subjectivist without translation rules, natural language semantics is statistical.

The debate about the objectivity of semantics is independent of the turf war be-

tween the ‘product’ and ‘action’ traditions in the study of language (Clark, 1992, xi-xiii). This may come as a surprise. After all, Grice’s psychological analysis of non-natural meaning (Grice, 1957) stands at the foundation of the action tradition (Clark, 1996; Stone, 2003). However, a reductive account of speaker meaning in terms of intentions is no bar to the autonomy of semantics. As the Gricean Stephen Neale, among others (the point is already made in (Grice, 1971)), makes clear, the communicative intentions constitutive of speaker meaning are tightly constrained by a speaker’s beliefs – about the circumstances of the utterance situation and the hearer’s capacities – but also about the rules of the language (Neale, 2005, 189). Thus the existence and efficacy of independent semantic rules (for instance principles of semantic composition, or the rule assigning a word its conventional meaning) are not threatened by Grice’s celebrated analysis. Indeed, in the course of this dissertation I follow the example of others (e.g. Matthew Stone) in defending an objectivist, Gricean view.

The objective perspective has a number of intuitive advantages over the subjective. We ordinarily speak as though language has rules and words have meanings; we do not, except as theorists, prefer to say that the exercise of one and utterances of the other merely have regular effects on our patterns of thought. It is natural to hold, for instance, that an infant learns the semantic rules of a language – what its words mean – a turn of phrase unavailable to the conceptual semanticist. While particular semantically-laden statements admit of reduction into (usually vague) psychological terms,¹ it is an ingrained lesson that such piecemeal reductions cannot be parlayed into the complete reduction of the higher science (Fodor, 1974).

The subjectivist concedes these (prodigious) advantages. Jackendoff himself demonstrates ‘that a mentalistic theory of language proves not so easy to reconcile with [the] commonsense position on reference and truth’ (2002, 294). His conclusion: ‘it is neces-

¹For example, the statement that *bald* means having 500 or fewer hairs on one’s head may be reduced to the claim that English-speakers generally use *bald* when they intend to indicate that someone has 500 or fewer hairs (Williamson, 1994).

sary to thoroughly psychologize not just language, but also “the world”’.² A common response (e.g. Laurence, 2003) on the part of the subjectivist is to disparage the role of commonsense intuitions in the pursuit of science, whose greatest discoveries are frequently unintuitive. Jackendoff effects a farcical comparison with the opposition of the ‘commonsense’ idea that the sun sets to the underlying truth of the earth’s rotation (2002, 294–5), although in the next breath he is calling our intuitions to mutiny against the notion that ‘Nicaraguan Sign Language [should have been] lying around in the abstract domain until the 1980s, when it was at last grasped by someone’ (ibid, 299).

While abstruse disciplines such as modern physics are notoriously unintuitive, there is reason to think that semantics might be exempt from this general tendency. Human locomotion doesn’t rely on our commonsense notion of physics, however communication *does* rely on our common understanding of semantic facts. In particular, the assumption that a particular semantic hypothesis is shared by one’s interlocutor plays an indispensable role in both the generation and interpretation of natural language utterances. Since semantic facts are themselves the products of negotiated convention (Lewis, 1969), it is self-defeating to argue that commonsense semantic hypotheses are convenient fictions or heuristics that facilitate coordination on an interpretation (because they are shared), yet belie the crypto-operations of scientific semantics. Yet this is what the analogy with theoretical physics amounts to.

2.2 Linguistic Context and Dynamic Meaning

It is a truism that interpretation depends on the context. Indeed, it would seem that something stronger is true. Since, in coming up with an interpretation, an audience’s resources are limited to the utterance made, the rules of the language (if such

²The fact that he adverts to a popular science interpretation of Special Relativity in explicating this ‘move’ provides an indication of the sheer radicalness of individual subjectivists.

are admissible), and the context it was made in; and, moreover, since an utterance is designed to yield a unique interpretation, it must be that an utterance in a context *determines* an interpretation (in accordance with the rules of the language spoken). I see no harm in thus deriving a ‘metaphysical’ conclusion from epistemological premises.

Context is crucial to semantics in explaining away apparently non-compositional phenomena, such as the fact that utterances of the same sentence can have different truth conditions (e.g. when the sentence contains an indexical). However, the understanding of context is particularly susceptible to biases of perspective. The objective notion of context, dubbed ‘utterance context’ in the literature, is primarily associated with the work of David Kaplan. Utterance context fits the template above by supplying referents to indexicals and demonstratives. Linguistic meanings may be treated (once more compositionally) as functions from contexts to objects and properties, or ‘characters’ (Kaplan, 1989).

The subjective notion of context – sometimes called ‘discourse context’, at others ‘common ground’ – is prominent in the work of Herbert Clark and Robert Stalnaker. It consists of the information (including facts about attention) available from a particular perspective. This perspective might be as simple as the beliefs of a single participant, or else as complex as the mutual presuppositions of *all* participants. The former is connected with the psychologicistic perspective of conceptual semantics, while the latter is closely allied with the action tradition and the treatment of language use as coordinated activity. Stalnaker’s work accounts for the apparently non-compositional phenomenon of presupposition projection in complex sentences and discourses (although for him presupposition is a pragmatic phenomenon).

The subjective notion of discourse context is thought (especially by Hamm et al., 2006) to be connected with the effect of *linguistic context* on interpretation. In addition to being dependent on factors of the external environment, the interpretation of a sentence is affected by the sentences that have come before it. The following serves as an illustration of this ubiquitous phenomenon:

- (1) a. You smoked all my weed, man. You got to call the green man.³
 b. At this crossing, the walk signal isn't periodic. You have to call the green man.

In each case, the (almost identical) second sentence has a substantially different interpretation. While the second sentence of (1a) is a command to the hearer to phone her marijuana dealer, its counterpart in (1b) asserts that in order to cue the walk signal (a striding green man in the UK and Australia), pedestrians must activate a mechanism.

So how does the empirical matter of linguistic context bear on the theoretical issue of the subjectivity of semantics? Hamm et al. (2006) claim that to account for the influence of the preceding discourse on interpretation, a semantic theory must avail itself of a notion of context that is *representational*.

Suppose we as objectivists wish to account for the effect of linguistic context on interpretation. At the very least, we must add an extra parameter – for the preceding linguistic context – to our contextual ‘index’. This will minimally satisfy the requirement on the context that it specify everything necessary for the interpretation of an utterance. However, something is still lacking. Suppose we have a sentence S (with character C) interpreted at an index c – including a parameter c_L for the linguistic context – modelling the utterance context. S expresses the proposition – has the truth conditions – $C(c)$ in the context. So far so good. Now consider the *sequence* of sentences $S \frown S'$, also interpreted at c . Let's suppose our semantic theory provides a character C' for S' . We can calculate, as before, the truth conditions of the first segment of the discourse as $C(c)$. However we also require an index at which to interpret the second segment, S' . What index should we use? It can't be c itself, since the intervening utterance of S has presumably *changed* the context in certain ways (as in example (1) above).

³From A. Winehouse, ‘Addicted’.

It is possible to show that the context in which S' is to be interpreted is not a function of the starting context c and the truth conditions $C(c)$ of the utterance of S . The demonstration uses Partee's marble example (Kamp et al., 1999; Heim, 1982):

- (2) a. Out of ten marbles, one is not in the bag. It must be under the couch.
 b. Out of ten marbles, nine are in the bag. ??It must be under the couch.

The first sentence of (2a) is truth-conditionally equivalent to the first sentence of (2b), yet the former differs from the latter in licensing the pronoun *it* in subsequent discourse. Clearly, truth-conditionally equivalent utterances can update the context differently.

If we are to arrive at an adequate semantics for connected discourses, we must specify semantic values at a more fine-grained level than characters. (The move is analogous to Frege's differentiation of the 'thought' expressed by a sentence from its truth-value in his account of attitude reports (Frege, 1892).) Dynamic semantics (Kamp, 1981; Heim, 1982; Groenendijk and Stokhof, 1991) takes the bull by the horns, and treats sentences as themselves denoting contextual updates. This way we always know, given the meaning of a sentence and a context in which it is uttered, the context in which the next sentence in line is to be interpreted – it is the result of updating the context of utterance with the denotation of the first sentence.

Examples like (2) require a context that is representational. It should (minimally) keep track of the anaphoric potentialities of the discourse in progress (Groenendijk and Stokhof, 1991). A Kaplan-style utterance context, which only specifies factors relevant to the truth conditions of utterances, is demonstrably insufficient. A further problem is that no (non-incantatory) utterance literally changes the context, conceived in Kaplan's way as a physical and logical environment. An utterance might update a receptive audience member's *representation* of that environment, however. While the world does not change to accommodate someone's assertion that p , the world *as someone sees it* might well do so.

2.3 DRT, a Study in Subjectivism

The incremental nature of interpretation is closely connected with a ubiquitous feature of discourse, its semantic cohesiveness. Typically the sentences that make up a coherent piece of discourse are connected by various kinds of cross-reference. As a consequence, it is often impossible to analyze the meaning of cohesive discourse as a simple conjunction of the separate meanings of the individual sentences that make it up. The meaning of the whole is more, one might say, than the conjunction of its parts. The connection between cohesiveness and incremental discourse processing is, in rough outline, this: to understand what information is added by the next sentence of a discourse to what he has learned already from the sentences preceding it, the interpreter must relate that sentence to the information structure he has already obtained from those preceding sentences. Thus his interpretation of the new sentence must rely on two kinds of structures, the syntactic structure of the sentence itself and the structure representing the context of the earlier sentences. (Kamp and Reyle, 1993, 59)

The proponents of Discourse Representation Theory (DRT) are subjectivists. They present an idealized picture of how an interpreter processes an utterance into an ongoing mental representation (modelled by a Discourse Representation Structure, or DRS). This part of the theory is known as the DRS construction algorithm. They complete the picture with a model-theoretic semantics for DRSs, supplying mental representations with truth conditions. This results in an indirect (and non-compositional) truth-conditional semantics for natural language logical forms interpreted relative to a DRS ‘context’.

DRT suffers from the subjective perspective’s intuitive defects. While a Tarskian truth theory can serve as a model of *what we know* when we understand and use a language – the semantic rules we have learned and employ in interpretation – a theory of processing, à la DRT, has a different status. Such a theory, as it is framed at least, is a theory of *how we do* something. Then again, perhaps we shouldn’t rule out of hand the (subtle) possibility that our knowledge of language amounts to knowing the algorithm our audience will use to process our utterance (or by which an interlocutor is expecting us to process his utterance).

DRT is a theory of *competent* processing. A theory of processing *in general* would

fail to allocate the intuitively correct semantic properties to natural language sentences, since processing is prone to error. As it happens, the adjective ‘competent’ smuggles in a great deal. For one way to read ‘competent processing’ is as processing *that obeys the rules of the language*. We continue to be dogged by Lewis’ modular vision of a linguistic semantics (rules of proper processing) detached from psycholinguistic happenstance (how utterances are processed).

We might, pursuing this thought, treat the DRT construction algorithm as a complex claim about how an utterance with a particular logical form *ought to be* processed. Thus s is true iff p (according to DRT), because s should (according to ‘processing’ laws of the language) be processed into a certain structure D and D is true iff p . Since the construction algorithm is a primitive component of the account, and independent of the model theory for DRSs, one cannot further break down normative judgments about acts of processing into, for instance, claims about the preservation of semantic properties (e.g. truth). One might use preservation of truth as a litmus for adherence to the ‘laws’ of processing, but never as a justification for them, since these laws are the (Saussurean) arbitrary primitives of the grammar.

Alas, our pre-theoretic intuitions revile this suggestion. An act of processing a sentence, in our ordinary understanding, is not primitively correct or incorrect by the rules of a language – this status *can* be broken down further. We have processed a sentence correctly just in case we have interpreted its *meaning*, attributed the correct semantic properties to it (which is not to say merely that we have attributed to it *this way* of attributing it with semantic properties, as the primitive theory would have it).

2.4 Context as Objective Representation

A different approach, which continues to treat the context as a representation *without* supposing that it must be subjective, first appears in (Lewis, 1979b), and is taken up

in (DeVault and Stone, 2006) and (Thomason et al., 2007).

Suppose the context contains an *objective* representation – like the changing scorecard of a baseball game – that stays abreast of the ‘language-game’ in progress. This representation is not physical, but should be taken as an abstract theoretical posit. DeVault and Stone compare it to the abstract ‘board’ in a game of correspondence chess. While there is no (authoritative) physical board in an email chess game, there is nevertheless an objective fact about the state of the game – which pieces occupy which squares and whose move it is – at every turn. This game state can be thought of as an abstract data structure and the emailed moves as updates that transition between game states.

So what kind of representation is the state of the game? Is it some representation in some player’s head? Of course, each player must have a mental representation that tracks the state of the game. However, a little reflection shows that no such cognitive structure could determine the state of the game. What determines the state of the game is the moves that have been played, and the rules of chess. It is in this sense that the state of this chess game is an objective, abstract social construction.

(Steedman and Stone, 2006, 74)

The notion of context as an objective representation forms the basis for a dynamic view of meaning without the shortcomings of subjectivism. If the context is a real entity, rather than a mental state, then we can treat semantic values as contextual updates without turning semantics into a theory of processing. There is a perfectly natural objectivist understanding of semantic rules: they associate sentences with contextual updates (and, derivatively, with truth conditions). In a parallel manner, the rules of chess associate different moves with different updates of the game state. Knowledge of semantic rules is naturally connected with, but theoretically distinguishable from, correct semantic performance.

The realist view gives not only a coherent account of what language is but enables a coherent, computational account of how we use language. We can know the rules of language: we can maintain mental representations

whose content accurately tracks the conventions of our community. We can know the state of the discourse: we can maintain mental representations whose content accurately tracks the real-world social representations of context. And we can draw inferences – computing new mental representations about the interpretation of utterances in context – which faithfully mirror the consequences of the real-world rules in the real-world context.

[...]

Only a realist can say that language acquisition is just a case of genuine learning. The story is simple. There are general facts about meaning. The child obtains linguistic experiences that give evidence about what these facts are, and thereby arrives inductively at an increasingly precise idea of them. (Steedman and Stone, 2006, 76)

2.5 The Nature of Discourse Referents

With the notion of an objective representational context in place, we can discuss the nature of discourse referents – introduced last chapter – in more detail.

Karttunen (1976) presents an engineering task (‘Consider a device designed to read a text in some natural language...’) which he thinks will shed light on the workings of natural language. He counsels a discipline of regimentation (‘It has to be able to build a file that consists of records of all individuals, that is, events, objects, etc. mentioned in the text, and, for each individual, record whatever is said about it’) that sets the agenda for DRT’s idealized model of human processing. Along the way, he mints an intriguingly concrete terminology:

Let us say that the appearance of an indefinite noun phrase establishes a discourse referent just in case it justifies the occurrence of a coreferential pronoun or a definite noun phrase later in the text.
(Karttunen, 1976, 366)

DRT treats Karttunen’s discourse referents as private mental symbols, components of the mental model (or DRS) constructed by a human interpreter. They exist, but are confined to a particular agent’s brain. Heim considers Karttunen’s innovation to be purely terminological:

It is clear from Karttunen’s own comments that discourse referents are not individuals and that to establish a discourse referent does not necessarily mean to refer to anything. In the context of our theory, I propose that we identify discourse referents simply with numbers: to establish a discourse referent means to carry a referential index in logical form.

(Heim, 1982, 250)

If we take her ‘logical forms’ to be mental representations, then Heim seems to follow Kamp and Reyle in consigning drefs to the head. Elsewhere (281), Heim identifies discourse referents with the ‘file cards’ on which information is inscribed, but shortly thereafter claims that this is no departure from her earlier doctrine since ‘...file cards can also be identified with numbers’ (clearly ‘inscription’ stands in need of a non-standard gloss).⁴

Heim also goes on (285) to suggest that Stalnaker’s ‘common ground’ (Stalnaker, 1979) ought to be understood as a ‘file’ of file cards. However, since a common ground is a body of shared information, she implicitly endorses an understanding of discourse referents that is somewhat different to Kamp and Reyle’s. Since the medium of information is the file card (dref), file cards must be held *in common* for that information to be mutual.

DeVault and Stone’s picture of context gives us a way of treating discourse referents as common currency, while fleshing them out beyond bare ‘numbers’. On their picture, drefs are components of the objective context. Indeed, they are the variables whose values *constitute* the state of the context, and so are the very weft of the representation – DeVault and Stone liken them to squares on a chessboard. Since the context is objective – independent of any agent’s mental state – drefs count as objects of independent means. An indefinite, for example, adds a new (abstract) object to

⁴Unsettlingly, Heim proceeds, ‘That is the way we have been referring to them anyway: “card number 1,” “card number 2,” etc.’ suggesting that the previous sentence should be read: ‘file cards can also be identified *by* numbers’. Of course, the fact that some cards are *indexed* by numbers is consistent with the cards themselves being something other than numbers. If file cards are meant to be distinct from numbers, then the proposal that discourse referents are file cards *is* an advance on the proposal that discourse referents are numbers. Moreover, the suggestion that discourse referents are only identified *by* numbers no longer ‘does away with questions as to their ontological status’ (282).

the (abstract) context. Moreover, we can treat discourse properties, like the stages of attention discussed in the last chapter, as metaphysical flora: properties of discourse referents. A dref acquires the property of being *in-focus* after being denoted by an expression in subject position. A dref's being *in-focus* is thus an objective fact (determined by the course of the conversation and the rules of language), which an interpreter could make a mistake about.

2.6 Names: Denotation

The discourse semantics of names depends on an objective understanding of discourse referents. In the last chapter we were exercised by discourse-initial names, which seemed to present the familiarity theory of definiteness with an anomaly: a definite expression with no indefinite antecedent (in the discourse). The dref denoted by a discourse-initial name must be familiar, yet it could not have acquired that status *in the course of the discourse*.

Our solution was to say that the discourse referent was *already familiar* at the beginning of the discourse. It was familiar to each of the participants (being present in their 'long-term' memory), since each had been introduced to it – the very same dref – in an earlier context. Indeed, it is possible that each participant was introduced to the dref in a *different* context. However, so long as they were all introduced to the same dref, an anaphoric path could be traced back from each of their introductions to a common point. This is because, as we may now put it, the *history* of each discourse referent is a single tree of continuous, branching anaphoric paths, and any two points on such a tree (corresponding to utterances denoting the dref) exist on paths that meet at some point (if nowhere else, at the root).

An utterance of a name α is therefore felicitous iff it enables the speaker and the audience to coordinate on a particular *discourse referent*. In other words, there must be a unique dref u , having (at least) the status α -*activated*, such that the

utterance denotes u (as determined by the speaker's intention) and the hearer is in a position to infer that the utterance denotes u . Apart from the requisite attentional status, this definition is the same for pronouns. Note, too, that coordination on a discourse referent is only possible if discourse referents are objective. Two agents cannot converge on a mental symbol (as drefs are in DRT) that at most one of them can be privy to. On the objectivist version of things, drefs are components of an exterior context (of which each participant carries a more-or-less accurate interior model), making it possible for two agents to latch onto the same one.

In order to predict the felicity of certain discourse-initial names, we must assume that the discourse context begins in a particular non-trivial state of information and attention. In particular, those drefs with which all the participants are familiar (in 'long-term' memory) will have a suitable attentional status.

There is only one starting position for a game of chess. The pieces face off in their customary phalanxes and white has the first move. There is, however, a constellation of different initial conditions for discourses. Conversations between intimates start out on a different footing to conversations between strangers. Among the latter, conversations between people from the same geographical location are likely to have different 'presets' to conversations between people from opposite sides of the globe (cf. Clark, 1996, Ch. 4). A social interaction between strangers often involves some plumbing of social commonalities ('What do you do for a living?', 'That's an interesting accent, where are you from?'), which can be seen as stage-setting for the remainder of the conversation.

What happens when interlocutors discover, midway through a conversation, that they have more in common than they originally supposed? Once again, the mid-conversation shift to a more comprehensive 'initial' state might appear to require accommodation. I propose, instead, to adhere to what I will call the 'secret society' model of context. You and I were members of the same underground league of Hanseatic oligarchs even before I produced the handshake that led to our mutual

awareness of this fact. On the secret society model, the scorecard has *hidden presets*, which public events in the course of our interaction may reveal, allowing us to augment our mental representations of the context. According to this model, some drefs may be available to us (those common to members of our ‘society’), though we don’t yet have mutual knowledge of the fact (and so will hardly be inclined to exploit their presence). Uttering the secret name of a Bavarian Illuminato (for example) unveils more of the context we share by demonstrating that a certain reference is available to us.

The secret society model demands an objective understanding of context, since it supposes that the discourse context might be a certain way without anyone realising it. Indeed, if we reflect that no two people are ever so intimate as to know *every one* of their common references, we are forced to say that there is always a certain amount missing from our representations of the context.⁵

2.7 Names: Reference

Discourse semantics is all about matters of coherence and co-reference. Truth-conditional semantics, on the other hand, is chiefly concerned with reference. While a name *denotes* an abstract object – a discourse referent – it also *refers to* some (usually non-abstract) individual.

Chastain (1975) and Kripke (1972) account for the reference of names with the help of the context-spanning anaphoric trees mentioned in the last section. For both,

⁵Contextual realism leaves room for less-than-full knowledge of the context – as when one is unsure about a past move in a game of correspondence chess – but hardly mandates it. In the account of (Thomason et al., 2007), being uncertain about the context is a defective state, which public acts (like utterances) are calculated to resolve. Steedman and Stone (2006), on the other hand, seem to countenance ineradicable uncertainty about the context when they allude to Williamson’s ‘epistemicist’ account of vagueness (Williamson, 1994). In suggesting that the context specifies the precise threshold at which a vague predicate is satisfied (e.g. the number of hairs necessary for one to qualify as ‘not bald’), they implicitly suggest that some parameters of the context are, by epistemicist principles, unknowable. Moreover, on the revised model of (DeVault and Stone, 2007), the periodic ‘checkpoints’ of (Thomason et al., 2007) and the corresponding requirement of *strong recognizability* are relaxed in favour of a certain tolerance for uncertainty.

an anaphorically continuous path may be traced from the use of a name to the referent of the name on that use. Kripke famously proposed that the path terminated in a special baptismal rite, which grounded the name in an object identified perceptually or by description. In the more sophisticated model of (Evans, 1982), itself an improvement on (Putnam, 1975),⁶ a name is grounded in its referent by a series of identificatory events in different contexts we might call (adapting Evans somewhat) ‘producer’ contexts. While anaphoric branches extend to ‘consumer’ (non-producer) contexts too, identifications local to these contexts do not have the authority to ground the name (compare Putnam’s ‘division of linguistic labour’). The theory is geared to handle diachronic reference-shift, as brought out in Evans’ Madagascar example, and as described below for *Rum*:

The Roman imperial mantle on Greek shoulders has led to a splendid confusion; for the word ‘Rum’, on Oriental tongues, referred not only to the Christian Byzantines – they are so styled in the Koran – but, for a century or two, to their conquered territory in Asia Minor; it designated the empire of the Seldjuk Turks in Anatolia with its capital at Konia (Iconium), reigned over by the ‘Sultans of Rum’. To tangle matters still further the word Romania was often used in the West, especially during the crusades, to specify the parts of the Eastern Empire which lay in Europe; the Turks extended ‘Rum’ into ‘Rumeli’, (‘land of the Rumis’) to cover the same area. One still finds the confusing word ‘Rumelia’ on old maps. (In Greece, Rumeli now specifically applies to the great mountainous stretch of continental Greece running from the Adriatic to the Aegean, north of the Gulf of Corinth and south of Epirus and Thessaly.)
(Fermor, 1966, 98)

I abstract from the accounts above the idea that a dref’s history (an anaphoric tree) determines its reference (if any), and so the reference of any utterance denoting that dref. In further speculative semantic naturalization, I do not at present indulge.

Returning once more to the debate between the subjective and objective takes on semantics, I would like to note the subjectivist Kamp’s approach to reference. Recall that DRT consists of an algorithm for converting pairs of logical forms and

⁶See also (Devitt, 1974) – in particular his claim that the ‘sense’ of a name is the causal network underlying it (204), which, while it doesn’t make a great deal of sense as it stands, might be seen as an early precursor of the view proposed here.

DRSs into new DRSs, and a model theory specifying the truth conditions of DRSs. A natural language sentence is true in a model M iff processing it in the ‘empty’ context produces a DRS that is true in M . As Kamp desires the derived natural language semantics to be materially adequate, he must ensure that sentences with referential expressions have truth conditions that depend on the expressions’ referents.

His solution is to *externally anchor* drefs to objects in the model. An ‘external anchor’ is formally just a pair of a discourse referent and an object (the discourse referent’s referent, just as Jeeves is a gentleman’s personal gentleman). The anchor is *not* understood as part of the representation constructed by the interpreter (in the sense that is not transparent to them), but nonetheless constrains the possible embeddings of the representation into the model (and so the truth conditions of the utterance). The external anchor is supposed to stand for the combination of psychological, sociological and causal factors that go into determining the reference of a mental symbol (as drefs are in DRT).

Thus DRT contexts are really hybrids of a subjective discourse context and a set of external anchors reminiscent of Kaplan’s utterance context. Note too that it is drefs – mental symbols – that are directly assigned referents by the theory. A name refers only derivatively, by giving rise to a mental symbol with a particular reference. However, intuitively (and in the accounts of Kripke and Putnam) the dependency operates in reverse: the reference of a (non-deictic) mental symbol piggybacks on the reference of a name, and is determined by the history and use of that name.

Chapter 3

Content

In the previous chapter, we brought our intuitions to bear on the proper understanding of semantics. I came down on the objectivist side of the debate, and decided that sentence meanings were updates to a representational, but nonetheless objective context. In this chapter, I will present an account of communication and communicative content that is equally beholden to commonsense intuitions, in particular to judgments about the success of acts of communication and interpretation. This topic is obviously related to the previous one. An act of communication is successful so long as the hearer has correctly identified the *meaning* of the utterance. In fact I will be extending the account of meaning from the last chapter in a natural way to address this further issue.

I begin by supposing that ‘referential isomorphism’ is the hallmark of successful communication (Stone, 2004a,b). I proceed to develop a counterexample due to Brian Loar (Loar, 1976), showing that the criterion fails to respect our intuitive normative assessments.

Next, I test the mettle of the semantic subjectivist (discussed in the previous chapter). Pure subjectivism can respond to Loar’s counterexample in a limited way, by pursuing a suggestion of Nicholas Asher’s (Asher, 1986). However, it is easy to show that this response is itself open to counterexample. I conclude (with Kamp) that

the subjective approach stands in need of supplementation of an objective sort. The combination of objective representational contexts and drefs that span those contexts (from the previous chapter) fills the breach admirably.

Each account of communication privileges a particular ‘grain’ of content (the level of content communicated in the act). Referential isomorphism corresponds to a fairly coarse grain of content: the Russellian proposition. The account settled on distinguishes content of a finer grain, which can moreover be identified with the (discourse semantic) denotation of the utterance. This notion of ‘discourse’ content is also central to the chapter to follow on attitude reports.

3.1 Referential Isomorphism

Consider the utterance, made by Bob to Alice:

(1) The red apple is bruised

Suppose (1) manifests Bob’s effort to alter the context to reflect the fact that a particular apple in plain view, a , is damaged in a particular way, F . I will further assume (to ease the exposition) the Representational Theory of Mind (Fodor, 1987), which connects psychological accounts given in the ‘semantic’ terms of belief/desire explanation to a hypothetical syntactic implementation in the brain. Concretely, suppose Bob is in possession of a mental symbol, *bobs-object-4*, deictically grounded in the apple a . He also has the symbol *marred-by-contusion*, grounded (in a more nuanced way) in the property F .

The message Bob wishes to convey to Alice by uttering (1) is, as he would put it to himself:

(2) *marred-by-contusion(bobs-object-4)*

In processing Bob’s utterance, Alice must formulate her own representation of Bob’s message, couched in her own mental vocabulary. Let’s suppose she has a different primitive symbol, *alices-object-15*, for the apple *a* and another, *spoiled-by-hitting*, for the property *F*. What must occur for Alice to construe the message behind Bob’s utterance correctly? Since Alice’s symbols differ from Bob’s, she has no chance of reproducing the *syntax* of Bob’s representation. Instead, she must construct a representation that is *referentially isomorphic* to Bob’s (Stone, 2004b). That is to say: the atomic symbols must refer to the same things and be syntactically composed in the same manner, as shown in (3):

- (3) a. *marred-by-contusion(bobs-object-4)*
 b. *spoiled-by-hitting(alices-object-15)*

Note that both representations in (3) determine the same *Russellian proposition*, which we write as an ordered pair of its components: $\langle F, a \rangle$. We can think of this proposition as the *content* communicated by the assertion; it started off as the content of Bob’s communicative intention and came to comprise the content of Alice’s construal. The requirement that Alice’s construal must be referentially isomorphic to Bob’s amounts to the claim that communication is the lossless transmission of a particular Russellian proposition.

3.2 Loar’s Counterexample

The assumption that referential isomorphism is sufficient for communication is controversial. In particular, there is an outstanding counterexample, due to Loar (1976, 357), which I will retread for the present discussion.

Bouvard and Pécuchet are up all night gazing at the heavens. At dusk, Pécuchet notices the first star and proposes to name it ‘Hesperus’, after the Latin. The friends, quite taken with the find, discuss it for the remainder of the night, failing to register all

but the most coruscating of subsequent developments. Finally, around dawn, Bouvard makes his own celestial discovery: a brilliant dot in the Eastern sky. Awestruck, he gasps:

(4) It's so bright!

Pécuchet, half-asleep, murmurs in agreement. He is, however, oblivious to Bouvard's discovery, and assumes that his friend is merely continuing an extended peroration on 'Hesperus'. Intuitively, they have failed to communicate: Pécuchet has interpreted the pronoun *it* as anaphoric to the name 'Hesperus', whereas Bouvard had intended deictic reference to an object visible in the lightening Eastern sky. As it happens, Bouvard has only rediscovered Pécuchet's planet on its morning leg, which means that, despite his intuitive failure to understand Bouvard's intention, Pécuchet has produced an interpretation that is referentially isomorphic to it.

Bouvard's intended message was that the newly discovered star was bright; Pécuchet construed him as asserting that 'Hesperus' was:

- (5) a. *high-luminance(bouwards-object-5)*
 b. *brilliant-colour(pecuchets-hesperus)*

Clearly there is a reference-preserving isomorphism (mapping *pecuchets-hesperus* to *bouwards-object-5* and *brilliant-colour* to *high-luminance*) between (5a) and (5b). It follows that referential isomorphism is not strong enough to characterize all and only cases of successful communication. As a corollary to this, Russellian propositions are too coarse-grained to serve as the distinguished contents of communicative acts.

One view of this example (a fruitful one, I will argue) is that Pécuchet's interpretation is problematic because it makes the wrong connection with the surrounding discourse. The symbol *pecuchets-hesperus* is an internal reflection of a particular role in discourse (the role occupied by utterances of 'Hesperus' and initialized by Pécuchet's dubbing), while the symbol *bouwards-object-5* is connected to a more re-

cent, perceptually-supplied ‘role’.

Of course other diagnoses are possible. Bouvard has another mental symbol – *bouwards-hesperus* – that is intuitively the proper correlate of *pecuchets-hesperus*. If we can somehow characterize the notion of ‘proper correlate’ internalistically, then we can produce an alternative account of communication that is satisfactory for the subjectivist.

An interesting proposal in this direction is due to Nicholas Asher, who utilizes the framework of DRT ‘to speak coherently of two distinct beliefs being about the same object in an “internal” sense[,] so that the truth of describing two beliefs as so related does not depend on there being an external object that these two beliefs are about’ (Asher, 1986, 128). His proposal brings out nicely the limits of the subjective approach, as applied to communication. While his work is aimed at the semantics of attitude reports (which won’t be discussed for another chapter), he formulates constraints that are supposed to mark the intuitive limits on the correct representation of another’s cognitive state (and so are relevant here).

Asher’s approach improves upon Stone’s. He suggests that fidelity of construal is not determined in isolation, as we have been assuming, but takes place against the backdrop of (a representation of) the speaker’s overall mental state (echoing last chapter’s discussion of interpretation relative to a ‘discourse context’). Thus the appropriate translation of a speaker’s symbol (such as *bouwards-object-5*) must conform to it, not just in reference, but also in its *internal connections* within the overall cognitive state (insofar as these are reproduced). Asher’s proposal is the internal analogue of the ‘discourse role’ suggestion above (but thought unfit for subjectivists).

Concretely, Asher proposes that a symbol x used to represent a symbol y in another’s mental life must have an identical (or at least similar) *internal anchor* to y , where an internal anchor is a privileged set of descriptive conditions. He also calls these ‘familiarity conditions’ and has the following to say:

If the processing of a definite α within a belief context leads a recipient

to assume the believer has some familiarity with the referent of α , then he must also assume that the believer has some way of determining which individual that might be. I shall call this sort of information the believer's *familiarity conditions*. (ibid., 141)

Asher's internal anchors clearly have an affiliation with acquaintance relations (Lewis, 1979a; Maier, 2006). His more stringent conditions on communication can be used to invalidate Pécuchet's interpretation. Let's assume Bouvard has the following internal anchors:

- (6) a. *dusk-appearing(bouwards-hesperus)*
 b. *dawn-appearing(bouwards-object-5)*

And Pécuchet's representation of Bouvard's cognitive state contains the internal anchor:¹

- (7) *evening-visible(pecuchets-hesperus)*

Now notice that Pécuchet's construal of Bouvard's message doesn't respect internal anchoring, since he uses the symbol *pecuchets-hesperus* to represent the symbol *bouwards-object-5*, and the two symbols have non-equivalent internal anchors. Thus Asher correctly predicts (in this instance) that Pécuchet's construal is incorrect.

Unfortunately there are notorious problems with this sort of account, which relies on the fortuitous coincidence of descriptive content in the mental lives of different parties (Kripke, 1972; Donnellan, 1972). What happens if the interpreter misrepresents the speaker's internal anchors? He might have no representation of them at all, or else have them reversed (i.e. the symbol he associates with the name 'Hesperus' has an internal anchor that describes it as rising in the morning). In such cases, Asher must make predictions contrary to our intuitions.

Unfortunately, subjectivists can do no better than this. By internalizing discourse roles, the subjectivist jettisons the means for the intuitive characterization of commu-

¹In Asher's system, the dref occurring in the internal anchor would need to be different to Pécuchet's own 'Hesperus' dref, but I am ignoring this complication for now.

nication presented earlier. Since the hearer’s internal model of the speaker’s internal model is sequestered from its target (i.e. the same discourse role cannot occur in both), the two may only be compared by their internal structure (internal anchors) or by reference (external anchors).

Hans Kamp, recognising this difficulty, adds a new widget to the DRT toolbox. I would argue that, in so doing, he ventures further beyond the ambit of pure subjective semantics.

Concretely, Kamp allows the relation of ‘proper correlation’ between discourse referents to be stipulated externally to the DRSs in which they appear. He writes:

The pairs [of the relation] stand here outside the representations of both attitudinal states. In view of this, they must be understood as external assessments of subjects and their states, which involve not only the individual psychologies of those subjects, but also certain causal relations between them (having to do with their present and past communications and more generally with their interactions within the larger information exchanging community of which they are a part), just as external anchors must be understood as claims based not only on a person’s psychology but also on the causal relations in which he stands to the things towards which his thoughts are directed. (Kamp, 1990, 81)

Once again, Kamp is hybridizing subjective and objective strains of context. It should be noted (Kamp, judging by the passage above, would not demur) that one of the effects of language use is to forge the ‘external’ bonds between mental symbols that correlate them. This means that utterances have a dynamic impact on the *objective* as well as the subjective components of Kamp’s hybrid context, showing that the arguments to the effect that dynamicism necessitates subjectivism (discussed in Chapter 2) rest on shaky ground.

3.3 Denotational Isomorphism

Bouvard’s utterance was a failure; Pécuchet misconstrued his communicative intention. The misconstrual was referentially isomorphic to the original message, but

this proved insufficient. What more could poor Pécuchet have done? Certainly, he couldn't reproduce the private *syntax* in which the message was indited. Nevertheless, it was a mistake to use *pecuchets-hesperus* as a construal of *bouwards-object-5*, because the symbol was connected in the wrong way to the surrounding discourse.

Pécuchet's discovery and subsequent dubbing created an 'antecedent' for a conga-line of anaphoric utterances, including occurrences of the name *Hesperus* and the pronoun *it*. This dubbing also forged Pécuchet's mental symbol *pecuchets-hesperus* as well as Bouvard's *bouwards-hesperus*. As a result, both belong on the same *anaphoric tree* (Pécuchet's initializing utterance corresponds to the 'root' of the tree). Both are connected to some (public) segment of the tree by communicative intentions of the sort already discussed. Thus, *pecuchets-hesperus* and *bouwards-hesperus* ply the same role in discourse, and also denote the same discourse referent.

I present the following revised condition on successful communication: there must be an isomorphism between the intention and its construal that takes each atomic symbol to one with the same *denotation*. We will call this requirement 'denotational isomorphism'. Note that any denotational isomorphism is also a referential isomorphism, as discourse reference determines reference (if any). Moreover, reference does *not* determine denotation, and so the new constraint is more restrictive, as desired. It also correctly predicts the failure of Pécuchet's construal, since *pecuchets-hesperus* has a different denotation (though not reference) to *bouwards-object-5*.

Representations that are referentially isomorphic determine the same Russellian proposition. Representations that are denotationally isomorphic determine a more fine-grained semantic object: a tree whose leaf nodes are decorated with the *denotations* (rather than the referents) of the atomic components of the representation. Once again, such trees are strictly more precise than Russellian propositions. Whereas (5a) and (5b) determine the same Russellian proposition – $\langle \textit{bright}, \textit{Venus} \rangle$ – they determine different denotational trees – $\langle \textit{bright}, u_{it} \rangle$ and $\langle \textit{bright}, u_{hes} \rangle$, respectively.

Grice (1957) identified the meaning of an utterance with the content of a speaker's

communicative intention. We have, in the foregoing, identified the content of the update that is the meaning of an utterance with the content of the utterer's *message*, the keystone of his communicative intention (Stone, 2004b). We have thus dovetailed this chapter with the previous one.

3.4 Transmission of Content

As we have seen in the preceding chapters, discourse referents can occur in more than one context. Indeed, the history of a long-lived discourse referent is a tree spanning a number of different contexts. It follows that different contexts can contain the same piece of information, so long as those contexts share some drefs. For instance, the content $\langle \textit{bright}, u_{hes} \rangle$ can occur in any context in which the dref u_{hes} is available.

Communication, as we have seen, involves the transmission of a particular content (the content of the speaker's communicative intention) from the speaker's subjective context (see Chapter 4) to the hearer's subjective context.

More dramatically, suppose that the anaphoric traces of our current uses of *Hesperus* and *Phosphorus* extend back into prehistory, and, as a consequence, the corresponding drefs (call them u_{hes} and u_{pho}) were available in Babylonian times. Next consider the following contents:

- (8) a. $\langle =, u_{hes}, u_{hes} \rangle$
 b. $\langle =, u_{hes}, u_{pho} \rangle$

The information content (8a) *was* present in Babylonian times. Indeed, it is the trivial piece of information that the referent of u_{hes} is identical to the referent of u_{hes} . The content (8b) is different – evidently – and non-trivial, since it is the information that the drefs u_{hes} and u_{pho} co-refer. This information was not present in Babylonian contexts. It was (allegedly) discovered by Pythagoras, subsequently disseminated around the globe, and eventually passed (always via contexts in which both u_{hes} and

u_{pho} are present) down to us.

On the other hand, an piece of content containing drefs that are not present in a context cannot (without further ado – see below) be transmitted to that context. Suppose (as in Shakespeare’s play) Thaisa meets Pericles, a prince of Tyre, who is disguised as a knight of the same name. Thaisa’s internal language contains two distinct symbols, both of which refer to Pericles, but which denote different drefs – call them u_{per} (the prince) and $u_{per'}$ (the knight). When she learns, shortly before her burial at sea, that Pericles the knight is in fact Pericles, Prince of Tyre, she cons a (non-trivial) piece of content: $\langle =, u_{per}, u_{per'} \rangle$. This content cannot be transmitted to a context in which either dref is unavailable (for instance the subjective context of a faithful servant accompanying Pericles through both his guises). This is an obvious consequence of the requirement of discursive isomorphism, but should sound intuitive; after all, what could Thaisa’s revelation mean to the retainer?

I should say a few words, at this point, about indefinites. Suppose a speaker would like to convey a message that contains a symbol s denoting a dref u , however the speaker suspects that the hearer doesn’t possess a mental symbol denoting u . Concretely, let’s suppose the hearer has never heard of the German critic and philosopher Walter Benjamin, and the speaker wants to tell him that Benjamin experimented with drugs on Ibiza. The best the speaker can do (since discursive isomorphism is compromised) is begin a new discourse role – bringing into being a new dref u' – and have the hearer form a new mental symbol that (initially at least) denotes u' . He might say, for instance, ‘You know, there was a German critic and philosopher in the twentieth century called Walter Benjamin’. Subsequent definites anaphoric to the indefinite will denote u' and will function semantically as existentially-bound variables; the speaker’s truth-commitments will be existential rather than particular, as will the information semantically conveyed to the hearer.²

²This claim is based on the judgment that the truth of the sentence above does not depend on Benjamin himself being a twentieth-century philosopher and critic, so long as there was someone of *that name* who was.

Nevertheless, it is possible in this way to put the hearer in communion with the extra-contextual discourse referent u .³ In particular, if the hearer learns enough to be able (in principle) to coordinate on u with arbitrary persons who themselves have u in their subjective contexts, then the hearer will count as having u in his subjective context (and thus the potential to have cognitive attitudes towards contents involving u). Despite the fact that the semantic content of the speaker's utterance does not include u , the hearer may still come into some knowledge involving u . There is no reason to think that the semantic content of an utterance will necessarily circumscribe the information one might glean from it.

3.5 Diagonalization

The sense of a proper name is grasped by everybody who is sufficiently familiar with the language or totality of designations to which it belongs; but this serves to illuminate only a single aspect of the *Bedeutung*, supposing it to have one. Comprehensive knowledge of the *Bedeutung* would require us to be able to say immediately whether any given sense attaches to it. To such knowledge we never attain.

Frege, On *Sinn* and *Bedeutung*

In this chapter, I replaced *referential isomorphism* by *denotational isomorphism* as the criterion on successful communication. I also replaced 'singular' content (Russellian propositions) with 'discourse' content (denotational trees). To conclude, I want to compare this chapter with its most successful competitor in the philosophical literature, Stalnaker's account of assertion and his notion of diagonalization (Stalnaker,

³Devitt (1974) is an early discussion of name-use as an 'ability' and of how this ability is passed on to others.

1979).

It is a testament to Stalnaker's ingenuity that he makes do, in his account of the pragmatics of information exchange, with a conservative notion of content. In fact, his propositions – sets of possible worlds – even efface the distinctions between logically equivalent Russellian propositions. For Stalnaker, there is just one necessary proposition – the set of all possible worlds – and one necessarily false one – the empty set. He also works with a more minimal contextual staging than that developed in Chapter 2. His discourse context is, once again, a set of possible worlds (the 'context set'). This is the set of 'live candidates' for the actual world, according to the presuppositions of the discourse participants.

According to Stalnaker, an assertive utterance U updates the context set by intersecting it with:

1. The proposition that U has taken place (*The Commonplace Effect*)
2. The proposition p expressed by U (*The Essential Effect*)

Stalnaker's account, though an evident simplification, makes precise certain intuitions related to information exchange. He views conversation as an enterprise whose object is to remove uncertainty present in the participants' presuppositions (i.e. to reduce, without eliminating entirely, the context set). An assertion with an ineffectual essential effect (one that expresses a proposition that is a superset of the current context set) can be of no use in this task – it is redundant, given the context. An assertion that conflicts with the presuppositions of the context (whose essential effect would eliminate all the worlds in the context set) is equally defective.

Only contingent propositions stand a chance of being assertible. A sentence that is true or false of necessity automatically fails one of the conditions above, as its semantic content is either trivial or contradictory with respect to *every* context set. However, certain necessarily true or false sentences (e.g. *Hesperus is Phosphorus*, or

more commonplace informative identities like *This is Bob*) are assertible and, indeed, informative. Stalnaker solves this problem by detaching the pragmatic effect of such utterances from their semantic content.

For Stalnaker, the two types of update listed above are not on a par. The commonplace effect always comes first (Stalnaker, 1998). The essential effect is more limited in scope. It is brought to bear whenever U expresses the same proposition at every possible world in the context set (meaning the participants presuppose a particular semantic content for U). If there is uncertainty in the context about the meaning of U , then the essential effect can't take place.

Instead, the pragmatic mechanism of diagonalization kicks in. The first step is the generation of a 'propositional concept' (a function from worlds to propositions, a.k.a. an accessibility relation on worlds) P from the utterance U . P maps each possible world in the context set onto the proposition expressed by U at that world (this will be determinate, since worlds are maximally specific). Rather than updating the context set with the semantic content of U , we update using the information in P . We determine, for each world w in the context set, whether the proposition that P projects at w includes w . If it does, then w remains in the context set, if it doesn't, w is ejected. The set of worlds that stay – all w such that $P(w, w)$ – is known as the *diagonal* proposition. Note that the diagonal proposition projected by U in a particular context set corresponds to the metalinguistic proposition that U is true (restricted to the context set).

It is worthwhile comparing Stalnaker's model to a Bayesian one. In a Bayesian model, there is no way to update on evidence that has a prior probability of zero. The Bayesian 'solution' is: don't attach zero probability to any proposition that might occur as evidence. Anything is possible, so one must leaven one's opinions with a pinch of uncertainty. Stalnaker's model is different, but the problem is the same. Unlike the Bayesian, he assumes there are propositions of which the conversational participants are certain (indeed, the necessarily true proposition is presupposed in

every context). However, his model doesn't fail in the problem cases (e.g. informative identities), thanks to the back-up mechanism of diagonalization.

The intuitive justification of diagonalization is, however, Bayesian in spirit. An informative identity sentence trips the diagonalization mechanism because there is some uncertainty in the context about *which* proposition the sentence expresses, which means that the reference of one of the terms must be in doubt (or, perhaps less plausibly, some other factor relevant to the computation of the sentence's meaning). The Bayesian idea in diagonalization is that, so as to avoid impossible updates, our semantic opinions must remain uncertain.

In contrast with the Bayesian, Stalnaker believes that uncertainty should be kept on a short leash. While uncertainty about reference is introduced whenever an informative identity (e.g. *Hesperus is Phosphorus*) is theorized about, Stalnaker nevertheless assumes that the reference of similar yet irrelevant terms (e.g. *Mars*) is not in doubt. This way of modelling things doesn't seem terribly realistic. If there is no difference between our grasp of the semantic facts pertaining to *Hesperus* on the one hand, and *Mars* on the other, why model them differently? Nor does referential uncertainty pertain to the names of far-off planets only. As the ancient conundrum of the Hooded Man reminds us, one's own brother can appear on either end of an informative identity.

Furthermore, the limited uncertainty model can provide only after-the-fact analyses. We can apply a Stalnakerian analysis to a discourse that has already happened, but never to one that is in progress, since we won't know ahead of time which terms will end up figuring in an informative identity. Realistic Bayesian models, as used in artificial intelligence, instill uncertainty at every point, so that they can be effective in reacting to any contingency.

So what would a Stalnakerian account of assertion look like if semantic uncertainty were modelled realistically? Since informative identities containing any constant⁴ can

⁴At least, until we have compared that constant with every other constant in the language.

crop up at any time, the context set must model the semantic interpretation function with robust uncertainty (‘robust’ in the sense that it cannot be (completely) eliminated – threatening the informativeness of future identity statements – by updating). Referential uncertainty will be passed up to the sentential level by compositionality, making it uncertain in the context which proposition a sentence expresses. It follows that no utterance will meet the preconditions for the essential effect, and diagonalization must be brought in at every turn.

On this alternative way of modelling the context set, the essential effect becomes a chimaera, and the pragmatic ‘intervention’ of diagonalization, the norm.⁵ An utterance U will now intersect the context set with:

1. The proposition that U has taken place
2. The proposition that U is true (i.e. the diagonal proposition derived from U)

Let’s move on to consider the merits of Stalnaker’s picture as an account of communication. First note that a speaker who has communicated that his utterance U is true has not necessarily communicated successfully. For the context set may contain worlds at which U expresses a proposition quite unlike the one the speaker had hoped to assert. An especially compliant hearer, who is nevertheless in the dark about the content of U , might readily accept the proposition that U is true (without, we would say, understanding the utterance).

Stalnaker’s own examples of diagonalization at work lump together cases of true communication (e.g. uses of informative identity sentences, such as the Zsa-Zsa Gabor example) with cases of evident miscommunication (as in the example where Bob utters ‘You are a fool’ and O’Leary thinks that he is addressing Daniels, rather than himself). Recourse to diagonalization therefore fails to line up with the intuitive demarcation of communicative success.

⁵This point is not new; cf. (Lepore and Stone, 2007).

In the account of this chapter, denotation trees corresponded to a communicable level of information content, by which we could compare an ancient Sumerian context to our own. On Stalnaker's account, while we have a standard of comparison of entire contexts (the worlds in their respective context sets); there isn't any means of comparing contexts proposition-by-proposition. The diagonal proposition expressed by a present-day utterance of *Hesperus is Phosphorus* doesn't correspond to anything Pythagoras claimed or Hammurabi failed to realise, since it makes reference to a contemporary utterance, something neither Pythagoras nor Hammurabi could have had an opinion about.

Stalnaker's theory of assertion (and diagonalization) was formulated with different aspirations to the account of communication in this chapter (though both have in common a treatment of Frege's puzzle). It doesn't aspire to distinguish acts of communication that intuitively succeed from those that fail. Nor does it – if extended in the manner I propose – allocate an intuitively harmonious content to acts of assertion. It succeeds in what may be fatally limited goals. Though it brings out the need to consider our hypotheses about reference as *uncertain*, the account cannot be extended to the full range of phenomena occurring under the umbrella of communication.

Chapter 4

Attitudes

We begin this chapter with Frege’s truth-conditional puzzle about attitude reports. I consider Putnam’s formal approach to the puzzle and Fine’s semantic version thereof. I demonstrate that the essential insight is captured neatly (and ‘intrinsically’) by discourse semantics, extended, as I have proposed, to cross-discourse anaphora.

Armed with Chastain (1975)’s understanding of an agent’s mental state as a type of context,¹ I treat attitude reports as attributions of ‘discourse’ content. I demonstrate how this predicts Frege’s intuitions in his puzzle. I also visit Geach’s phenomenon of intentional identity, and Kripke’s puzzle about belief.

I demonstrate empirically that names in attitude reports have two different readings, analogous to the *de re* and *de dicto* readings of definite descriptions. I adapt (indeed simplify) ’s analysis of the *de re* reading.

The full implications of a *de dicto* reading depend on the dref’s status. If it is free in the discourse context, the report implies that there is an anaphoric connection between the discourse context and the subject’s subjective context. If instead it is bound, then the claim amounts to an existential claim over drefs.

This complex range of options is captured with relative ease by an intensional version of the compositional semantics of (Muskens, 1996), provided in the Appendix.

¹Inverting DRT’s treatment of the discourse context as a mental state.

4.1 Attitudes and Truth

Last chapter, we accounted for the Fregean intuition that utterances of (1) and (2) communicate different contents (corresponding to the pairs below):

(1) Hesperus is bright

$\langle \textit{bright}, u_{hes} \rangle$

(2) Phosphorus is bright

$\langle \textit{bright}, u_{pho} \rangle$

Sentences (3) and (4) also communicate different contents, and this could be accounted for in the same way.

(3) Biron thinks Hesperus is bright

(4) Biron thinks Phosphorus is bright

However, there is a further difference between (3) and (4); notoriously, they have different (apparent) *truth conditions* – (3) could be true while (4) remained false, and vice versa. (1) and (2), by contrast, are truth-conditionally equivalent.

Now the content $\langle \textit{bright}, u_{hes} \rangle$ is true iff the referent of u_{hes} is bright; likewise, the content $\langle \textit{bright}, u_{pho} \rangle$ is true iff the referent of u_{pho} is bright. The two drefs are distinct, yet they co-refer; thus, while the contents of (1) and (2) differ, their truth conditions coincide.

Plainly, we can't reason from the fact that *Hesperus* and *Phosphorus* co-refer to the logical equivalence of (3) and (4). For example, we can't simply allocate the content $\langle \textit{believes}, u_{bir}, \langle \textit{bright}, u_{hes} \rangle \rangle$ to (3) and propose that it is true iff the referent of u_{bir} stands in the belief relation to the complex consisting of the referent of u_{hes} and the property of being bright. For if we do likewise for (4), then we do not predict the truth-conditional difference between the two reports.

There are three approaches of note to this puzzle. One is just to deny that there is a truth-conditional difference between (3) and (4). It devolves upon the proponent of this view to explain our intuitions to the contrary, perhaps as a pragmatic side effect (Salmon, 1986; Soames, 2002). We return to this approach in §4.7.

Another approach is Frege's. He assumes that (3) and (4) have different truth conditions, and suggests that it is because the names falling in the scope of the attitude verbs contribute something other than their customary reference to the determination of those truth conditions (on his view they contribute their customary 'sense').

The third approach is due to Putnam (1954). He agrees with the intuition that the attitude reports are non-equivalent, but departs from Frege in maintaining that names only make a contribution of their reference to the truth conditions of sentences. How does he reconcile these two claims? By discovering a loophole in the principle of compositionality (below):

- (5) The meaning of a complex expression is a function of:
 - a. the meanings of its constituents
 - and
 - b. its structure

Part (b) of the principle is meant to allow *Man bites dog* to differ from *Dog bites man*. Putnam uses it to wrangle a truth-conditional difference between attitude reports that differ by a substitution of co-referring names (and so, like that pair, receive the same contributions from their respective constituents).

(5) is usually thought of as a high-level constraint that has some ground-floor implementation. For instance, Montague (1973) organizes the constituents of a complex expression into a tree structure, and computes the denotation of a node (if it is not a leaf) from its daughters by function application.

Putnam, unfortunately, doesn't provide an implementation of his scheme. He assumes that a sentence has a 'logical structure' which, among other things, specifies

which positions in the sentence are filled by the same term. For instance, (6a) and (6b) differ in logical structure, since the former, but not the latter, has the same term filling both argument slots. On the other hand, (6b) and (6c) have the same logical structure.

- (6) a. $R(a, a)$
 b. $R(a, b)$
 c. $R(c, d)$

Putnam, responding to Mates (1952), had attitude report pairs like (7a) and (7b) in his sights:

- (7) a. Biron thinks Hesperus is Hesperus
 b. Biron thinks Hesperus is Phosphorus

The reports differ in logical structure, since in (7b) the arguments of the embedded verb are different, whereas in (7a) they are the same. This leaves room for a truth-conditional difference between the two reports. Note that unwanted ‘room’ is also left for a truth-conditional difference between *Hesperus is Phosphorus* and *Hesperus is Hesperus*. It is left to the semantic implementation to explain why (7a) and (7b) differ while the latter do not.

Putnam’s account does not *prima facie* allow for a truth-conditional difference between our original examples (3) and (4). For in each case the embedded predicate takes only one term, and so (it would seem) their logical structures must coincide.²

While (3) and (4) coincide in logical structure, they differ in *discourse structure* (Chapter 1). To extend Putnam’s general gist to the examples we are interested in, we might adjust the principle of compositionality as follows: the meaning of a sentence is a function of the contributions of its parts and its *discourse structure* (where the discourse structure of a sentence includes the anaphoric resolutions of its component

²Soames (1988, fn. 20) attributes this point to David Lewis.

expressions – how it ‘plugs in’ to the surrounding discourse).

This might seem an abuse of the principle of compositionality, whose august compositors were surely only thinking (when inscribing (5b)) of structure *local* to the sentence; and indeed were endeavouring to keep the factors relevant to the determination of meaning within reasonable bounds. In the end, however, we will treat the problem without tinkering with the principle of compositionality, and will factor in discourse-sensitivity in an entirely Montagovian manner (after Muskens, 1996).

I would also mention William Taschek (1998), who extends Putnam’s notion of logical structure to cover our example. The logical structure of a sentence, for Taschek, consists of how it ‘plugs in’ to a surrounding *argument*. Two sentences have the same logical structure if and only if they are interchangeable in the context of any argument. Since (8) is valid and (9) is not, and they differ (only) by the substitution of *Phosphorus is bright* for *Hesperus is bright*, it follows that those sentences have different logical structures.

- (8) Hesperus is bright
 Hesperus is yellowish

 Hesperus is bright and yellowish
- (9) Phosphorus is bright
 Hesperus is yellowish

 Hesperus is bright and yellowish

Note a few things. Since it is difficult to find anything that follows by logic from attitude reports, it is nigh on impossible to illustrate this notion of logical structure with examples that are relevant. Keep in mind that we don’t particularly want there to be room for *Hesperus is bright* and *Phosphorus is bright* to differ in their truth

conditions.

Moreover, Taschek's account of logical structure must itself be underwritten – presumably by an account of discourse structure that includes anaphoric connections among names. Logical structure cannot simply be read off of surface form (as he emphasises in fn.7). (10) is almost identical on the surface to (8), and yet it is invalid:³

- (10) Pericles is a knight
 Pericles is a prince

 Pericles is a knight and a prince

It is not abundantly clear what (if not discourse structure), distinguishes (10) from (8). Taschek himself views logical structure as primitive, and in excess of sentence meaning, but here he is in the minority.

Kit Fine makes use of the same line of argument against the formalist approach. He champions a 'relationist' semantics that treats the facts Putnam relegates to logical structure as semantic ones (of 'representing as the same'). He writes:

I take it that we all have some intuitive grip on this notion of coordination or *representing as the same*. But a good test of when an object is represented as the same is in terms of whether one might sensibly raise the question of whether it *is* the same. An object is represented as the same in a piece of discourse if no one who understands the discourse can sensibly raise the question of whether it is the same. Suppose that you say 'Cicero is an orator' and later say 'Cicero was honest', intending to make the very same use of the name 'Cicero'. Then anyone who raises the question of whether the reference was the same thereby betrays his lack of understanding of what you meant. (Fine, 2007, 40)

Fine proclaims an additional layer of 'relational semantics', which, in conjunction with the usual (i.e. static) referential semantics, assigns 'coordinated' sequences of propositions to discourses. A coordinated sequence of propositions is a sequence of (Russellian) propositions plus a relation defined over the occurrences of objects in

³Thaisa's failure to recognise Pericles the knight as Pericles, Prince of Tyre is not a logical error.

those propositions. The relation holds between any two object-occurrences that are ‘represented as the same’ in the discourse. Coordination can extend, according to Fine, between discourses, and also between propositions expressed by utterances and those entertained in the mind. Fine exploits this versatility to solve Frege’s puzzle in a myriad of forms, and develops a particularly baroque account of Kripke’s tale of Paderewski that is, at the same time, a *tour de force*.

Fine distinguishes his approach from Putnam’s by its semantic, as opposed to formal, basis:

The relationist understanding of the phenomenon requires two further theses. First, the phenomenon is indeed semantic. When a piece of discourse represents an object as the same, then this is a semantic feature of the expressions by which reference to the object is made. Second, the phenomenon is essentially relational; there are no intrinsic semantic features of the individual expressions in virtue of which they represent the object as the same. (Fine, 2004, Lecture 3)

In the second point above, Fine distances himself from ‘intrinsicist’ accounts. It is clear, from what he goes on to say, that he is thinking of sense theorists:

The more serious challenge is from the intrinsicist. He agrees that the phenomenon is semantic but argues that when two expressions represent an object as the same it is because they represent that object in the same way. . . .

There are two main objections to the intrinsicist. The first is that his belief in these ways of representing the object is a myth. We might think of them as the ghost of images under the old imagist theory of meaning; these images may have been drained of their psychological content but they still operate as intrinsic aspects of meaning. The other objection is that these ways of representing an object are not, in any case, able to do the job required of them since, even if the ways were the same, there would still arise the question of whether they were represented as the same. (Fine, 2004, Lecture 3)

The arguments Fine weighs against the sense theorist are familiar (they are related to the arguments brought against Asher in Chapter 3). More importantly, we know that a different sort of intrinsicism is possible, one that makes use of a structured notion of *context*, and encapsulates anaphoric dependencies in denotations that

are defined against this structure. Discourse semantics prefers to translate Fine's relational superstructure into the context, rather than have it further complicate the products of linguistic expression.

The beauty of Reinhard Muskens' orchestrated détente between Montague semantics and DRT (see Muskens, 1996, and the Appendix) is that Putnam's original insight, semanticized by Fine, may be incorporated into an 'intrinsic' denotational semantics whose notion of grammatical structure is conservative (i.e. Montagovian sub-sentential structure, plus sentence concatenation).⁴ The outré ideas that semantic calculations are sensitive to 'logical structure' in addition to grammatical structure, or that denotational semantics must be overlain with an extra tier of relational semantics, are entirely dispensable. On the discourse semantic approach, we only require attitude verbs to be sensitive to the denotation of the clause in their scope, since that denotation *encapsulates* the relevant discourse structure.

The discourse semantic treatment has the effect of annexing Putnam's approach to Frege's, since a name in attitude-scope no longer contributes its referent to the calculation of truth conditions (instead it contributes a *dref*). A standard objection to Frege is the Davidsonian one of declaring such a semantics 'uninnocent'. Note that on Frege's account, an expression in the scope of an attitude verb *refers* to its customary sense. Davidson declared it absurd to think that, for instance, the name *Aristotle* no longer refers to Aristotle when it occurs in the sentence *Biron thinks Aristotle was wise*.

The present account is not open to this charge. *Aristotle* still refers to Aristotle wherever it appears. I make a distinction between denotation and reference, where denotation is an expression's semantic contribution, and reference is determined by that denotation. An expression, no matter where it appears, always contributes its denotation to the denotation of larger expressions of which it is a part.

⁴This is not particularly to endorse such a notion of textual structure (cf.).

While often – though not in attitude reports – the truth-value of a sentence will *only depend on* the referent determined by an expression’s denotation (i.e. the dref’s referent), this has to do with the rules for determining the truth of a sentence-level denotation (see §4.2 below), and is not because the expression *contributed* its referent in place of its denotation on that occasion.

It seems reasonable to me to preserve natural sounding judgments – such as that *Aristotle*, as it occurs in *Biron thinks Aristotle was wise*, refers to Aristotle – and I have endeavoured to do so. However it is not my brief to maintain the truth of ingrained yet quasi-theoretical statements like: the semantic contribution of *Aristotle* in the sentence *Biron thinks Aristotle was wise* is its referent – which raise a host of difficulties.

4.2 The Semantics of Attitude Reports

The next question is: what does it take for an attribution like (3) to be *true*? On what sorts of occasion does Biron count as believing the content $\langle \textit{bright}, u_{hes} \rangle$ (i.e. the content attributed to him in (3))? After the discussion in Chapter 3, this should be straightforward. First, Biron must have some private symbol, *birons-hesperus*, that is connected in an appropriate way⁵ with the dref u_{hes} . Second, he should have a mentalese sentence in his ‘belief box’ such as *high-effulgence(birons-hesperus)*, where *high-effulgence* is Biron’s concept designating the property of being bright.⁶

I follow Chastain (1975) in treating an agent’s cognitive state as a context, on a par with the socially-constructed contexts discussed in Chapter 2. An agent’s

⁵More carefully (and reprising the discussion from §3.4) this symbol must be associated with an ability to coordinate with others on the dref u_{hes} – i.e. to produce and interpret utterances that denote this dref. This participation condition prevents agents with ‘fractured’ concepts (like Peter’s concept of Paderewski – see §4.6) from possessing the relevant drefs.

⁶In all likelihood, we should treat property-level concepts as denoting (higher-type) drefs. After all, property-level expressions (at least those belonging to open classes) are introduced and passed on in much the same manner as names. Since I have not done the necessary empirical work to properly defend this claim, I leave it only as a (very plausible) hypothesis.

personal context is a DRS (with a universe of bound drefs, atomic conditions, and sub-DRSs corresponding to embedded contexts). The drefs free in the DRS – such as u_{hes} in Biron’s context – correspond to those mental symbols that are connected to anaphoric trees external to the context. Finally, the conditions of the DRS correspond to particular mentalese sentences: Biron’s private sentence *high-effulgence*(*bironshesperus*) shows up in his personal context as a condition constraining the reference of u_{hes} .

Let’s compare, formally, the two sentences:

(11) Hesperus is bright

(12) Biron thinks Hesperus is bright

(11) is true iff the referent of the dref u_{hes} is bright (at the actual world). I introduce the assignment function I^* which maps each dref onto its (actual) referent. If we additionally let w^* designate the actual world, we can perspicuously render the truth conditions of (11) as:

$$bright_{w^*}(I^*(u_{hes}))$$

Next consider (12). This sentence is true iff Biron believes the content $\langle bright, u_{hes} \rangle$ (i.e. his belief context contains a condition requiring the referent of u_{hes} to be bright). This in turn is true iff all pairs of assignments of referents to drefs and worlds *consistent with what Biron believes* verify the content.

If we let I range over possible assignments of referents to drefs, and let w range over possible worlds, then (3) is true iff:

$$\text{For all pairs } \langle I, w \rangle \text{ consistent with what Biron believes, } bright_w(I(u_{hes}))$$

Note that we predict Frege’s intuition that there is a truth-conditional difference between the attitude reports (3) and (4), since it is possible for Biron’s context to

contain a condition requiring u_{hes} to denote something bright without carrying an analogous condition for u_{pho} . For Biron might have in his ‘belief box’ the mentalese sentence *high-effulgence(birons-hesperus)* but not *high-effulgence(birons-phosphorus)*.

4.3 Indefinites and Attitude Reports

An indefinite always introduces a new dref. This means that the dref denoted by an indefinite won’t appear in any context *earlier than* the context in which the indefinite was uttered. We can use this clear consequence of the theory to raise a problem for the semantics of attitude reports just sketched. For though (13) may be true, its evident consequence (14) cannot be, if the theory is on track:

(13) Hammurabi thought Hesperus was bright

(14) Hammurabi thought something^{*u*} was bright

The rationale for this startling claim is that, while the denotation of *Hesperus* – u_{hes} – is free in the present context, and so could be a survival from pre-history, the dref u is *freshly minted*, and so *a fortiori* could not have occurred in Hammurabi’s ancient subjective context.

To solve this problem, some wiggle-room must be introduced into the truth definition. In effect, I treat indefinites as existentials *over drefs*, so that (14) is true iff there is a dref δ such that Hammurabi’s believes the content $\langle \textit{bright}, \delta \rangle$. This repair means that (13) once again entails (14). Dynamically speaking, indefinites still introduce new drefs, however the truth definition allows drefs introduced in the discourse context to be *isomorphically mapped* to drefs outside the context (while those drefs that are ‘free’ in the context are left to represent themselves).

Let f range over ‘schemes of replacement’ (i.e. automorphisms on the set of drefs). (14) is true iff:

For some f , for all pairs $\langle I, w \rangle$ in Hammurabi's belief set, $bright_w(I(f(u)))$

4.4 The Dynamics of Attitude Reports

Take the (marked-up) sentences:

(15) Something ^{u} is _{ω} bright

(16) Biron thinks _{ω'} something ^{u} is _{ω'} bright

Let's write the update contributed by (15) as:

$$[u \mid bright_\omega(u)]$$

This update may be cleft into two parts. First $[u]$, and then $[bright_\omega(u)]$ (contributed by the subject and predicate, respectively).

Due to the presence of *de re* readings (cf. §4.7), we follow the usual route (Hintikka, 1969; Cresswell, 1990) of trading in an operator syntax that perspicuously renders different layers of embedding (i.e. the usual DRT syntax) for a *binding* syntax that explicitly (or 'extensionally') represents the underlying model theory, in this case quantifying over pairs of assignments and worlds. In particular, all atomic conditions are indexed by a 'marker' (ω, ω' , etc.).

An *infostate* is a gizmo that maps drefs to drefs (like f above) and also maps 'markers' to *assignment-world pairs*.

The (non-deterministic) update $[u]$ ('random reassignment') corresponds to the relation that holds between any infostates i and j for which i only differs from j in (possibly) which dref it maps u onto. $[u]$ indicates that u is not free in the context (i.e. u is only a proxy for some other dref when it comes to truth conditions).

Allow me to introduce a piece of notation: $l(\langle \alpha, \beta \rangle) = \alpha$, and $r(\langle \alpha, \beta \rangle) = \beta$.

The update $[bright_\omega(u)]$ maps i onto *itself* iff $bright_{r(i\omega)}(l(i\omega)(iu))$ – i.e. just in case the referent assigned by the *left* member of the assignment-world pair $i(\omega)$ to

the dref $i(u)$ is bright at the world that is the *right* member of $i(\omega)$. It is a *test* to check whether u 's replacement under i occurs in a condition constraining its referent to be bright in the context indicated by the marker ω .

Finally, ‘;’ stands for ‘sequencing’ (i.e. relation composition). $[u]; [bright_\omega(u)]$ is the composition of the two relations just mentioned. We also write this as $[u | bright_\omega(u)]$. Note that a discourse context is just a DRS of the form: $[drefs | conditions]$.⁷ To the left of the line is the list of drefs bound in the context, and to the right is the list of conditions occurring in the context. The result of updating a context D with a DRS D' is $D; D'$.

Summarizing, $[u | bright_\omega(u)]$ denotes the relation that holds between infostates i and j such that i differs from j at most in what it maps u onto and $l(j\omega)(ju)$ is bright at $r(j\omega)$.

Next, consider the update contributed by (16), which I write as:

$$Bel_\omega^{\omega'}(u_{bir}, [u | bright_\omega(u)])$$

Once again, this update splits into two. First, the update $Bel_\omega^{\omega'}(u_{bir}, [u])$ which functions just like $[u]$ despite being embedded, since $[u]$ is not indexed by a marker. Second, $Bel_\omega^{\omega'}(u_{bir}, [bright_\omega(u)])$ which, like $[bright_\omega(u)]$ is a test (i.e. a restriction of the identity map). It maps i onto itself iff for every assignment-world pair $\langle I, w \rangle$ in the belief set assigned to the dref $i(u_{bir})$ at the assignment-world pair $i(\omega)$, $I(j(u))$ is bright at w .

How do we transition from updates (associated compositionally with English sentences in the Appendix) to the truth conditions adduced earlier? A DRS D is true at an infostate i iff there is some j , such that iDj (Groenendijk and Stokhof, 1991). Let i^* be the union of the identity map on the set of drefs and the function that takes any world-marker to the pair $\langle I^*, w^* \rangle$ (i.e. the actual assignment and world). D is true *simpliciter* iff it is true at i^* .

⁷We will supplement this shortly with a means for representing embedded contexts.

For instance, $[u \mid \text{bright}_\omega(u)]$ is true iff $\exists j(i^*[u \mid \text{bright}_\omega(u)]j)$. That is to say, there must be some dref δ (corresponding to the value of $j(u)$) such that $I^*(\delta)$ is bright at w^* . The dref u_{hes} would be a possible witness to this claim.

$Bel_\omega^\omega(u_{bir}, [u \mid \text{bright}_\omega(u)])$ is predicted to be true iff there is a dref δ such that for every assignment-world pair $\langle I, w \rangle$ in the belief set assigned to the dref u_{bir} at the assignment-world pair $\langle I^*, w^* \rangle$, $I(\delta)$ is bright at w .

I have treated contexts both as DRSs (for updating purposes) and also as sets of assignment-world pairs (in the treatment of attitude reports). We therefore require some means of extracting sets of assignment-world pairs from DRSs, so that we can look at the DRS corresponding to an individual's belief context (say) and determine the assignment-world pairs consistent with that DRS. There is only one free marker in any well-formed DRS (see Appendix), representing 'reality'. Call it ϖ . We determine the set of assignment-world pairs consistent with a DRS D to be $\{\pi \mid \exists j(i^{*[\varpi/\pi]} D j^{[\varpi/\pi]})\}$.⁸

I should also say something about the connection between denotational trees (or discourse 'propositions') and DRSs. In the Appendix, we will see how to combine the denotations of the basic components of a sentence (i.e. the leaves on a denotational tree) into a DRS. Provisionally, we can make the connection via assignment-world pairs. An agent who believes the content $\langle \text{bright}, u_{hes} \rangle$ is one for whom the assignment-world pairs consistent with what he or she believes all *verify* that content (i.e. for all $\langle I, w \rangle$ consistent with what A believes, $I(u_{hes})$ is bright at w). It follows that the set of pairs consistent with the DRS representing A 's belief state is a subset of the pairs that verify that content. This in turn means that (assuming once again that ϖ is free in the DRS), A 's belief context is equivalent to a DRS of the form: $[\dots \mid \dots, \text{bright}_\varpi(u_{hes}), \dots]$

⁸Where $i^{[\omega/\alpha]}$ is the infostate just like i except in mapping ω onto α .

4.5 Hob and Nob

We now have all the formal components in place to deal with Geach's famous Hob-Nob example (Geach, 1967). Do I dare to Peter Geach?

- (17) Hob thinks a^u witch has blighted Bob's mare, and Nob wonders whether she_u killed Cob's sow.

Our intuitions about the sentence are very robust. There are no witches, so there couldn't be an actual witch that (17) is about (this tunes out the otherwise available *de re* reading). Intriguingly, while the witch must be a figment of their imaginations, the sentence is intuitively false in a situation in which Hob and Nob independently dream up a tormenting beldam. In that case, we prefer to say:

- (18) Hob thinks a^u witch has blighted Bob's mare. Nob is also wondering whether a^{u'} witch killed Cob's sow.

In his article, Geach suggests that in (17), but not (18), we have an intuition of a 'common focus' of Hob and Nob's beliefs:

Etymology is more often a hindrance than a help in philosophy, but in this case it may be a help to remember the metaphor that underlies the words 'intention' and 'intentional': '*intendo arcum in...*', 'I draw a bow at...' For a number of archers may all point their arrows at one actual target, a deer or a man (real identity); but we may also be able to verify that they are all pointing their arrows the same way, regardless of finding out whether there is any shootable object at the point where the lines of fire meet (intentional identity). We have intentional identity when a number of people, or one person on different occasions, have attitudes with a common focus, whether or not there actually is something at that focus. (Geach, 1967, 627)

Edelberg (1986) gives two scenarios that tease our intuitions out further. In the first, a superstitious Hob, concerned at the failing health of Bob's mare, informs his good buddy Nob of his suspicions: he thinks a witch has blighted her. Nob then wonders whether the witch Hob is on about could be responsible for *sus*-icide.

In an alternative scenario, Hob and Nob live in the same town, but are unacquainted. A salacious gazette to which they both subscribe runs a story about a witch, whom both Hob and Nob soon suspect of separate agricultural misdemeanours.

(17) is true in either scenario (though untrue in a third situation in which there is no common cause of their beliefs). Readers of this dissertation should already be conceiving of the difference in terms of anaphoric connection! In both of Edelberg's situations, Hob's concept for the witch belongs on the *same anaphoric tree* as Nob's concept (in the second tale, the anaphoric traces meet in the newspaper article). In a third situation, in which there is no such 'anaphoric triangulation', the sentence is infelicitous. It seems there is an external criterion on cross-attitude anaphora that mirrors the external criterion on successful communication discussed in Chapter 3.⁹ For the report to be felicitous, the subjects of the two reports must hold the relevant attitude *towards the same dref*. Thus, for (17) to be true, there must be some dref δ , such that Hob believes the content $\langle \textit{blighted-Bobs-mare}, \delta \rangle$ and Nob is wondering whether the content $\langle \textit{killed-Cobs-sow}, \delta \rangle$ is true.

The semantics of this chapter already predicts this difference between (17) and (18) (it arises out of the combination discourse-sensitive attitude verbs and indefinites that quantify over drefs). This empirical adequacy is the more theoretically virtuous for arising out of general concerns (related to communication, context and content), rather than the specific application of 'intentional identity'. The account's closest cousin is the counterpart-theoretic account of (van Rooy, 2000).¹⁰

(17) denotes the following update:

$$Bel_{\omega}^{\omega'}(u_{hob}, [u \mid \textit{witch}_{\omega'}(u), \textit{blighted-Bobs-mare}_{\omega'}(u)]);$$

$$Bel_{\omega}^{\omega''}(u_{nob}, [\textit{killed-Cobs-sow}_{\omega''}(u)])$$

⁹Dekker and van Rooy (1998) also make this connection.

¹⁰Van Rooy is to be commended for his insight (*contra* Edelberg) that the difference between (17) and (18) rests on an *external* condition, connected (at least in paradigm cases) to communication.

(17) is true iff there is some dref δ such that for every assignment-world pair $\langle I, w \rangle$ in the belief set of u_{hob} at $\langle I^*, w^* \rangle$, $I(\delta)$ is a witch at w , $I(\delta)$ blighted Bob's mare at w , and for every assignment-world pair $\langle I', w' \rangle$ in the belief set assigned to u_{nob} at $\langle I^*, w^* \rangle$, $I'(\delta)$ killed Cob's sow at w' .

Note that these truth conditions require Hob and Nob to have beliefs pertaining to the same dref, implying that the relevant beliefs are anaphorically connected. Thus Geach's intuition of 'intentional identity' is captured.

(18), on the other hand, denotes the update:

$$\begin{aligned} &Bel_{\omega}^{\omega'}(u_{hob}, [u \mid witch_{\omega'}(u), blighted-Bobs-mare_{\omega'}(u)]); \\ &Bel_{\omega}^{\omega''}(u_{nob}, [u' \mid witch_{\omega''}(u'), killed-Cobs-sow_{\omega''}(u')]) \end{aligned}$$

It follows that (18) is true iff there is some dref δ such that for every assignment-world pair $\langle I, w \rangle$ in the belief set assigned to u_{hob} at $\langle I^*, w^* \rangle$, $I(\delta)$ is a witch at w and $I(\delta)$ blighted Bob's mare at w , and there is some (possibly distinct) dref δ' such that for every assignment-world pair $\langle I', w' \rangle$ in the belief set assigned to u_{nob} at $\langle I^*, w^* \rangle$, $I'(\delta')$ is a witch at w' and $I'(\delta')$ killed Cob's sow at w' .

In short, it is predicted that (18), unlike (17), reveals nothing about anaphoric connections between Hob and Nob's beliefs. This concludes my discussion of the Hob-Nob example.¹¹

4.6 Paderewski

In Kripke's Paderewski puzzle (Kripke, 1979), the protagonist Peter hears of the musician-statesman Paderewski twice, in different circumstances, and comes to believe

¹¹There is a voluminous literature on the nuanced intuitions related to attitude-crossing anaphora, whose backbone is an excellent series of articles by Edberg. Unfortunately, the account just given doesn't predict, by itself, some of the intriguing empirical data discussed in this literature, including the asymmetry remarked on in (Edberg, 1986). Moreover, the literature even questions what I earlier called the 'robust' intuition that Hob and Nob's concepts must be connected by a path on an anaphoric tree (Edberg, 1992), and so challenges the necessity of the present analysis. I believe my analysis is substantially correct, but do not propose to ponder these developments here (I hope to do so in future work).

there are two persons answering to ‘Paderewski’ – one a musician, and the other a politician. The situation is analogous to one from literature. Cast your mind back to Chapter 3, where I claimed that the content of Thaisa’s realization that Pericles the knight was Pericles the prince (which I represented as the content $\langle =, u_{per}, u_{per'} \rangle$) could not be added to the (apocryphal) faithful servant’s subjective context. Likewise, Kripke’s character Peter has beliefs whose content is not transmissible beyond his own subjective context.

Specifically, there are two drefs – u_{pad_1} and u_{pad_2} – that occur in Peter’s subjective context and nowhere else. Discourse ‘propositions’ containing these drefs (for instance, $\langle politician, u_{pad_1} \rangle$ and $\langle musician, u_{pad_2} \rangle$) are perspicuous to Peter alone, and in particular cannot be attributed to him in a belief report, since it is impossible to pick out these contents from the discourse context (where the drefs are unavailable – i.e. they don’t have the minimum ‘familiar’ status).

Furthermore, Peter does not possess the common-currency dref that refers to the musician-statesman (call it u_{pad}). Why not? We can narrate the situation as follows. At his first introduction to the name (under its politician guise) he acquired the dref u_{pad_1} (remember, from §3.4, that one is always first introduced to a *new* dref, and only later, by learning the skill of coordination, acquires the one already in currency). The next time he heard the name, rather than connecting it with his old symbol, he forged a new symbol (thus betraying the ability he still lacked) and attached another new dref u_{pad_2} to it.

So long as his concept of Paderewski was ‘fractured’ (the terminology is Fine’s) in this way, his ability to coordinate with others on the dref u_{pad} was compromised. For an interlocutor could not rely on Peter accessing the same mental symbol on successive uses of the name *Paderewski*.¹² It follows that we cannot (truly) attribute

¹²I imagine I’ll be taken to task for the regrettable sin of *false precision*. Whereas dref possession is an all-or-nothing affair, coordination ability belongs on a gradient and is unduly affected by the circumstances and the nature and abilities of the other participants. As in other areas of language study, I must cordon off certain complications to develop a model that accounts for the phenomena I’m interested in. I hope you’ll submit to the necessary simplifications long enough to see the pay-off.

any beliefs involving the dref u_{pad} to Peter.

We can still say *something* about Peter's belief state. For instance, we can report that there are these two drefs and that Peter has certain beliefs about them. That is to say, we can *describe* his beliefs, without *specifying* their (discourse) content. This corresponds to how we tend to set up the case, anyway. The matter is set out using *indefinites* – see my own description in the first paragraph of this section – which I have given reason to think effectively quantify over drefs.¹³ Concretely, we might formalize the usual Paderewski *mise-en-scène* as the following update:

$$Bel_{\omega'}^{\omega'}(u_{pet}, [u u' \mid politician_{\omega'}(u), musician_{\omega'}(u), u \neq u'])$$

As you might expect, this update is true so long as there are two drefs δ and δ' such that Peter believes that the referent of δ is a politician, the referent of δ' is a musician, and that they refer to different individuals. This update is therefore true on my reckoning of Peter's situation, as laid out in the preceding discussion.

4.7 *De Re* and *De Dicto*

I'd like to return to the first response (Salmon, 1986; Soames, 2002) to Frege's puzzle, which was to claim that (19) and (20) do not, when properly understood, have different truth conditions.

(19) Biron thinks Hesperus is bright

(20) Biron thinks Phosphorus is bright

There is something to this claim; there are certain contexts in which (19) appears to *entail* (20). In fact, there is a valid reading of (19) and (20) on which they are equivalent.

¹³We might also, as Kripke does in his article, say that Peter believes conflicting things about Paderewski (the musical politician). In this case we are making a *de re* report (see the following section), and implying that Peter's drefs both *refer* to Paderewski (we might say that their reference results from their *deference* to the dref u_{pad} (cf. Putnam, 1975)).

Consider, in this connection, a possible problem for the foregoing account. It seems that I can successfully attribute a belief to someone without knowing *which* dref their concept manifests (that is to say: without making an anaphoric connection to the discourse context). If I think there are sentient beings on Venus, then I can attribute beliefs about Venus to them despite our having no dref in common.

In such cases it is the object – the referent of the dref invoked in the attitude report – that is relevant to the characterization of the belief state, rather than the dref itself. We must suppose that the subject is connected to the object by *some* means, but we can leave this nicety out of the report.

Speaking in this mode, (19) will be equivalent to (20), since u_{hes} and u_{pho} co-refer. Slightly less obviously, the argument (21) is invalid. This is because Biron could satisfy the premises with beliefs that single out Hesperus (a.k.a. Venus) via different drefs – drefs that he doesn't realise co-refer. For instance, he might believe the contents $\langle \textit{bright}, u_{hes} \rangle$ and $\langle \textit{yellowish}, u_{pho} \rangle$ and not be aware that u_{hes} and u_{pho} co-refer.

- (21) Biron thinks Hesperus is bright
 Biron thinks Hesperus is yellowish
-
- Biron thinks Hesperus is bright and yellowish

While there is a genuine reading of (19) on which the referent of *Hesperus* (rather than its denotation) is used to characterize Biron's belief, it isn't the only reading. I call the object-dependent reading *de re* (after a familiar fashion), and the dref-dependent reading (the reading assumed prior to this section) *de dicto*.¹⁴ Though

¹⁴Quine (1956) distinguishes between:

- (22) Ralph believes Ortcutt is a spy
 (23) Ralph believes ($\lambda x.x$ is a spy) of Ortcutt

The former entails the latter, and the latter has the formal consequences (and inconsequences) discussed in this section. I require, in addition, that (23) can be pronounced the same as (22).

Salmon and others go to great lengths to defend the uniqueness of the *de re* reading from apparent absurdity, I see no reason to emulate their intellectual gyrations.

The *de re* reading is analysed (along the lines of) as an existential quantification over *de dicto* attributions. (24) is true if and only if there is a dref δ that co-refers with u_{hes} and Biron believes the content $\langle \textit{bright}, \delta \rangle$.

(24) Biron thinks Hesperus (*de re*) is bright

$$Bel_{\omega}^{\omega'}(u_{bir}, [u \mid u =_{\omega} u_{hes}, \textit{bright}_{\omega'}(u)])$$

Note that (24) is entailed by the *de dicto* reading of (19) (and the fact that u_{hes} refers), but not vice versa.

The simplest (and most venerable) argument for the existence of the *de dicto* reading is that (25a) has an interpretation on which it attributes a non-trivial content to Biron:

(25) a. Biron thinks Hesperus is Phosphorus


b. $Bel_{\omega}^{\omega'}(u_{bir}, [u_{hes} =_{\omega'} u_{pho}])$

c. $Bel_{\omega}^{\omega'}(u_{bir}, [u \ u' \mid u =_{\omega} u_{hes}, u' =_{\omega} u_{pho}, u =_{\omega'} u'])$

While (25b), on which both names are read *de dicto* is non-trivial, the (doubly) *de re* reading (25c) is entailed by the trivial (26):

(26) Biron thinks Hesperus is Hesperus

Here is a further argument for ambiguity. Rosaline, Maria and Katherine are en route to a masked ball. Each lady anticipates (tepidly) the attentions of a different suitor, and each suitor has given his lady a favour to wear, by which he hopes to recognize her during the masque. The ladies, with malice aforethought, decide to swap favours and so cause the suitors to ‘woo contrary’. In this (counterfactual) version of *Love’s Labour’s Lost*, the exchange of favours goes as follows:



At the first *lavolta*, Biron – Rosaline’s swain – takes Katherine by the busk. Maria wryly comments, ‘Why—

(27) Biron thinks Katherine is Rosaline!

Her utterance is true, yet the following, in the same circumstances, would be false:

(28) Biron thinks Rosaline is Katherine

Not convinced? Reflect that (29) is true (and so (28) couldn’t be):

(29) Biron thinks Rosaline is Maria

On the assumption that the *is* in the sentences above is the *is* of identity (what other interpretation of *is* takes individual-level relata?), it follows that neither univocal account of proper names in attitude reports will suffice. For if *Katherine* and *Rosaline* are read the same way in (27) and (28) (i.e. either both *de re* or both *de dicto*), then by the symmetry of identity,¹⁵ each sentence entails the other.

Concretely, suppose all names occur *de dicto*. The *de dicto* reading of (27) states that Biron believes the content $\langle =, u_{kat}, u_{ros} \rangle$. If this is true, then by symmetry he also believes $\langle =, u_{ros}, u_{kat} \rangle$, the content attributed to him by the *de dicto* reading of (28). If instead names are perennially read *de re*, then (27) says that there are some drefs δ and δ' such that δ co-refers with u_{kat} and δ' co-refers with u_{ros} and Biron believes $\langle =, \delta, \delta' \rangle$. Again, by symmetry, he must also believe $\langle =, \delta', \delta \rangle$ and so verify the *de re* reading of (28).

¹⁵I’m (innocently enough) assuming Biron’s beliefs are closed under this inference

Once we avail ourselves of a *de re/de dicto* ambiguity for names, it is legitimate to suppose that the subject in each embedded clause is interpreted *de re*, while the object is interpreted *de dicto*. This yields the interpretations below:

(30) Biron thinks Katherine is Rosaline
 $Bel_{\omega}^{\omega'}(u_{bir}, [u \mid u =_{\omega} u_{kat}, u =_{\omega'} u_{ros}])$

(31) Biron thinks Rosaline is Katherine
 $Bel_{\omega}^{\omega'}(u_{bir}, [u \mid u =_{\omega} u_{ros}, u =_{\omega'} u_{kat}])$

There is a truth-conditional difference between (30) and (31), so interpreted. (30) states that there is some dref that refers to Katherine but which Biron thinks co-refers with u_{ros} . In the situation described, the witness for this existential is presumably the dref denoted by the deictic mental symbol (call it *birons-object-1*) that Biron is using to track Katherine at the ball.

(31), on the other hand, is true iff there is some dref that refers to Rosaline but which Biron thinks co-refers with u_{kat} . There is no such dref (Biron connects the ‘deictic’ dref that tracks Rosaline with u_{mar} , not u_{kat}), and so (31) is false.

This analysis lines up with the homonymous ambiguity for definite descriptions. Of the two sentences below, only the first (on its most natural interpretation) is true in a situation where Katherine entered the ballroom first, yet Biron thinks she entered last:

(32) Biron thinks the one who came in first is the one who came in last

(33) Biron thinks the one who came in last is the one who came in first

The difference between (32) and (33) is explained by a natural tendency to read the embedded subject *de re* and the embedded object *de dicto*. The readings are formally captured by the patterns of intensional binding below:¹⁶

¹⁶What I’m temporizing as a subject-object asymmetry must ultimately be understood in terms of information structure (Ramsey, 1925; Büring, 1997). Note that if we embed (33) in the following

(36) Biron thinks $_{\omega'}^{\omega}$ the one who came in first $_{\omega}$ is $_{\omega'}$ is the one who came in last $_{\omega'}$

(37) Biron thinks $_{\omega'}^{\omega}$ the one who came in last $_{\omega}$ is $_{\omega'}$ the one who came in first $_{\omega'}$

We can make the analogy between the readings of names and definite descriptions tighter if we think of (38b) as a simplification of (38c) (one more dref is introduced in the latter, but they are truth-conditionally equivalent). This allows us to treat the intensional binding as in (38a).

- (38) a. Biron thinks $_{\omega'}^{\omega}$ Katherine $_{\omega}$ is $_{\omega'}$ Rosaline $_{\omega'}$
 b. $Bel_{\omega'}^{\omega}(u_{bir}, [u|u =_{\omega} u_{kat}, u =_{\omega'} u_{ros}])$
 c. $Bel_{\omega'}^{\omega}(u_{bir}, [u|u =_{\omega} u_{kat}]; [u'|u' =_{\omega'} u_{ros}]; [u =_{\omega'} u'])$

Furthermore, this means we can translate occurrences of names uniformly by the scheme below:¹⁷

$$\text{Name}_{\omega} \rightsquigarrow \lambda F.[u|u =_{\omega} u_{name}]; Fu$$

4.8 Previous Accounts

This section is of necessity selective, since there have been more papers written about proper names in attitude reports than there are pages in this dissertation. Many of those accounts have certain features in common, and it will prove advantageous to bring out those features while ignoring myriad differences.

mini-dialogue it *can* sound true, which suggests that, in the proper circumstances, the subject is interpreted *de dicto* and the object *de re*:

- (34) a. Who does Biron think the one who came in last is?
 b. Biron thinks the one who came in last is the one who came in first

Note that the same goes for (28):

- (35) a. Who does Biron think Rosaline is?
 b. Biron thinks Rosaline is Katherine

¹⁷This analysis of names is argued for on independent grounds in (Brasoveanu, 2006, Ch. 3).

Many accounts begin by pointing out that attitude reports are context-sensitive. This is often accompanied by admission of the failure of what one might call ‘conventional compositionality’, the principle that the meaning of a complex utterance is arrived at by combining the *conventional* meanings of its parts. Some accounts violate the principle by having atomic expressions in the scope of an attitude verb contribute something other than their conventional meaning (e.g. Frege has expressions contribute their customary sense). In others, the principle of compositionality itself is violated (e.g. Crimmins (1992) enriches the meaning of the subordinate clause with ‘unarticulated constituents’).

In cases where there is no technical violation of conventional compositionality, the law may nevertheless be bent by the postulation of hidden syntactic structure (as with (Forbes, 1996)’s metalinguistic variables). Quite often, the exculpation for these posits is purely semantic, rather than syntactic (as Stanley (2002) stresses it must ultimately be).

Here are some representative passages:

... even if we eternalize the contained sentence and also rid the containing sentence of such sources of truth-value variation as inadequate descriptions, indicator words, and the like, still the whole [attitude report] may remain capable of varying in truth value from occasion to occasion...
(Quine, 1960, 218)

The idea that there is something ‘intrinsic’ to the content of Jeff’s sentence that would make it represent one but not another of the mediators of Odile’s belief is a mistake. ‘Twain’ can in principle represent any name, demonstrative, or indexical that Odile uses to refer to Twain. Of course, given contextual factors, there may be something that makes ‘Twain’ a more apt representative of ‘that one’ than other expressions the speaker might have used. (Richard, 1990, 135)

When we utter a belief sentence, we are talking *about* an agent’s ideas and notions, and these notions and ideas become unarticulated constituents of what we say. (Crimmins, 1992, 152)

In more detail, attitude attributions are taken to impute psychological relations towards propositions, which in turn I take to be ways of thinking of *situation-types*. There can be different ways of thinking of the same situation-type (‘STP’), since the latter are merely complexes of objects and properties, and an unambiguous attitude attribution adverts, via the

hidden ‘so’ [which Forbes postulates at logical form], to a particular one of the ways of thinking of the STP. (Forbes, 1996, 350)

Of the accounts worth taking seriously, many adhere to the following scheme: whereas the sentence S uttered in context c expresses the proposition p , the sentence ‘ A thinks that S ’ uttered in c expresses the proposition that A stands in a certain relation to another proposition, or *propositionoid*, p' . Whereas the proposition p is the result of applying the conventional character of S to the context c , p' is the product of a more esoteric ceremony. In particular, while proper names occurring in S contribute their usual (constant) characters to the determination of p , they may contribute something else (or something additional) to the determination of p' .

That something else might be syntactic: it could be a natural language form (Richard, 1990), or it could be a glyph of A ’s mentalese. It could be a ‘mental dossier’ belonging to A , or an internal ‘variable’ (van Rooy, 2000). It could be a sense, or an individual concept (Aloni, 2000). It could be some combination of the above.

Whatever the non-standard component, call it α , and whatever the particular mechanism settled on (be it free pragmatic enrichment or the resolution of a hidden pronoun), every account of this ilk is committed to an additional transaction between the speaker and hearer to coordinate on α . These accounts betray what one might call *processing uniformity*, the thesis that expressions are processed the same way both inside and outside attitude contexts. For when *Aristotle* is processed outside an attitude context, the speaker and hearer are only required to settle on the reference of the name; however, when it occurs *inside* such a context, the speaker and hearer must additionally settle on *Aristotle*’s contribution to the ‘propositionoid’, be it a notion, dossier or symbol.

I prefer this more general way of characterizing the views I am interested in (to the one employing ‘conventional compositionality’), since I don’t believe that names have conventional meanings. Instead, I think of names on the model of pronouns

(more generally, definites). While the semantic contribution of a name differs from occasion to occasion (too much for any one contribution to count as its conventional meaning), they are all processed in a uniform way, using the standard techniques of anaphora resolution. In my case, uniformity implies that, if the same factors relevant to anaphora resolution impinge on an occurrence of a name, then, whether it occurs in the scope of an attitude report or not, it will be resolved the same way (i.e. denote the same dref).

The proponents of DRT are led by their theoretical attachments to the view that anaphoric resolution is conducted differently inside and outside the scope of an attitude verb. DRT is a theory, at least in part, of how a sentence gets processed into an ongoing mental model. It handles the recursion of attitude reports by supposing that the audience's mental model contains embedded DRSs which are the audience's representations of other agents' mental models.

DRT is a subjectivist theory, and treats the drefs populating DRSs as private mental symbols. As a result, it makes no sense to have drefs occurring free in either the main DRS or a sub-DRSs representing another agent's cognitive state (since by definition private symbols cannot span contexts). A DRS representing an entire mental state must be 'proper', in the terminology of Kamp and Reyle (1993).

Let's say that an utterance of a sentence S contributes the update D to the main DRS. A report ' A thinks S ' would then update the sub-DRS representing A 's cognitive state (in the audience's model) using D . In particular, if definite expressions occur in S , they must be anaphorically resolved (bound) within the *sub*-DRS. They could not be bound in the main DRS, for they would then appear free in a sub-DRS representing a mental state, which is taboo in DRT.

When definite NPs occur in belief contexts, a presupposition of familiarity extends to the agent. . . So the recipient of a report containing a definite NP α should, according to the principle of definite reference identification, identify the discourse referent introduced by α with some discourse referent already declared in the universe of the delineated DRS representing

what the recipient already knows about the agent's beliefs.
(Asher, 1993, 120)

For example, while the pronoun *she* in (39b) is resolved to the indefinite antecedent in (39a), the same pronoun occurring in (40b) cannot be, since it must be resolved within the audience's representation of Nob's belief state. This requires accommodation, since – in this discourse – there is no indefinite scoped under the attitude verb to serve as an antecedent.

- (39) a. A witch is at large.
b. She killed Cob's sow.

- (40) a. A witch is at large.
b. Nob thinks she killed Cob's sow.

Even though it is in many ways a different approach to attitude reports, DRT, in common with many static theories, attaches a *processing differential* to the scope of a report verb. By comparison, the theory developed in this dissertation treats the pronouns in both discourses above as anaphoric to the preceding indefinites. No alternative resolution (or accommodation) is necessary in (40b).

Processing uniformity results in a neater, more modular theory. Anaphora (and reference) can be kept apart from theorizing about attitude reports, and their principles may be presumed to have general application (e.g. we don't need to provide a special version of centering theory for use in attitude reports). Also, data within attitude reports has the same evidential bearing on theories of anaphora (or reference) as data occurring outside.

In this chapter, we have witnessed the numerous benefits of uniformity.

On my account, the *de re* and *de dicto* readings of proper names are completely analogous to those of definite descriptions, in that they correspond to the same patterns of *intensional binding*. On other accounts, the situation is not so happy. We may divide the accounts I've been discussing into referential and quantificational camps,

depending on whether the extra piece supplied to the ‘propositionoid’ by a name is an object of some sort, or a piece of quantificational apparatus (generally a restricted existential quantifier).

Referential accounts are only able to capture *de dicto* readings, as each name in attitude-scope supplies a *particular* notion, dossier, etc. (analogous to the particular dref employed in *de dicto* characterization). Quantificational accounts are more flexible. Indeed, they are too flexible. For instead of a binary distinction, there is a continuum of different restrictions that could be supplied by context, ranging from the *de re* (no restriction) to the *de dicto* (restriction to a unit set). The analogy with definite descriptions (for which there are two readings) is lost.¹⁸

Finally, we saw in §4.5 that the social connection thought to be implied by Geach’s Hob-Nob example (and cross-attitudinal anaphora generally) fell out on our (uniform) semantics. The same result in a DRT-based account, where anaphora must be resolved locally to the cognitive state reported on, requires special stipulation (van Rooy, 2000).

¹⁸Cf. Crimmins (1992), who argues that the *de re* reading is empirically unattested.

Appendix

Introduction

Muskens' Compositional DRT is a formalism combining the dynamic insight that the meaning of a sentence is a contextual update with Montague's application of the typed lambda calculus to sub-sentential compositionality. On a static Montagovian semantics, the meaning of a sentence is a proposition, and the meaning of a name (simplifying somewhat) is an entity. Thus the meaning of a predicate is a function from entities to propositions, captured by a lambda expression of the form $\lambda x.\phi$. In Muskens' dynamic system, the meaning of a sentence is an update, while the meaning of a name is (simplifying somewhat) a discourse referent. The meaning of a predicate is therefore a function from drefs to updates, captured by a lambda expression of the form $\lambda\delta.D$.

Syntax

There will be two 'syntaxes' corresponding to two different notions of possible expression. On the one hand, there is the syntax of expressions allowed in the typed lambda calculus that will form the basis of the system ICDRT (Intensional Compositional DRT). On the other hand, there is the more circumscribed syntax of expressions generated by the lexicon (of which I provide only a suggestive soupçon). It is in the second sense in which statements like 'there can be only one free marker in any DRS'

from Chapter 4 should be understood.

Let there be five basic types:

$$B ::= e \mid d \mid m \mid w \mid t$$

These are the types of *entities*, *drefs*, *markers*, *worlds* and *truth-values*.

There are three ways of forming complex types. Two are given by the grammar:

$$C ::= B \mid C \rightarrow C \mid C \times C$$

$T_1 \rightarrow T_2$ is the type of (total) functions from entities of type T_1 to entities of type T_2 . $T_1 \times T_2$ is the type of ordered pairs whose left member is of type T_1 and whose right member is of type T_2 .

Moreover, if T_1, T_2, T_3 and T_4 are types and $T_1 \neq T_2$, then $(T_1 \rightarrow T_3) + (T_2 \rightarrow T_4)$ is a type. Intuitively, $(T_1 \rightarrow T_3) + (T_2 \rightarrow T_4)$ is the type of functions that are the union of some function of type $T_1 \rightarrow T_3$ and some function of type $T_2 \rightarrow T_4$ (the requirement that the domains of the unionized functions be distinct ensures that the union is indeed a function).

I will adopt the following abbreviations:

$$\begin{aligned} \mathbf{g} &:= d \rightarrow e \\ \mathbf{s} &:= (d \rightarrow d) + (m \rightarrow (g \times w)) \end{aligned}$$

The first, \mathbf{g} , is the complex type of *assignments*, or total functions from drefs to entities. The second, \mathbf{s} , is the (more) complex type of *infostates*, or unions of automorphisms on the set of drefs and functions from markers to assignment-world pairs.

The syntax includes the following basic expressions ($e : T$ means e is an expression of type T).

Variables –

$$\begin{aligned}
x : e \quad \delta, \delta' : d \quad \varpi, \varpi' : m \quad w, w' : w \quad I, I' : g \quad h, i, j, k : s \\
C, C' : s \rightarrow t \quad (\text{condition}) \quad D, D' : s \rightarrow s \rightarrow t \quad (DRS) \\
F, G : d \rightarrow s \rightarrow s \rightarrow t \quad (\text{dref} \rightarrow DRS) \\
P, Q : (d \rightarrow s \rightarrow s \rightarrow t) \rightarrow s \rightarrow s \rightarrow t \quad ((\text{dref} \rightarrow DRS) \rightarrow DRS)
\end{aligned}$$

Constants –

$$\begin{aligned}
Venus : e \quad u, u', \text{ etc.} : d \quad \omega, \omega' : m \\
farmer, donkey : w \rightarrow e \rightarrow t \quad (\text{static predicate}) \\
owns, beats : w \rightarrow e \rightarrow e \rightarrow t \quad (\text{static relation}) \\
bel : e \rightarrow (g \times w) \rightarrow (g \times w) \rightarrow t \quad (\text{entity} \rightarrow \text{relation on assignment-world pairs})
\end{aligned}$$

More complex expressions may be formed from the basic ones, as follows:¹⁹

If $e' : T_1$ and if $e : T_1 \rightarrow T_2$ or $e : T_1 \rightarrow T_2 + T_3$ or $e : T_3 + T_1 \rightarrow T_2$, then $e(e') : T_2$

If v is a variable of type T_1 and $e : T_2$, then $\lambda v.e : T_1 \rightarrow T_2$

If $e : T_1$ and $e' : T_2$, then $\langle e, e' \rangle : T_1 \times T_2$

If $e : T_1 \times T_2$ then $l(e) : T_1$ and $r(e) : T_2$

If $e : T_1 \rightarrow T_3$ and $e' : T_2 \rightarrow T_4$ and $T_1 \neq T_2$, then $e \cup e' : (T_1 \rightarrow T_3) + (T_2 \rightarrow T_4)$

Semantics

An ICDRT frame \mathcal{F} is a set of sets labelled by type. There are five disjoint sets \mathcal{D}_e , \mathcal{D}_d , \mathcal{D}_m , \mathcal{D}_w , \mathcal{D}_t constituting the domains of the basic types. $\mathcal{D}_t = \{\top, \perp\}$; the others are non-empty. For all T_1, T_2 , the domain $\mathcal{D}_{T_1 \rightarrow T_2}$ is the set of all total functions from \mathcal{D}_{T_1} to \mathcal{D}_{T_2} ; $\mathcal{D}_{T_1 \times T_2}$ is the set of all ordered pairs whose left member is of type T_1 and whose right member is of type T_2 ; and $\mathcal{D}_{T_1 + T_2} = \{f \cup g \mid f \in \mathcal{D}_{T_1} \wedge g \in \mathcal{D}_{T_2}\}$.²⁰

¹⁹Relevant cases to prove are – $i(u) : d$ and $i(\omega) : g \times w$.

²⁰For technical reasons, we require that for every $x \in \mathcal{D}_e$ there is a dref δ that refers to it (i.e. $x = I^*(\delta)$). Assume this as an axiom constraining M . The elements δ of \mathcal{D}_d that refer to some entity, but are not ‘objects of thought’ (i.e. δ is not a component of any content that is believed by someone, or believed to be believed by someone, etc.) are technical vehicles that do not correspond

A model M is a quadruple $\langle \mathcal{F}^M, \|\cdot\|^M, w^M, I^M \rangle$, where $\|\cdot\|^M$ is a function taking each primitive non-logical constant c of type T to an element $\|c\|^M$ of \mathcal{D}_T . w^M and I^M (the actual world and the actual assignment) are members of \mathcal{D}_w and \mathcal{D}_g , respectively. $\|w^*\|^M = w^M$, $\|I^*\|^M = I^M$ and $\|i^*\|^M = \|\lambda\delta.\delta \cup \lambda\varpi.\langle I^*, w^* \rangle\|^M$. Furthermore, for each constant c of type d or m , the value $\|c\|^M$ must be distinct from the interpretation of any other constant of that type (cf. Muskens, 1996).

A ‘meta-assignment’ θ is a function on the domain of variables that maps each variable v of type T onto an element of D_T . $\theta^{[v/\alpha]}$ is the meta-assignment that takes v to α but is otherwise identical to θ (note: I use the same notational device for *infostates*).

The interpretation function $\|\cdot\|^{M,\theta}$ is defined as follows:

$$\begin{aligned} \|c\|^{M,\theta} &= \|c\|^M \text{ (for any constant } c) \\ \|v\|^{M,\theta} &= \theta(v) \text{ (for any variable } v) \\ \|e = e'\|^{M,\theta} &= \top, \text{ iff } \|e\|^{M,\theta} = \|e'\|^{M,\theta} \\ \|e(e')\|^{M,\theta} &= \|e\|^{M,\theta} (\|e'\|^{M,\theta}) \\ \|\lambda v.e\|^{M,\theta} &= \langle \|e\|^{M,\theta^{[v/d]}} : d \in D_T \rangle, \text{ where } v : T \\ \|l(e)\|^{M,\theta} &= \alpha, \text{ iff there exists } \beta \text{ and } \|e\|^{M,\theta} = \langle \alpha, \beta \rangle \\ \|r(e)\|^{M,\theta} &= \alpha, \text{ iff there exists } \beta \text{ and } \|e\|^{M,\theta} = \langle \beta, \alpha \rangle \end{aligned}$$

An expression $e : t$ is true iff, for every θ , $\|e\|^{M,\theta} = \top$.

Next, I will introduce DRT’s linearized box notation in a series of abbreviatory conventions (Muskens, 1996; Brasoveanu, 2006).

to drefs in the sense (of socially- or cognitively-constructed abstract objects) developed in these chapters. Rather they correspond to the *objects* they themselves ‘refer to’ (i.e. the semantics makes use of them as ‘Lagadonian’ denotations (Lewis, 1986)).

A further technical hitch is that the treatment of I^* as a total assignment means that every dref has a referent (a value at I^*). Since we want to countenance non-referential drefs (for instance, the witch-dref that Hob and Nob are thinking about), we must add a ‘dummy’ referent to $\mathcal{D}_e - \boxtimes$. Assume as a further axiom that every n -place relation constant denotes a function that maps any (well-typed) sequence of n arguments which includes \boxtimes onto \perp .

$$\forall I, w, \delta, \delta' (I(\delta) = I(\delta') \rightarrow bel(\delta, \langle I, w \rangle) = bel(\delta', \langle I, w \rangle))$$

First, *conditions* abbreviate lambda expressions of type $\mathbf{s} \rightarrow \mathbf{t}$:²¹

$$\begin{aligned} R_\omega(u_1, \dots, u_n) &:= \lambda i. R_{r(i\omega)}(l(i\omega)(iu_1), \dots, l(i\omega)(iu_n)) \\ u_1 =_\omega u_2 &:= \lambda i. l(i\omega)(iu_1) = l(i\omega)(iu_2) \\ \sim D &:= \lambda i. \neg \exists j (iDj) \end{aligned}$$

Second, *linearized boxes* abbreviate expressions of type $\mathbf{s} \rightarrow \mathbf{s} \rightarrow \mathbf{t}$:

$$\begin{aligned} [u] &:= \lambda ij. \forall \delta (\delta \neq u \rightarrow i(\delta) = j(\delta)) \wedge \forall \varpi (i(\varpi) = j(\varpi)) \\ [C] &:= \lambda ij. i = j \wedge Ci \\ Bel_\omega^\omega(u, D) &:= \lambda ij. \forall I \forall w (bel_{l(i\omega)(iu)}(i(\omega), \langle I, w \rangle) \rightarrow i^{[\omega'/\langle I, w \rangle]} D j^{[\omega'/\langle I, w \rangle]}) \\ D; D &:= \lambda ij. \exists k (iDk \wedge kDj) \\ [u_1 \dots u_n \mid C_1, \dots, C_m] &:= [u_1]; \dots; [u_n]; [C_1]; \dots; [C_m] \end{aligned}$$

$D : \mathbf{s} \rightarrow \mathbf{s} \rightarrow \mathbf{t}$ is true at an infostate i iff $\exists j (iDj)$.

D is true *simpliciter* iff D is true at i^* .

Lexicon

$$\begin{array}{ll} \text{farmer}_\omega & \rightsquigarrow \lambda \delta. [\text{farmer}_\omega(\delta)] \\ \text{donkey}_\omega & \rightsquigarrow \lambda \delta. [\text{donkey}_\omega(\delta)] \\ \text{owns}_\omega & \rightsquigarrow \lambda Q \lambda \delta. Q \lambda \delta'. [\text{owns}_\omega(\delta, \delta')] \\ \text{is}_\omega & \rightsquigarrow \lambda Q \lambda \delta. Q \lambda \delta'. [\delta =_\omega \delta'] \\ \text{thinks}_\omega^\omega & \rightsquigarrow \lambda D \lambda \delta. \text{Bel}_\omega^\omega(\delta, D) \\ \mathbf{a}^u & \rightsquigarrow \lambda F \lambda G. [u]; Fu; Gu \\ \text{every}^u & \rightsquigarrow \lambda F \lambda G. [[u]; Fu \Rightarrow Gu] \end{array}$$

Additional Rule I: No two superscripts in a discourse can match.

Additional Rule II: The index of a predicate may either be the marker ω or else the superscript of any attitude verb taking scope over the predicate. The main verb of the clause that constitutes the scope of an attitude verb Att_ω^ω must be indexed by ω' (Percus, 2000).

²¹Note that the other dynamic connectives may be defined in terms of ' \sim ' and ' $;$ ':

$$D_1 \vee D_2 := \sim ([\sim D_1]; [\sim D_2]) \quad D_1 \Rightarrow D_2 := \sim (D_1; [\sim D_2])$$

An Example

Every^u farmer_ω owns_ω a^{u'} donkey_ω

$$\begin{aligned} \rightsquigarrow & (\lambda F \lambda G. [[u]; Fu \Rightarrow Gu]) \\ & (\lambda \delta. [farmer_{\omega}(\delta)]) \\ & ((\lambda Q \lambda \delta. Q \lambda \delta' [owns_{\omega}(\delta, \delta')])((\lambda F \lambda G. [u']; Fu'; Gu')(\lambda \delta. [donkey_{\omega}(\delta)]))) \end{aligned}$$

This term reduces to:

$$[[u]; [farmer_{\omega}(u)] \Rightarrow [u']; [donkey_{\omega}(u')]; [owns_{\omega}(u, u')]]$$

For which the box below is ‘syntactic sugar’:

$$[[u \mid farmer_{\omega}(u)] \Rightarrow [u' \mid donkey_{\omega}(u'), owns_{\omega}(u, u')]]$$

Names

In Chapter 4, I gave the following analysis of names as both retrieving a familiar dref (and so qualifying as a definite expression) and introducing a new dref:

$$\text{Name}_{\omega} \rightsquigarrow \lambda F. [u \mid u =_{\omega} u_{name}]; Fu$$

Given this analysis, we can derive a mixed *de re-de dicto* reading of an attitude report via the configuration of world-binding in (41):

(41) Biron_ω thinks_ω^{ω'} Katherine_ω is_{ω'} Rosaline_{ω'}

$$\begin{aligned} \rightsquigarrow & (\lambda F. [u \mid u =_{\omega} u_{bir}]; Fu) \\ & ((\lambda D \lambda \delta. Bel_{\omega}^{\omega'}(\delta, D)) \\ & ((\lambda F. [u' \mid u' =_{\omega} u_{kat}]; Fu') \\ & ((\lambda Q \lambda \delta. Q \lambda \delta' [\delta =_{\omega'} \delta']) (\lambda F. [u'' \mid u'' =_{\omega'} u_{ros}]; Fu'')))) \end{aligned}$$

Which boils down to:

$$[u|u =_{\omega} u_{bir}]; Bel_{\omega}^{\omega'}(u, [u'|u' =_{\omega} u_{kat}]; [u''|u'' =_{\omega'} u_{ros}]; [u' =_{\omega'} u''])$$

Which is truth-conditionally (though not quite dynamically) equivalent to:

$$Bel_{\omega}^{\omega'}(u_{bir}, [u|u =_{\omega} u_{kat}, u =_{\omega'} u_{ros}])$$

Bibliography

- Abbott, B. (2002). Definiteness and proper names: Some bad news for the description theory. *Journal of Semantics*, 19:191–201.
- Algeo, J. (1973). *On Defining the Proper Name*. University of Florida Press, Gainesville.
- Aloni, M. (2000). *Quantification under Conceptual Covers*. PhD thesis, University of Amsterdam.
- Asher, N. (1986). Belief in discourse representation theory. *Journal of Philosophical Logic*, 15(2):127–89.
- Asher, N. (1993). *Reference to Abstract Objects in English*. D. Reidel, Dordrecht.
- Asher, N. and Lascarides, A. *Logics of Conversation*. Cambridge University Press, Cambridge.
- Ashworth, E. J. (Fall 2006). Medieval theories of singular terms. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*.
- Brasoveanu, A. (2006). *Structured Nominal and Modal Reference*. PhD thesis, Rutgers, The State University of New Jersey.
- Burge, T. (1973). Reference and proper names. *Journal of Philosophy*, 70(14):425–439.
- Büring, D. (1997). *The Meaning of Topic and Focus – The 59th Street Bridge Accent*. Routledge, London.
- Carroll, J. M. (1985). *What’s in a Name?* Freeman and Company, New York.
- Chastain, C. (1975). Reference and context. In Gunderson, K., editor, *Minnesota*

- Studies in the Philosophy of Science*, volume 7, pages 194–269. University of Minnesota Press, Minnesota.
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge University Press, Cambridge.
- Christoffersen, P. (1939). *The Articles: A Study of their Theory and Use in English*. E. Munksgaard, Copenhagen.
- Clark, E. and Clark, H. (1979). When nouns surface as verbs. *Language*, 55:767–811.
- Clark, H. (1992). *Arenas of Language Use*. Chicago University Press, Chicago.
- Clark, H. (1996). *Using Language*. Cambridge University Press, Cambridge.
- Clark, H. and Gerrig, R. (1983). Understanding old words with new meanings. *Journal of Verbal Learning and Verbal Behavior*, 22:591–608.
- Cresswell, M. (1990). *Entities and Indicies*. Kluwer, Dordrecht.
- Crimmins, M. (1992). *Talk about Beliefs*. MIT Press, Cambridge, MA.
- Dekker, P. and van Rooy, R. (1998). Intentional identity and information exchange. In Cooper, R. and Gamkrelidze, T., editors, *Proceedings of the Second Tbilisi Symposium on Language, Logic and Computation*. Tbilisi State University.
- DeVault, D. and Stone, M. (2006). Scorekeeping in an uncertain language game. In Schlangen, D. and Fernandez, R., editors, *BRANDIAL 06: Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, pages 139–146.
- DeVault, D. and Stone, M. (2007). Managing ambiguities across utterances in dialogue. In Artstein, R. and Vieu, L., editors, *DECALOG: The 2007 Workshop on the Semantics and Pragmatics of Dialogue*.
- Devitt, M. (1974). Singular terms. *The Journal of Philosophy*, 71(7):183–205.
- Donnellan, K. (1966). Reference and definite descriptions. *Philosophical Review*, 75(3):231–304.
- Donnellan, K. (1972). Proper names and identifying descriptions. In Davidson, D. and Harman, G., editors, *Semantics of Natural Language*, pages 356–379. D. Reidel, Dordrecht.

- Edelberg, W. (1986). A new puzzle about intentional identity. *Journal of Philosophical Logic*, 15(1):1–25.
- Edelberg, W. (1992). Intentional identity and the attitudes. *Linguistics and Philosophy*, 15(6):561–596.
- Elbourne, P. (2005). *Situations and Individuals*. MIT Press, Cambridge, MA.
- Evans, G. (1982). *Varieties of Reference*. Clarendon Press, Oxford.
- Fermor, P. (1966). *Roumeli*. Parker & Row, New York.
- Fine, K. (2003–2004). John Locke Lectures. Oxford University.
- Fine, K. (2007). *Semantic Relationism*. Blackwell Publishing, Oxford.
- Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28:97–115.
- Fodor, J. A. (1975). *The Language of Thought*. Thomas A. Crowell Co., New York.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press, Cambridge, MA.
- Forbes, G. (1996). Substitutivity and the coherence of quantifying in. *Philosophical Review*, 105(3):337–372.
- Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, pages 25–50. English Translation: *On Sense and Meaning*, in Brian McGuinness (ed.), *Frege: Collected Works*, pages 157–177. Basil Blackwell, Oxford.
- Frege, G. (1953). *The Foundations of Arithmetic*. Philosophical Library, New York. Originally published 1884.
- Geach, P. (1967). Intentional identity. *Journal of Philosophy*, 74(20).
- Geurts, B. (1997). Good news about the description theory of names. *Journal of Semantics*, 14:319–348.
- Grice, P. (1957). Meaning. *The Philosophical Review*, 66:377–88.
- Grice, P. (1971). Intention and uncertainty. *Proceedings of the British Academy*, 57:263–279.

- Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100.
- Gundel, J., Hedberg, N., and Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 69(2):274–307.
- Hahn, E. A. (1969). *Naming Constructions in Some Indo-European Languages*. Case Western Reserve University Press, Cleveland, OH.
- Hamm, F., Kamp, H., and van Lambalgen, M. (2006). There is no opposition between formal and cognitive semantics. *Theoretical Linguistics*, 32.
- Heim, I. (1982). *On the Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts at Amherst.
- Hintikka, J. (1969). Semantics for propositional attitudes. In *Models for Modalities*, pages 87–112. D. Reidel, Dordrecht.
- Hobbs, J., Stickel, M., Martin, P., and Edwards, D. (1993). Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142.
- Husserl, E. (1900–1901). *Logische Untersuchungen*. Halle/Saale.
- Jackendoff, R. (1994). *Patterns in the Mind: Language and Human Nature*. Basic Books.
- Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press, Oxford.
- Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language, Part 1*, volume 135, pages 277–322. Mathematical Centre Tracts, Amsterdam.
- Kamp, H. (1990). Prolegomena to a structural account of belief and other attitudes. In Anderson, C. A. and Owens, J., editors, *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*. Center for the Study of Language and Information, Stanford, California.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.
- Kamp, H., van Genabith, J., and Reyle, U. (1999). Discourse representation theory.

- Draft of an article to appear in the *Handbook of Philosophical Logic*.
- Kaplan, D. Quantifying in. In *Words and Objections*. D. Reidel, Dordrecht.
- Kaplan, D. (1989). Demonstratives. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 481–564. Oxford University Press, Oxford.
- Karttunen, L. (1976). Discourse referents. In McCawley, J., editor, *Syntax and Semantics: Notes from the Linguistic Underground*, volume 7, pages 363–386. Academic Press, New York.
- Kripke, S. (1972). Naming and necessity. In Davidson, D. and Harman, G., editors, *Semantics of Natural Language*, pages 253–355. D. Reidel, Dordrecht.
- Kripke, S. (1977). Speaker's reference and semantic reference. In French, P., Uehling, T., and Wettstein, H., editors, *Midwest Studies in Philosophy vol. II: Studies in the Philosophy of Language*, pages 255–276. University of Minnesota Press, Morris, MN.
- Kripke, S. (1979). A puzzle about belief. In Margalit, A., editor, *Meaning and Use: Papers Presented at the Second Jerusalem Philosophy Encounter*, pages 239–88. D. Reidel, Dordrecht.
- Lambrecht, K. (1994). *Information Structure and Sentence Form: A Theory of Topic, Focus, and the Mental Representations of Discourse Referents*. Cambridge University Press, Cambridge.
- Larson, R. and Segal, G. (1995). *Knowledge of Meaning*. MIT Press, Cambridge, MA.
- Laurence, S. (2003). Is linguistics a branch of psychology? In Barber, A., editor, *Epistemology of Language*. Oxford University Press, Oxford.
- Lepore, E. and Stone, M. (2007). Logic and semantic analysis. In Jacquette, D., editor, *Handbook of the Philosophy of Logic*. Elsevier.
- Lewis, D. (1969). *Convention*. Harvard University Press, Cambridge, MA.
- Lewis, D. (1972). General semantics. In Davidson, D. and Harman, G., editors, *Semantics of Natural Language*. D. Reidel, Dordrecht.

- Lewis, D. (1979a). Attitudes *de dicto* and *de se*. *Philosophical Review*, 88:513–43.
- Lewis, D. (1979b). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell, Oxford.
- Loar, B. (1976). The semantics of singular terms. *Philosophical Studies*, 30:353–377.
- Maier, E. (2006). *Belief in Context: Towards a unified semantics of de re and de se attitude reports*. PhD thesis, Radboud University, Nijmegen.
- Mani, I., Macmillan, T. R., Luperfoy, S., Lusher, E. P., and Laskowski, S. J. (1993). Identifying unknown proper names in newswire text. In *Proceedings of the Workshop on Acquisition of Lexical Knowledge from Text*, pages 44–54, Columbus, OH.
- Mates, B. (1952). Synonymity. In Linsky, L., editor, *Semantics and the Philosophy of Language*, pages 111–36. University of Illinois Press, Urbana, Illinois.
- Matushansky, O. (2006a). Call me an ambulance. In Bateman, L. and Ussery, C., editors, *Proceedings of NELS*, volume 35, pages 419–434.
- Matushansky, O. (2006b). Why Rose is the Rose. In Bonami, O. and Hofherr, P. C., editors, *Empirical Issues in Formal Syntax and Semantics*, volume 6, pages 285–308.
- McDonald, J. (1996). Internal and external evidence in the identification and semantic categorization of proper names. In Boguraev, B. and Pustejovsky, J., editors, *Corpus Processing for Lexical Acquisition*, pages 21–39. MIT Press, Cambridge, MA.
- Miklosich, F. (1874). *Vergleichende Grammatik der Slavischen Sprache IV*. Vienna.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In et. al., H., editor, *Approaches to Natural Language*, pages 221–242. D. Reidel, Dordrecht.
- Muskens, R. (1996). Combining Montague semantics and discourse representation. *Linguistics and Philosophy*, 19:143–186.
- Neale, S. (1990). *Descriptions*. MIT Press, Cambridge, MA.

- Neale, S. (2005). Pragmatism and binding. In Szabó, Z. G., editor, *Semantics versus Pragmatics*, pages 165–285. Oxford University Press, Oxford.
- Payne, J. and Huddleston, R. (2002). Nouns and noun phrases. In Pullum, G. K. and Huddleston, R., editors, *The Cambridge Grammar of the English Language*. Cambridge University Press, Cambridge.
- Percus, O. (2000). Constraints on some other variables in syntax. *Natural Language Semantics*, 8:173–229.
- Prince, E. (1979). On the given/new distinction. In Hanks, W., Hofbauer, C., and Clyne, P., editors, *Papers from the Fifteenth Regional Meeting, Chicago Linguistic Society*, pages 267–278.
- Putnam, H. (1954). Synonymity, and the analysis of belief sentences. *Analysis*, 41:114–122.
- Putnam, H. (1975). The meaning of ‘meaning’. In Gunderson, K., editor, *Language, Mind, and Knowledge. Minnesota Studies in the Philosophy of Science, Vol. 7*, pages 131–93. University of Minnesota Press, Minneapolis, Minnesota.
- Quine, W. V. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy*, 53:101–111.
- Quine, W. V. (1960). *Word and Object*. MIT Press, Cambridge, MA.
- Ramsey, F. (1925). Universals. *Mind*, 34:401–417.
- Richard, M. (1990). *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*. Cambridge University Press, Cambridge.
- Russell, B. (1905). On Denoting. *Mind*, 14:479–493.
- Salmon, N. (1986). *Frege’s Puzzle*. MIT Press, Cambridge, MA.
- Sanford, A. J. and Lockhart, F. (1990). Description types and method of conjoining as factors influencing plural anaphora: a continuation study of focus. *Journal of Semantics*, 7(4):365–378.
- Soames, S. (1988). Substitutivity. In Thomson, J. J., editor, *On Being and Saying: Essays in Honor of Richard Cartwright*, pages 99–132. MIT Press, Cambridge, MA.

- Soames, S. (2002). *Beyond Rigidity*. Oxford University Press, Oxford.
- Stalnaker, R. (1979). Assertion. In Cole, P., editor, *Syntax and Semantics*, volume 9: Pragmatics, pages 315–32. Academic Press, New York. Reprinted in Stalnaker 1999, 78–95.
- Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language, and Information*, 7(1):3–19.
- Stalnaker, R. (1999). *Context and Content*. Oxford University Press, Oxford.
- Stanley, J. (2002). Making it articulated. *Mind & Language*, 17(1&2):149–168.
- Steedman, M. and Stone, M. (2006). Is semantics computational? *Theoretical Linguistics*, 32(1):73–89.
- Stone, M. (2003). Linguistic representation and Gricean inference. In *IWCS-5: Proceedings of the Fifth International Workshop on Computational Semantics*, pages 5–21.
- Stone, M. (2004a). Communicative intentions and conversational processes in human-human and human-computer dialogue. In Trueswell and Tanenhaus, editors, *Approaches to Studying World-Situated Language Use*, pages 39–70. MIT Press, Cambridge, MA.
- Stone, M. (2004b). Intention, interpretation and the computational structure of language. *Cognitive Science*, 28(5):781–809.
- Tarski, A. (1956). The concept of truth in formalized languages. In Woodger, J., editor, *Logic, Semantics, Metamathematics*, pages 152–78. Clarendon Press, Oxford.
- Taschek, W. (1998). On ascribing beliefs: Content in context. *Journal of Philosophy*, 95:323–353.
- Thomason, R. H., Stone, M., and DeVault, D. (2007). Enlightened update: A computational architecture for presupposition and other pragmatic phenomena. In Byron, D., Roberts, C., and Schwenter, S., editors, *Presupposition Accommodation*.
- Van Langendonck, W. (1998). *Theory and Typology of Proper Names*. Mouton de Gruyter, Berlin and New York.

van Rooy, R. (2000). Anaphoric relations across attitude contexts. In von Heusinger, K. and Egli, U., editors, *Reference and Anaphoric Relations*, pages 157–181. Kluwer, Dordrecht.

Williamson, T. (1994). *Vagueness*. Routledge, London and New York.

Curriculum Vitae

Samuel J. Cumming

Education

- 1995–1998 BA (Hons.) in English and Philosophy,
University of Melbourne
- 1999–2001 MA in Philosophy, University of Melbourne
- 2002–2007 PhD in Philosophy (with Certificate in Cognitive Science),
Rutgers University

Positions

- 2002–2005 Graduate Fellow, Rutgers University
- 2005–2006 Teaching Assistant, Rutgers University
- 2006–2007 University and Bevier Fellow, Rutgers University

Publications

- 2003 Two Accounts of Indexicals in Mixed Quotation.
Belgian Journal of Linguistics, 17: 77–88.