REFERENCE AND RELEVANCE

by

PIERRE BAUMANN

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, the City University of New York.

2008

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date	Barbara Montero Chair of Examining Committee
Date	Iakovos Vasiliou Executive Officer

Stephen Neale

Robert Fiengo

Michael Levin

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

REFERENCE AND RELEVANCE

by

Pierre Baumann

Adviser: Stephen Neale

Implicit in much of contemporary philosophy of language is a certain picture of linguistic communication, which may be called the "code model." Linguistic communication on the code model works like mailing a package: the sender (the speaker) places the item (the thought) in a box (the sentence) and the recipient (the hearer) takes out the very item from the box (understands the thought encoded by the sentence). Thus, if the hearer masters the language (knows the code), she should have no problem identifying the thought or proposition the speaker wants to express by her utterance. Non-literality, indexicality, and ambiguity complicate this simple picture, but in general it is assumed that sentence meaning and syntax suffice to determine the proposition expressed by a literal utterance.

Hand-in-hand with the code model of linguistic communication goes the idea that certain expressions are inherently referential: they pick out or represent

entities in the world as a matter of their linguistic meaning and the syntactic category to which they belong. That is, these "referring expressions" refer to things as a matter of their semantics. So, a hearer familiar with the referring expressions of her language should be able to determine the object referred to by a speaker's use of a referring expression (the expression's truth-conditional value) simply on the basis of her linguistic competence, on her knowledge of the code.

Various authors, notably Dan Sperber and Deirdre Wilson, and following them, Robyn Carston, François Récanati, Anne Bezuidenhout, and Stephen Neale, have argued in recent years that the code model is deeply flawed. In broad agreement with them, this dissertation criticizes one important element of the code model, the notion of semantic reference. The main thesis it defends is that semantic reference is a dubious notion and that reference is best understood pragmatically. The dissertation does not attempt to *prove* this thesis, since to do so it would have to consider every single type of noun phrase (NP)—too ambitious a task. Instead, the thesis is supported through two case studies. The first case study, developed in Chapters 2-4, concerns a type of NP that is standardly held to have a referential semantics: proper names. The second

case study, developed in Chapter 5, concerns a type of NP that is standardly held to have a non-referential semantics: NPs of the syntactic form “every F.” In both cases, it is shown that:

(a) The expression in question may be literally interpreted referentially or non-referentially;

(b) Depending on whether it is interpreted referentially or non-referentially, the expression may have a truth-conditional value that is not the one predicted by the standard semantic analysis of the expression;

(c) The type of truth-conditional value the NP may have on a given occasion of use is pragmatically determined in accordance with Relevance-Theoretic principles of interpretation.

These results parallel recent work on demonstratives suggesting that they admit of literal referential and non-referential interpretations. The basic view on reference expressed in this dissertation agrees with the spirit of Peter Strawson’s famous dictum that “referring is not something an expression does; it is something that someone can use an expression to do.” The dissertation attempts to validate the letter of the remark by exploiting the resources of a particular pragmatic framework, Sperber and Wilson’s Relevance Theory.

Acknowledgements

I would like to thank Stephen Neale for his generosity, open-mindedness, and support as adviser of this dissertation. It was a privilege to work with him. Thanks also go to Michael Levin, Michael Devitt, Robert Fiengo, and Barbara Montero, for their valuable feedback, and to John Greenwood and Rosemarie Iannuzzi, for their friendly advice and help with the administrative side of things. I am grateful to my parents, Daniel and Myrna Baumann, and sister, Anne-Monique Baumann Ransdell, for their untiring moral support, and to my wife, Heather, whose love and faith sustained me through this process.

Table of Contents

Chapter 1

Background

1. Introduction

This dissertation is concerned with the notion of semantic reference. The main thesis it defends is that semantic reference is a dubious notion and that reference is best understood pragmatically. The dissertation does not attempt to *prove* this thesis, since to do so it would have to consider every single type of noun phrase (NP)—too ambitious a task. Instead, the thesis is supported through two case studies. The first case study concerns a type of NP that is standardly held to have a referential semantics: proper names. The second case study concerns a type of NP that is standardly held to have a non-referential semantics: NPs of the syntactic form "every F." In both cases, it is shown that:

(a) The expression in question may be literally interpreted[1] referentially or non-referentially;

(b) Depending on whether it is interpreted referentially or non-referentially, the expression may have a truth-

[1] Here, and throughout the dissertation, by "a literal interpretation of an expression" is meant an interpretation that is (i) sanctioned by the expression's linguistic meaning and (ii) not derived through Gricean principles of communication, i.e. an interpretation according to which the information conveyed by the expression is not part of an implicature. The notion of linguistic meaning is explained in §3. Grice's notion of an implicature and his principles of communication are discussed in §4-6.

conditional value[2] that is not the one predicted by the standard semantic analysis of the expression;

(c) The type of truth-conditional value the NP may have on a given occasion of use is pragmatically determined in accordance with Relevance-Theoretic principles of interpretation.[3]

The basic view on reference expressed in this dissertation agrees with the spirit of Strawson's famous dictum that "referring is not something an expression does; it is something that someone can use an expression to do." (1950, p. 326) The dissertation attempts to validate the letter of the remark by exploiting the resources of a particular pragmatic framework, Sperber and Wilson's (1986/1995) Relevance Theory.

The aim of this first chapter is to explain certain notions and assumptions that are presupposed in later chapters.

2. Reference

Reference is a property that has been attributed to, among other things, mental states, uses of linguistic expressions, and the expressions themselves. The target of

[2] The notion of a truth-conditional value will be explained in §§2-3.
[3] Relevance Theory is discussed in §6.

this dissertation is reference viewed as a property of words, or "semantic" reference, as it is customarily called. We shall also use the term "referential" to characterize a type of use of words—or more precisely, a hearer's interpretation of a speaker's use of words; we will return to referential uses/interpretations in §4 below. The relevant understanding of semantic reference is described further in what follows.

Semantic reference is taken to be an inherent feature of certain expressions: namely, "referring" expressions. Normally included in this class are proper names ("Aristotle," "New York"), demonstratives ("this," "that"), and pronouns ("she," "I"). Other expressions, such as definite descriptions ("the table"), indefinite descriptions ("a chair"), and certain nouns and adjectives—so-called natural kind terms ("water," "tiger," "red")—are sometimes included in this category as well. The main function of a referring expression is supposed to be "to represent," "to stand for," "to designate," or "to pick out" an object in the world.[4] The name "John Smith," for

[4] Or several objects, if the term is plural (e.g. "those"). The topic of plural reference is comparatively underdeveloped in the philosophy of language; a couple of influential formal accounts are those of Link (1983) and Boolos (1984). See also (Cameron, 1999). From our point of view, it is of course no problem to acknowledge that people can and do refer to things using plural expressions just as they are able to do so via singular ones.

example, is supposed to refer to an individual so named, some John Smith.

For some philosophers,[5] a central question concerning reference has been: How do words come to refer? They have sought to identify the mechanism whereby an expression diachronically comes to represent an object. In this dissertation, we will not be concerned with this question, nor with the various existing proposals on how to answer it. If our main thesis is correct, this question is simply misplaced: words never refer to things on their own; people, using words, do.

In any case, the notion of semantic reference to be criticized here is a synchronic one: it is a property that certain linguistic expressions are supposed to possess, where "linguistic" refers to language viewed abstractly and statically. The notion of a "referring expression" that interests us here, the notion we aim to challenge, may be defined as follows:

Referring Expression

An expression is a referring expression if (according to the standard semantic view) it has an object as truth-conditional value.

[5] E.g., Evans (1973), Devitt (1981), Reimer (2003).

The key term in this definition is "truth-conditional value." An expression's truth-conditional value is the contribution the expression is supposed to make to the truth-conditional content of a sentence that contains it.[6, 7] According to standard truth-conditional semantic theory, as presented, for example, in (Lewis, 1970/1983), (Chierchia & McConnell-Ginet, 1990), and (Heim & Kratzer, 1998), different types of expression have different types of truth-conditional value. Referring expressions have objects as truth-conditional values; adjectives, nouns, and intransitive verbs properties (or sets); transitive verbs relations, etc. An expression's having a certain type of truth-conditional value is held to be a function of the expression's syntactic type and its linguistic meaning.

The assumption, in other words, is that the truth-conditional value of an expression is mechanically turned out by language. Presumably, then, according to this picture, a hearer who masters the language would be able to identify the correct truth-conditional value of a referring expression used by a speaker simply on the basis of her

[6] The term "semantic value" is also often used for this contribution. This term will not be used in the dissertation, for reasons that will become apparent in §4.

[7] Throughout the dissertation, by "sentence" is meant "declarative sentence in the indicative mood."

knowledge of the expression's meaning and syntactic type, that is, on the basis of her linguistic competence.

Along with the standard semantic notion of reference, this dissertation rejects this assumption as to how the truth-conditional value of an expression is decided on an occasion of use. (The reasons for this rejection are given in §§4-6.) However, we do want to retain the notion of a truth-conditional value itself. This notion will be interpreted pragmatically rather than semantically, in the way outlined in §4 below. The next section characterizes the notion of linguistic meaning, its relationship to the notion of a truth-conditional value, and the concept of a sentence's truth-conditional content. Once again, it should be kept in mind that "truth-conditional value" will be given a different, pragmatic interpretation later.

3. Two Distinctions: Expression Meaning/Truth-conditional Value, Sentence Meaning/Truth-conditional Content

An expression's truth-conditional value must be distinguished from its linguistic meaning.[8] The linguistic meaning of an expression is the context-invariant information encoded by the expression. Linguistic meanings are the sorts of meanings that lexicographers compile and

[8] The distinction being made here corresponds to Kaplan's (1989) character/content distinction.

that any competent speaker of the language knows; they are based on the ordinary usage of an expression within a linguistic community. In this dissertation, by "meaning" we shall always mean "linguistic meaning." By "to encode" we shall mean "to have as part of its linguistic meaning."[9] The following example illustrates the difference between an expression's truth-conditional value and its meaning.

Suppose you enter a classroom and find (1) written on the blackboard.

(1) I am thirsty.

As a competent speaker of English, you have no problem understanding this sentence. You know the meaning of each of its component expressions, including the meaning of the pronoun "I," viz., "the person currently speaking or writing." However, in the present circumstances you are unable to say whether (1) is true or false. In particular, you cannot fully specify the conditions that would make it true, since you don't know who "I" is being used to refer to. This individual, whoever he or she is, represents the truth-conditional value of "I": it is a condition for the truth of (1) (taken seriously and not as some sort of joke) that the individual who wrote the sentence be thirsty.

[9] The insertion of "part" is intended to register that affixes such as "-ed" in "walked" also have meaning, which they contribute to the expression as a whole.

Thus, the context-invariant linguistic meaning of "I" ("the person currently speaking or writing") cannot be identified with its truth-conditional value on an occasion of use, since a competent speaker will understand the meaning of the word even if she doesn't know to whom a particular utterance or inscription of it refers, and since the word may be used to refer to different individuals in different contexts.

Just as an expression's meaning must be distinguished from its truth-conditional value on an occasion of use, so must a whole sentence's meaning be distinguished from the proposition it expresses in a context, its truth-conditional content.[10] Consider (1) again. As just noted, (1) has a meaning graspable by anyone sufficiently proficient in English. Considered in the abstract of a context, however, the sentence fails to express anything definite. In particular, we don't know under which exact circumstances an utterance of (1) would be true or false, since we are missing a key element: the individual who represents the truth-conditional value of "I."[11] In contrast, when someone, say A, utters (1), she expresses

[10] The terms "truth-conditional content," "proposition" and "thought" will be used interchangeably throughout the dissertation, as will "object" and "individual."
[11] In addition, "is" and "thirsty" would need to be disambiguated.

the determinate proposition that A is thirsty; and what she expresses is true just in case A *is* thirsty.

Another reason why a sentence's meaning and its truth-conditional content must be distinguished is because the relation between sentences and propositions is many-to-many. One and the same sentence may be used to express multiple propositions: e.g. (1) uttered by B expresses the proposition that B is thirsty. And conversely, one and the same proposition may be expressed via different sentences of one or several languages: e.g. if B utters the sentence "A is thirsty" or the Spanish sentence "*A tiene sed*" he expresses the same proposition A did in uttering (1).

At this point two qualifications need to be made regarding propositions. The first is that throughout the dissertation propositions will be construed in a neutral manner, without a commitment to a particular view on their nature or their theoretical eliminability.[12] Propositions, for us, simply represent a convenient and familiar way of putting things. The second point has been presupposed by the remarks of the preceding paragraphs, but it is important to make it explicit. We will be assuming, with Strawson (1950, 1952) and a number of prominent

[12] The same disclaimer applies to truth-conditional values. For instance, the terms "object" and "property" will be used without presupposing a particular view on the nature of objects and properties.

contemporary authors (including, e.g., Searle [1980], Sperber and Wilson [1986/1995], Perry [2001], Neale [1990], Bach [1987/1994a], Récanati [1993]), and Carston [2002]) that sentences, qua abstract types, do not express propositions or have determinate truth conditions, but that only sentence tokens used in actual contexts do.[13] We shall use the term "utterance" to refer to a sentence-token-used-in-a-context,[14] and say that utterances (and not sentences) are true or false, have truth conditions, and express propositions.[15]

With the expression meaning/truth-conditional value and sentence meaning/truth-conditional content distinctions in place, it is now possible to state the problem of this dissertation in a precise fashion. The problem is that,

---

[13] Here we are alluding to the type/token distinction that is commonly made in philosophy and linguistics. The distinction is between expressions considered as abstract entities (types) or as concrete instances (tokens). For example, the sentence "The book is on the table" contains two instances or tokens of the word "the;" these are tokens of one English word, "the," considered abstractly as a type. The type/token distinction applies not just to words, but also to more complex expressions, such as sentences. The dissertation is neutral with respect to the metaphysics of types and tokens.

[14] By "used" is meant "spoken:" all examples to follow involve imaginary face-to-face conversations.

[15] This is the official position of this dissertation. For expository purposes, however, the phrasing will frequently be used (particularly in the next chapter) that a sentence expresses a proposition, or that a sentence has truth conditions, or that a sentence is true or false. Similarly, we will often say that an expression "has a truth-conditional value." This, it should be stressed here, is really just shorthand for "the truth-conditional value the interpreter identified the expression as having in the context it was used in," a form of words too long and cumbersome to be used repeatedly. The official position concerning sentences and expressions shall be recalled whenever it is important to the discussion.

for our two target expressions at least, the meanings of NPs appear to sanction truth-conditional values other than those specified by the standard semantics of these expressions. As will be argued in Chapters 2-4, proper names may be literally interpreted as having not just objects as truth-conditional values, but also properties (or sets); in Chapter 5 the case will be made that NPs of the form “every F” may have as truth-conditional value not just a complex consisting of a property and the subset relation, but also objects. Utterances containing these expressions may thus have truth conditions different from those predicted by the standard semantic account.

To show that this problem exists—i.e. to demonstrate that these two types of NP may indeed have these non-standard truth-conditional values—is the principal task to which the dissertation is dedicated. But in addition, it will be argued that the existence of this problem casts doubt upon the mainstream assumption that an expression’s meaning and syntactic type determine its truth-conditional value, and in particular, that reference is a semantic matter. It shall be proposed instead that reference is best understood pragmatically, and in articulating this proposal a particular pragmatic framework, Relevance Theory, will be adopted. Before introducing this framework, however, it is

necessary to say a word about how the terms “semantic” and “pragmatic” should be interpreted.

4. “Semantic” and “Pragmatic”

The foregoing discussion raises the question: What is the proper object of semantics? For instance, is it linguistic meaning or truth-conditional value? The tradition initiated by Frege and Russell and developed by Tarski, Carnap, Montague, Lewis, Kripke, Kaplan, Davidson, and many others takes semantics to be centrally concerned with truth conditions and truth-evaluability. “Semantics without a treatment of truth conditions is not semantics,” as Lewis (1970/1983, p. 190) put it. According to this paradigm, truth-conditional value and truth-conditional content are the ends of semantic analysis, with expression meaning and sentence meaning representing intermediate steps in this analysis.

In contrast to the Strawsonian stance we endorsed above, according to which utterances and not sentences express truth-evaluable propositions, this semantic tradition generally assumes that sentences may be ascribed truth-conditional contents independently of a context of utterance and the speech acts the sentences may be used perform. The exceptions to this rule are sentences that

contain ambiguous or context-sensitive expressions, in which case it is admitted that contextual information is to be taken into account to disambiguate between different meanings and to assign truth-conditional values to the context-sensitive expressions. Ambiguity is a common, everyday affair: e.g. "bank" may mean "financial institution" or "side of a river." Expressions traditionally recognized as context-sensitive include pronouns, demonstratives, tensed verbs, and certain temporal and spatial expressions, such as "today" and "here." Collectively such expressions are denominated "indexicals." For all indexicals, the contextual contribution of the expression is held to be controlled by its linguistic meaning, in the manner illustrated by our example regarding "I." Not just any feature of the extra-linguistic context may be the truth-conditional value of "I:" solely the individual who is doing the speaking or writing.[16]

Aside from disambiguation and reference assignment to indexicals, and ignoring non-literality, the semantic approach under consideration supposes that meaning and syntax suffice to encapsulate a determinate and context-

[16] Or, as Neale points out (conversation), the individual *perceived* to be doing the speaking and writing: e.g. "me" in "Vote for me!", written on a political poster on a wall, would be taken to refer to the candidate promoted in the poster.

invariant truth-conditional content for every well-formed sentence of a natural language. Heim and Kratzer's slogan, "to know the meaning of a sentence is to know its truth conditions," (1998, p. 1) succinctly captures the mainstream view on knowledge of such truth-conditional content. This meaning-and-syntax-determined proposition is "what is said" when a sentence is used in a concrete communicative situation.

The term "what is said" alludes to a particular conception of pragmatics, to which semantics as we've just outlined it is allied. This conception originates in the work of Grice. In (1967/1989), Grice introduced a distinction between what is said by a speaker who assertively utters a sentence and what is "implicated" by him in using the sentence. What is said by a speaker, according to Grice, is "closely related to the conventional meaning of the words (the sentence) he has uttered" (p. 25); it is determined "in virtue of the particular meanings of the elements of [the sentence], their order, and their syntactic character." (p. 87) As currently used, the term "what is said" refers to the proposition, thought, or truth-conditional content expressed by a sentence uttered in a conversation.

Grice was interested in giving a precise account of what it is for a speaker to *say* one thing and *mean* another, since it is an obvious fact of ordinary communication that speakers intend to convey, and mean for hearers to infer, all sorts of information that does not strictly or logically follow from the meaning of the words that they use. That a speaker will use a sentence seemingly expressing a particular proposition to imply a different proposition (or several different propositions) without logically entailing the additional proposition(s) is certainly a familiar enough phenomenon; otherwise irony and sarcasm would be inexplicable, for example. From a logical point of view, the sentence "That is a brilliant answer" does not entail its negation; yet in a given conversational situation the opposite of its meaning is precisely what is implied. Grice was the first to systematically study this phenomenon, to which he gave the technical name "implicature."

The dominant view has been that Grice's distinction between what is said and what is implicated demarcates two separate spheres of linguistic study, semantics and pragmatics. What is said, on the view derived from (1967/1989) is a semantic notion—indeed, it is the *prime* semantic notion—since it is the truth-conditional,

propositional content expressed by the uttered sentence. Pragmatics then studies what is implicated. More generally, pragmatics is supposed to be concerned with speakers' communicative intentions or "speaker meaning."

This dissertation does not disagree with this last supposition about the main task of pragmatics: it accepts that pragmatics has to do with speakers' intentions and uses of expressions that reveal those intentions. Here we will call "pragmatic" any explanation that appeals to speakers' intentions and uses of expressions. So, for example, an explanation that appeals to a speaker's intention to refer to something, or to her referential use of some expression to make her intention known to her audience, is a pragmatic one.[17]

In this connection it should be noted that a speaker may use an expression to refer to an object even though the expression is not standardly classified as a referring expression. This would be the case, for instance, if

[17] The explanations offered in this dissertation will always be from the point of view of the audience, not the speaker. This is a methodological constraint imposed by Relevance Theory, which describes communication from the point of view of the audience. Thus the speaker's intentions and uses of expressions are the intentions and uses ascribed to her by the audience. This constraint is without consequence, however, since the examples we will be discussing are highly idealized ones, in which the audience always correctly recognizes the speaker's intentions, and so is able to understand or interpret exactly what the speaker means. It is stipulated that the audience's interpretation will always match the speaker's intended meaning.

someone used the definite description “the table” to say something about a particular table. The dominant view on the semantics of definite descriptions is the Russellian (1905) one according to which definite descriptions are not referring but “denoting” or quantificational expressions.

While aligning to this extent with the traditional understanding of pragmatics, this dissertation nevertheless disagrees with the Gricean view that uses, intentions, and other broadly contextual factors such as unstated facts and assumptions about the knowledge, beliefs, expectations, and personal characteristics and histories, etc., of the people involved in a communicative situation, are relevant only for inferring implicatures. Our position is that such factors may affect an utterance’s truth-conditional content as well; they enter into what is said. This position, which is sometimes labeled “pragmatism,” “contextualism,” or “truth-conditional pragmatics,” is shared by a number of influential theorists, including, e.g., Searle (1979, 1980), Travis (1981), Anscombre and Ducrot (1983), Ducrot (1984), Sperber and Wilson (1986/1995), Carston (1988, 2002), Neale (1990, 2001, 2004a, 2004b, 2007, 2008 forthcoming), Récanati (1993, 2004), Bezuidenhout (1997a,

1997b, 2002), Moravcsik (1998), and Perry (2001).[18] Our claim that reference (in the case of names and "every F") is pragmatically determined is grounded on the pragmatist idea that linguistically unmarked contextual factors such as speakers' intentions and assumptions about the audience's beliefs come into play in determining[19] the

[18] Early hints of this view may be found in (Quine, 1940, pp. 146-7) and in Sellars' (1954) discussion of "ellipsis" (a notion that does not seem equivalent to the contemporary syntactic notion of ellipsis, but rather more akin to semantic or truth-conditional underdetermination, a notion to be discussed in §5 below).

[19] As Neale (2004b) and Bach (2006) have emphasized, there are two senses of the term "to determine" that must be distinguished but are often not in contemporary discussions on the semantics/pragmatics distinction. The first is an epistemological sense (= to identify, figure out, ascertain, decide). The second is constitutive or metaphysical (= to constitute, compose). In this dissertation, "to determine" is used in the epistemological sense, not the constitutive sense. Although it is important to distinguish the two senses, it should be noted that the challenge of pragmatism can be formulated in either sense:

1) Epistemological "to determine:" if sentence meaning (plus disambiguation and reference assignment to indexicals) does not provide sufficient evidence on which to identify the truth-conditional content of an utterance, how then does the hearer identify it?

2) Metaphysical "to determine:" if sentence meaning (plus values yielded by disambiguation and reference assignment) does not provide sufficient material to constitute the proposition recovered by the interlocutors in a context, what extra (non-compositional) material goes into it, and how does it combine, exactly, with the semantic and syntactic information thrown up by the sentence?

As mentioned earlier, our aim is to stay away as much as possible from metaphysical issues having to do with the nature and structure of propositions, and in this case also from questions having to do with language processing and the architecture of the mind. Our discussion can be carried out without taking a stand on these issues. So, though recognizing the important challenge posed by 2), when we speak of pragmatism we will mean 1), and "to determine" will be used in the epistemological sense only.

Having said this, it should be stressed that, as with all terms used in linguistic theory that have epistemological connotations (e.g. "to know," "to understand," "to produce"), "to determine" should not be

proposition literally expressed, and not just in inferring an implicature.

From a Gricean point of view—and from the point of view of truth-conditional semantics as well—to speak of an expression's truth-conditional value as pragmatically determined is incoherent, since the notion of a truth-conditional value is a technical one belonging to truth-conditional semantics and not pragmatics. Truth-conditional values, to repeat, are viewed as decided by meaning and syntax,[20] not by contextual factors (with the exceptions of ambiguous and recognized context-sensitive expressions, which, however, possess linguistic meanings that mandate specific contextual contributions). It is true that the notion of a truth-conditional value is stipulated to be a semantic notion. Nevertheless, the notion can be detached from truth-conditional semantics as traditionally conceived. For the idea that communicators interpret expressions as making specific contributions to the truth conditions of whole utterances is intuitively plausible on its own, irrespective of which discipline

---

understood as implying that the hearer is consciously and occurrently trying to figure out an expression's truth-conditional value as he hears an utterance. Utterance understanding is a fast and mostly unconscious process according to Relevance Theory. More on this in §6 below.

[20] More on this mechanical "code model" of understanding and communication in §6 below.

should be assigned to study it.[21] To characterize this idea as semantic or pragmatic is an extra step. So, in order to get our discussion about the truth-conditional values of names and quantifiers off the ground, a certain leeway should be granted us to interpret "truth-conditional value" as separate from a particular conception of semantics.

The contrast between the mainstream and our preferred pragmatist understanding of the semantics/pragmatics distinction may be schematized as follows:

Mainstream View

Semantics	*Pragmatics*
Linguistic Meaning	Implicature
Truth-conditional Value	
Truth-conditional Content	
(What is Said)	

[21] That the identification of truth-conditional content is a task for pragmatics and not semantics is implied by the term "truth-conditional pragmatics," employed by some authors (e.g Récanati [1993]; Bezuidenhout [2002]) to describe their theoretical enterprise.

Pragmatist View

Semantics	*Pragmatics*
Linguistic Meaning	Truth-conditional Value
	Truth-conditional Content (What is Said)
	Implicature

A defender of the mainstream view on the division of labor between semantics and pragmatics might react in two ways to the arguments concerning names and the NP “every F” that will be developed in subsequent chapters.

First, she might counter that the standard semantic account of these two expressions may be expanded to include the non-standard truth-conditional values we argue these expressions may have. For example, it might be proposed that the two expressions are ambiguous. Each name and each NP of the form “every F” (considered as types in both cases) would have two entries (corresponding to the two sorts of truth-conditional value) in our “lexicon” or “mental dictionary”—to invoke the accustomed metaphor of Chomskyan linguistics. Associated with each mental entry

for a name and instantiation of “every F” there would be distinct concepts.

Alternatively, the non-standard truth-conditional values may be explained away by appealing to Grice’s notion of implicature. The propositions containing these truth-conditional values as constituents, the argument would go, should be seen as pragmatically conveyed in the Gricean sense; they are not literally expressed by the sentence that was uttered. On this second approach, therefore, the non-standard truth-conditional values do not count as truth-conditional values at all—they are not part of what is said, the truth-conditional content of the sentence uttered—but are constituents of some other proposition that is implicated by the utterance of a sentence containing these expressions in a given situation.

The chapters to follow contain specific arguments against these two alternatives. In the next section, we briefly look at different kinds of linguistic evidence that motivate the pragmatist view in general.

5. Support for Pragmatism

Once more, pragmatism’s main idea is that “what is said”—the proposition, thought, or truth-conditional content expressed by a declarative sentence uttered in a

conversation—is underdetermined by the sentence's syntax and linguistic meaning (taking into account indexicality and ambiguity), and is instead fully determined only in context, by factors that cannot be straightforwardly traced to syntactic elements in the sentence. Pragmatism obtains intuitive support from examples such as the following:

A: "How's Maria coming along?"

B: "Oh, great; Maria finished the novel."

B's reply to A's question includes the complete and grammatically well-formed sentence,

(2) Maria finished the novel.

Suppose that upon hearing B's utterance, the two interlocutors understand or recover the proposition *Maria finished* ***writing*** *the novel*. How is this proposition connected to the sentence's meaning and syntactic structure? The specific information that Maria finished writing the novel cannot simply be read off (2). The sentence doesn't specify explicitly whether Maria finished writing, as opposed to reading, editing, proofreading, typewriting, designing, shredding, etc., the novel. Note that (2) could have been used to express the different proposition *Maria finished* ***reading*** *the novel*, for instance, in a different context. Thus the meaning and syntax of the sentence do not determine the proposition *Maria finished*

writing *the novel*. So it would seem that, contra Heim and Kratzer, to know the meaning of this sentence is not to know its truth conditions.

Instead, we may suppose that the two interlocutors recovered the proposition that they did because of various unstated facts and assumptions manifest to them both, for example:

(a) Maria was writing a novel.

(b) Writing is a process that can come to a conclusion.

(c) A and B are meeting to discuss Maria's writing progress.

(d) A is a publisher who commissioned the novel Maria finished writing and would thus be interested in Maria's progress with the novel.

(e) B is Maria's agent and can be presumed to be in the know concerning Maria's progress.

(f) B had recently mentioned to A that Maria was struggling with a mild case of writer's block.

(g) The deadline for the completion of the book is fast approaching.

…etc.

Here's a second example from (Searle, 1980):

> Consider the following sequence of rather ordinary English sentences, all containing the word "cut."

> [3] Bill cut the grass.
> [4] The barber cut Tom's hair.
> [5] Sally cut the cake.
> [6] I just cut my skin.
> [7] The tailor cut the cloth.
> …
>
> The feature of this list which interests me for present purposes, and which I will try to explain is this. Though the occurrence of the word "cut" is literal in utterances of [3]-[7], and though the word is not ambiguous, it determines different sets of truth conditions for the different sentences. The sort of thing that constitutes cutting the grass is quite different from, e.g., the sort of thing that constitutes cutting a cake. One way to see this is to imagine what constitutes obeying the order to cut something. If someone tells me to cut the grass and I rush out and stab it with a knife, or if I am ordered to cut the cake and I run over it with a lawnmower, in each case I will have failed to obey the order. That is not what the speaker meant by his literal and serious utterance of the sentence. (pp. 221-3)

As Searle suggests, manifold unstated facts and assumptions (what he calls "the Background") about the various activities we call "cutting" seem to play a role in the identification of the different truth-conditional contents that would be expressed by utterances of (3)-(7). They do so in ways that circumvent traditional semantic explanation, since "to cut," like "to finish" in our first example, is not ambiguous or traditionally viewed as context-sensitive.

To give a name to the phenomenon just illustrated, we may say that these expressions and the sentences that contain them are *truth-conditionally underdetermined*. The syntax

and linguistic meaning of these expressions and the sentences that contain them on their own do not suffice to yield a definite proposition (or the proposition actually recovered by the communicators in the situation). Expressions of virtually all syntactic types (nouns, verbs, adjectives, adverbs, prepositions, function expressions, etc.) exhibit this sort of truth-conditional underdetermination, as the following examples demonstrate:

NPs

(8) The table [in this room] is covered with books.

(9) We were playing baseball in the backyard [where the game we were playing bears little resemblance to the game of baseball as described by the Major League Baseball Rulebook].[22]

VPs

(10) She opened the door [with a key].

(11) Alex is writing the list [on a laptop computer].

Adjectives

(12) That catcher is talented [defensively].

(13) Sam is ready [to go on stage].

[22] This example is from (Bezuidenhout, 2002, p. 106).

Adverbs

(14) Chris is merely a good goalie [as opposed to having a striker's athleticism].

(15) Pat dresses stylishly [for a grad student].

PPs

(16) The cat is on the mat [attached by thin wires in a zero-gravity environment, as in Searle's (1979) example].

(17) John is waiting for us at the post office [outside, saving a parking spot for us while we drive around the block].

Function words

(18) He's not [what I'd call] a shrink, he's a psychiatrist.

(19) Paul and Mary got married and [then] had children.[23]

Each of these sentences is perfectly well-formed and meaningful as it stands. Yet it would appear that it is only when the sentences are uttered in contexts where the bracketed information is supplied—against a Background of

[23] These last two examples are from (Bach, 2006). For a state-of-the-art discussion on the semantics and pragmatics of "and," see (Carston, 2002, Ch. 3).

appropriate facts and assumptions—that they have the truth-conditions the communicators in the context would say they have or that they have any truth conditions at all. That is to say, without this Background, the truth conditions of most of these sentences seem indeterminate, or are simply not the truth conditions the hearers would intuitively give for them.

Such examples constitute a prima facie challenge for the mainstream view that meaning plus syntax determines truth conditions, as well as for the dominant Gricean perspective on the semantics/pragmatics divide. It is not clear that a single comprehensive strategy can succeed in explaining such widespread truth-conditional underdetermination from a traditional point of view.[24] To date, perhaps the most sophisticated attempt has been carried out by Stanley (2000; 2002a; 2002b; Stanley & Szabó, 2000; King & Stanley, 2004). Stanley's proposed solution to the underdetermination problem would treat the various types of truth-conditionally underdetermined expressions on the model of indexical expressions such as pronouns and demonstratives, which are well-understood vehicles of

[24] See (Bezuidenhout, 2002) for a review and critique of treatments that focus solely on certain types of expression and that purport to explain the underdetermination problem on the basis of notions like non-literality, ambiguity, vagueness, polysemy, incompleteness, or ellipsis.

contextual input. The gist of Stanley's proposal is that associated with each seemingly underdetermined expression there are hidden variables getting their values from the context. Stanley has worked out his account in detail only in the case of NPs, however. He suggests that the account can be extended to cover other types of expression. Nevertheless, there are reasons not to share Stanley's optimism, as even the fullest version of his proposal, concerning NPs, is fraught with difficulties.[25] Here we shall not stop to consider Stanley's work or its defects; we will simply assume that any attempt along those lines is unsuccessful.

The next section describes a more promising way of explaining these examples.

6. Relevance Theory[26]

We have elected to articulate our claim that reference is pragmatic in Relevance-Theoretic terms for two reasons. First, Sperber and Wilson's (1986/1995) Relevance Theory

[25] For criticisms of Stanley's strategy, see (Bezuidenhout, 2002), (Récanati, 2002), (Rett, 2005), (Collins, 2007), and (Neale, 2007a). In evaluating Stanley's approach, it is helpful to keep Neale's warning in mind: "we shouldn't get hooked on aphonics." (2004b, p. 188)

[26] This section is just a bare-bones presentation of Relevance Theory, discussing only those notions that we will need in the other chapters. For more details and considerations in favor of the theory, see the more comprehensive introductory accounts on which this section is based: (Sperber & Wilson, Ch. 1), (Wilson & Sperber, 2002), and (Wilson & Sperber, 2004). (Unger, 2001, Ch. 1) also provided a helpful guide.

(RT) offers a deep and general answer as to why and how speakers' intentions and other linguistically unmarked contextual factors may come into play in determining utterances' truth conditions. And second, from an empirical point of view, RT is currently the best grounded pragmatic theory.[27]

This section briefly goes over the main ideas of RT; certain claims will be taken up again and discussed in greater detail in other chapters.

6.1. Inference and Communication

In (1986/1995) Sperber and Wilson seek to answer the question: How does communication work? They distinguish two different possible answers, two rival models of communication: the "inferential model" and the "code model." The inferential model would explain communication roughly as follows. A communicator, A, intends to express some information, I, and produces a stimulus that would enable the addressee, B, to identify this information by recognizing A's intention to express it. The stimulus is designed to provide evidence on the basis of which B can infer the information A wishes to convey. For example, if

[27] For a literature survey of recent wide-ranging experimental developments based on RT assumptions, see (Carston & Powell, 2006).

A wants B to take out the trash, she might touch B's arm and point to the garbage can, wave the trash bag at B, or say to him, "Take out the trash."

According to Sperber and Wilson, A must be attributed two intentions: 1) an intention to inform B of something, which they call an *informative intention*; and 2) the intention to inform B of this informative intention, which they call a *communicative intention*. The two intentions are necessary, since one can intend to inform someone of something without communication taking place as a result.[28] For instance, suppose that A wants to inform B that the trash needs to be taken out, but does not want to ask him openly to do it. In order to make her intention known, A places the recycling bin next to the front door. A hopes that upon seeing the recycling bin B will understand that he is to take out the trash. However, no real communication between A and B has occurred: B might see the recycling bin and not realize that the trash needs to be taken out, and A might rationalize his inaction by telling herself that perhaps B thought that A was in the process of taking it out herself. Sperber and Wilson emphasize that communication is intuitively thought of as

[28] Cf. (Grice, 1957/1989). On the fundamental role of intentions in communication, see also (Neale, 2004b, p. 181).

something we do overtly; communication should be distinguished from covert forms of information transmission.[29] (1986/1995, p. 30)

In contrast to the inferential model just sketched, the code model would describe communication as follows. A communicator, A, encodes a piece of information, I, in a signal, S, which is then decoded by a second communicator, B. A's intentions are irrelevant for the recognition and understanding of I by B. What matters is that B be in possession of the code in which I is encrypted. Computers exchanging information encoded in ASCII numbers would be an example of communication via a code.

In essence, linguistic communication on the code model works like mailing a package: the sender (the speaker) places the item (the thought) in a box (the sentence) and the recipient (the hearer) takes out the very item from the box (understands the proposition expressed by the sentence). Comprehension is thus really the duplication of the thought of the speaker in the mind of the hearer.

Sperber and Wilson credit Grice with being the first to question the code model, which according to them "is very well entrenched in the Western scholarly tradition."

[29] The idea that communication is essentially overt derives from Grice; see, e.g., (1957/1989).

(1986/1995, p. 24) Grice challenged the code model by drawing attention to the fact that in normal conversations interlocutors implicate propositions that do not follow logically from the words that they say, and very often such implicatures are what the communicators are most interested in. In (1975/1989) Grice explicitly typified "implication" as a form of inference, which he contrasted to logical entailment.

However, Grice did not seem to regard the interpretation of *what is said* as a matter of inference, but rather of decoding.[30] As we saw above, he characterized "what is said" as a proposition compositionally determined by an utterance's linguistic meaning and syntactic structure. Thus for Grice commonality of language between speaker and audience would appear to be all that is required to communicate verbally in a direct manner (without implicating).

Sperber and Wilson take Grice's challenge to the code model one step further: inference is required not only for determining what is implicated, as Grice proposed, but also for determining what is said. They do not deny that verbal communication involves decoding. Words' linguistic meanings encode specific information, for example. To understand an

[30] Here we agree with Carston, who says: "When it came to what is explicitly communicated [what is said], Grice was pretty much of a code theorist." (2002, p. 107)

utterance in a language that one knows entrains the decoding of the information encoded by the sentence's meaning. But decoding is not all there is to identifying the proposition expressed by an utterance, Sperber and Wilson argue.

Two phenomena that would at first sight appear to pose a problem for the code model are indexicality and ambiguity. Indexicals do not encode their truth-conditional values and ambiguous expressions do not, on their own, indicate which of their meanings is the appropriate one in a given situation. Also, the fact that comprehension is possible and fairly easy despite problems in the "transmission of the code" produced, for example, by ungrammaticality, code-switching,[31] and malapropisms and other slips of the tongue, provides further evidence against the code model. In addition, Sperber and Wilson cite examples like those in §5 to show that the discrepancy between what is linguistically encoded in an utterance and the proposition expressed in a context is even greater than suggested by these issues.

Sperber and Wilson conclude that verbal communication is a matter of both decoding and inference. They view the information encoded by a sentence's meaning as *part* of the evidence upon which to infer the information a speaker

[31] Code-switching is the mixing of phonological, syntactic, and lexical elements of two or more languages in a single utterance or conversation. Examples would be Franglais, Spanglish, and Portuñol.

intends to get across. Together with other evidence, including specific contextual information, the decoded material serves to justify a particular conclusion about the speaker's message.

Now, generally speaking all evidence can be interpreted in a variety of ways, warranting different conclusions, even mutually incompatible ones. Contextual facts are no exception. For instance, in the example in §5, the knowledge that Maria was writing a novel, which A and B share, could warrant the different conclusion that in uttering (2) B is expressing the proposition *Maria finished* ***reading*** *the novel*. A could conclude this, for example, if A and B also shared the knowledge that A had recently given Maria a novel she thought would offer inspiration for her writing, wanted her to read it right away, and feared her writing might get in the way of her reading.

The problem for Sperber and Wilson is then to explain how people are able to understand each other, how the audience is able to identify—seemingly effortlessly and with great accuracy—the precise proposition a speaker intends to express by her utterance, given that the evidence available to the audience is open to interpretation and that the proposition's linguistic garb is insufficient for this purpose.

Sperber and Wilson's explanation is founded in an account of cognition. The mind, they argue, is spontaneously able to pick out the right meanings, truth-conditional values, and linguistically unmarked contextual factors to identify the proposition intended by a speaker. That is, the mind has a natural tendency to hone in on information that is relevant in a technical sense. The next subsection discusses Sperber and Wilson's notion of relevance.

6.2. Relevance

According to Sperber and Wilson, cognition has one overriding goal: the improvement of a subject's knowledge or representation of the world. The mind seeks to construct an ever truer picture of reality. At every moment, though, the mind is bombarded with information; how does it select which stimuli to attend to? Sperber and Wilson answer that the mind spontaneously attends to that information that furthers its main goal of enhancing knowledge of the world. Information that does this is relevant information. This idea they express as the Cognitive Principle of Relevance:

> Cognitive Principle of Relevance
>
> The mind tends to be geared to the maximisation of relevance. (1986/1995, p. 260)

Sperber and Wilson characterize relevance as a property of an input to the mind; the input may be perceptual/sensory (e.g. an utterance) or psychological (e.g. a belief). The degree of relevance of an input is a matter of benefits versus costs. The more an input aids in gaining knowledge about the world, the more relevant it is. But the more effort is expended by the mind to process the input, the less relevant it is. Thus relevance may be construed as a sort of ratio.[32]

The benefits of an input Sperber and Wilson call its *positive cognitive effects*.[33] They identify three main types: 1) *contextual implication*, 2) *contextual strengthening*, and 3) *contextual contradiction and elimination*. A contextual implication is new knowledge that results from conjoining existing knowledge and assumptions with contextually available information. For example, from glancing at her watch and her knowledge of her flight's departure time, A derives the conclusion that the flight is delayed. Contextual information may also strengthen or provide further evidence for existing beliefs and knowledge. For instance, the flight's delayed status might confirm A's impression that service on this airline is deteriorating.

[32] Wilson and Sperber stress that relevance should not be understood as an "all-or-none matter." (Wilson & Sperber, 2004, p. 252)
[33] They also use the term "positive contextual effects."

And contextual information may lead to the revision or abandonment of mistaken beliefs and assumptions. For example, the delay might force A to change her dinner plans upon arriving home.[34]

An input's relevance to an individual is not absolute, but varies according to such factors as the individual's interests, knowledge, abilities, and time at which the stimulus occurs. It is also relevant in relation to other stimuli in the environment. The difference between an input that is relevant to an individual and other stimuli that are not is always one of degree, not kind. A relevant input is simply *more* relevant than other stimuli. That is, it has more positive cognitive effects than the other stimuli. For instance, the sight of the plane taxiing to the gate is more relevant to A than the sight of a half-eaten sandwich in the waiting room seat next to hers, since it allows her to forecast better the extent to which she is going to have to reorganize her evening.

As mentioned above, relevance is also a matter of costs. Sperber and Wilson cash out the idea of an input's costs in terms of the time it takes the mind to process the input. The longer it takes to process, the lower the input's

[34] These examples of positive cognitive effects are based on those offered in (Wilson & Sperber, 2004, §2).

relevance. Two examples of types of inputs that would generally be of low relevance are 1) new information that is entirely unconnected to existing knowledge and assumptions and 2) information that is already known. These two sorts of inputs would make the mind spin its wheels gratuitously and without end, since they cannot add to our knowledge, provide new evidence for our beliefs, or cause us to revise or eliminate previously held assumptions. In short, they do not increase our knowledge of the world.

To sum up, Sperber and Wilson's Cognitive Principle of Relevance says that the mind spontaneously attends to information that yields the greatest positive cognitive effects for the least processing effort. The principle is put forth as descriptive of human cognition, and not as a normative ideal. What it affirms is an overall tendency; Sperber and Wilson emphasize that the principle and the concepts of a cognitive effect and processing effort may not always apply in individual cases. (1986/1995, pp. 261-3) And they explicitly deny that their theory is intended to offer an analysis of the ordinary concept of relevance. In the next subsection, we look at how Sperber and Wilson's cognitive notion of relevance transfers over to the case of verbal communication.

6.3. Relevance and Verbal Communication

In §6.1 we mentioned Sperber and Wilson's point that true communication is overt. A communicator openly seeks her addressee's attention. Now, if according to the Cognitive Principle of Relevance attention tends automatically to go to what is most relevant to an individual at a given moment, a communicator will succeed in capturing the addressee's attention only if her utterance is perceived as relevant by the addressee at that moment. In this regard, however, the speaker has an advantage over the non-linguistic parts of the environment. Sperber and Wilson argue that unlike other inputs, utterances always involve a presumption of relevance. The speaker, by the very act of making an utterance, raises the expectation that her utterance is relevant to the audience. Sperber and Wilson dub this claim the Communicative Principle of Relevance:

> Communicative Principle of Relevance
>
> Utterances create expectations of optimal relevance.[35]
> (Wilson & Sperber, 2004, p. 249)

[35] A more general statement of the principle, not in terms of utterances, can be found in (1986/1995, p. 158).

Sperber and Wilson note that the idea that utterances raise certain expectations comes from Grice. Grice held that in making an utterance, a speaker should be presumed to be following a general Cooperative Principle and certain Maxims of Conversation:

Cooperative Principle

> Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

Maxim of Quantity

1. Make your contribution to the conversation as informative as necessary.
2. Do not make your contribution to the conversation more informative than necessary.

Maxim of Quality

1. Do not say what you believe to be false.
2. Do not say that for which you lack adequate evidence. (In general, try to make your contribution one that is true.)

Maxim of Relevance

> Be relevant [i.e., say things related to the current topic of the conversation].

Maxim of Manner

1. Avoid obscurity of expression.
2. Avoid ambiguity.
3. Be brief (avoid unnecessary prolixity).
4. Be orderly. (Grice, 1975/1989, pp. 26-7)

In presuming that the speaker is following the Cooperative Principle, the audience develops the expectation that the speaker's utterance will be appropriate to the ends of the conversation in which they are engaged with the speaker. And in presuming that the speaker is following the Maxims, the audience expects that the utterance will be as informative (Maxim of Quantity), true (Maxim of Quality), relevant (Maxim of Relation), and clear (Maxim of Manner) as the speaker can possibly make it (within the constraints of the goals of the conversation). Perceived violations to the Cooperative Principle and the Maxims by the speaker, together with contextual factors, may lead the audience to infer certain implicatures, or may simply result in miscommunication.

Sperber and Wilson disagree with Grice that utterances carry these four additional kinds of expectation; they believe that the Cooperative Principle and the Maxims can be subsumed under a single Communicative Principle of Relevance. Here we cannot address Sperber and Wilson's arguments on this score.[36] It is important to point out, however, that Grice's seemingly ordinary notion of relevance has little to do with the technical cognitive

[36] See (1986/1995, §§4-7) for criticisms of Grice's approach. Some of their first arguments in favor of a single principle appeared in (Wilson & Sperber, 1981). In (Wilson & Sperber, 2002), they demonstrate how some of Grice's maxims conflict with one another.

notion of relevance elaborated by Sperber and Wilson. "Relevance" in the Communicative Principle should be understood in this technical sense, so that the principle may be interpreted as saying that an utterance creates the expectation that it will generate the most cognitive effects for the least processing effort for the addressee.

Sperber and Wilson contend that an utterance's guarantee of relevance motivates a particular comprehension strategy for the audience. That is, if each utterance promises its own optimal relevance, then the audience is justified in following a path of least effort in interpreting the utterance and identifying the proposition expressed by it. (Given that the more relevant the input, the less processing effort required.) Sperber and Wilson call this the Relevance-Guided Comprehension Heuristic.

Relevance-Guided Comprehension Heuristic

> (a) Follow a path of least effort in constructing an interpretation of the utterance (and in particular in resolving ambiguities and referential indeterminacies, in going beyond linguistic meaning, in supplying contextual assumptions, computing implicatures, etc.).
>
> (b) Stop when your expectations of relevance are satisfied. (Wilson & Sperber, 2004, p. 260)

This comprehension heuristic is presented as a reconstruction of a psychological process that is supposed

to take place mostly unconsciously and very fast. Sperber and Wilson do not claim that the audience works out consciously and step-by-step what the speaker is saying by following this procedure. As an illustration of how this interpretation heuristic is supposed to work, consider the following example.

A and B are tasting wines. A says to B:

(20) This riesling is sweet.

B interprets A's utterance of (20) as expressing the proposition *the riesling at the tasting is not dry*. B's interpretation process may be reconstructed as follows. By the Communicative Principle of Relevance, A's utterance should be optimally relevant to B, i.e. it must not require much processing effort and must have positive cognitive effects. The following information is available from the context and is manifest to both A and B: they are at a wine tasting; both have a basic knowledge of wines; rieslings can be either dry or sweet; B dislikes sweet rieslings. So an interpretation in which "this riesling" refers to the riesling offered at the tasting and "is sweet" means that the wine has perceptibly retained some of the fruit's natural sugar appears salient and easily inferable by B. Such an interpretation has a specific positive cognitive effect; it yields the contextual

implication that B will probably not like the riesling. The implication is as follows: B dislikes sweet rieslings (contextually given assumption); the riesling at the tasting is a sweet riesling (what B understands A to have said by her utterance); therefore B will probably dislike the riesling at the tasting (contextual implication). Furthermore, this contextual implication is likely to lead to other implications for B, e.g. B would do well to avoid the riesling and focus on other wines. So the presumption of relevance appears satisfied in this case.

However, it is possible to interpret A's utterance in other ways. For instance, by "sweet" B could have taken A to mean "fortified with sugar." A's utterance could therefore have been interpreted by B as expressing the proposition that the riesling at the tasting has the sweetness of a fortified wine. However, in this case A's utterance would not be as relevant to B, since this latter proposition would not be the first to occur to B on an interpretive path of least effort. To come up with an interpretation of A's utterance according to which it expresses the proposition *the riesling at the tasting is fortified with sugar* B would have had to consider and

reject[37] the more salient and common meaning[38] of "sweet" of "not dry." B would also have had to ignore a contextual fact that was stipulated to be manifest to both A and B, namely, that rieslings are either dry or sweet, but seldom fortified—a fact they would know given their basic oenological knowledge, which is also contextually presupposed. In other words, B would have had to construct and then discard the interpretation sketched in the previous paragraph, according to which A's utterance expresses the proposition that the riesling at the tasting is not dry. This action, however, would demand extra processing effort. So even if the second interpretation were to produce positive cognitive effects for B, A's utterance still fails at being optimally relevant for B, inasmuch as it would have required additional processing effort. Hence, such a second interpretation, though conceivable, is unwarranted from a Relevance-Theoretic point of view.

Let us now see how RT would handle utterances of truth-conditionally underdetermined sentences. Consider the

[37] Again, comprehension is fast and largely unconscious according to Sperber and Wilson. Our explanation, however, employs terms that connote conscious deliberation. This connotation should be ignored.
[38] "Not dry" is the more common meaning of "sweet" in the context of wine. The word does have this linguistic meaning, as can be checked by consulting a dictionary.

first example of §5 once again. From B's utterance of a truth-conditionally underdetermined sentence,

(2) Maria finished the novel,

A understands the more specific proposition *Maria finished* ***writing*** *the novel*. A's comprehension process may be outlined as follows.

The Communicative Principle of Relevance dictates that A will find B's utterance optimally relevant to her; it will have cognitive effects for little processing effort. Certain contextual information is manifest to both A and B, in particular:

(a) Maria is writing a novel.

(b) Writing is a process that can come to a conclusion.

Following an interpretive path of least effort, A takes "Maria" to refer to the Maria they both know and "the novel" to signify the novel being written by this Maria. From contextual assumptions (a) and (b), in conjunction with her knowledge of the linguistic meaning of "to finish," A interprets "finished" to mean "finished writing." Thus, with little effort, A interprets B's utterance of (2) to express the proposition *Maria finished* ***writing*** *the novel*.

Interpreted to express this proposition, B's utterance of (2) has positive cognitive effects for A. For instance, it

would lead her to abandon her belief that it will still take Maria a few more months to complete the manuscript. In other words, B's utterance of (2), taken to express the proposition *Maria finished* ***writing*** *the novel*, enhances A's knowledge of Maria's writing progress, which had been the motivation for her question. In generating this positive cognitive effect while at the same time requiring little processing effort, B's utterance satisfies the expectation of relevance the Communicative Principle of Relevance says it has.

7. Reference and Relevance

Reference, this dissertation argues, is pragmatically determined. By "pragmatically determined" we mean "determined (= identified) in accordance to the principles and concepts of Relevance Theory," the most important of which were discussed in this chapter. The explanations offered in this dissertation of the varying truth-conditional values of names and the phrase "every F" will be modeled on the explanations of relevance-guided interpretation delineated in §6.

In relying on Relevance Theory, the dissertation makes substantive empirical assumptions. In particular, it presumes that Sperber and Wilson's Cognitive Principle of

Relevance is true of human cognition. Neither this principle nor the other main concepts of Relevance Theory will be explicitly argued for in the dissertation, however. They will simply be taken for granted. It is not up to us, but is the job of the originators and proponents of the theory, to continue to provide evidence for it. Relevance Theory is merely an instrument that we will use to present more sharply and fully the two case studies that make up the heart of the dissertation. Our goal is not to contribute to the development of Relevance Theory, but to show that these two cases constitute counterexamples to the widely held view, neatly expressed by Borg, that "all noun phrases are either quantifier expressions or referring terms."[39] (2000, p. 230)

[39] This is also the view, expressed in virtually the same terms, that Neale held early on in (1990) and (1993). See, e.g., (1990, p. 6).

Chapter 2

A Critique of Recent Defenses of the Direct Reference Theory of Proper Names

1. Introduction

A wide consensus exists that proper names[40] are referring expressions. Central to this view is the idea that a name, if used literally, has an object as truth-conditional value. There is less agreement about what further information names may encode, if any. The aim of the next three chapters is to challenge this consensus view of names. It is argued, first, that names encode various sorts of descriptive information; second, that names, when interpreted literally, may have truth-conditional values other than objects (in particular, they may have properties or sets as truth-conditional values); and third, that the type of truth-conditional value a name may have on a given occasion of use is pragmatically determined by relevance-guided processes of interpretation.[41] This chapter concerns a theory that contradicts the first two of these theses, and which represents the most radical expression of the

[40] This chapter and the next two focus on ordinary personal names such as "John Smith." It is reasonable to suppose that most, if not all, of what will be concluded concerning personal names would apply to names of other types of object as well.

[41] See Chapter 1, §6, for an overview of Relevance Theory.

consensus view: Direct Reference (DR). The following chapter presents positive arguments for the three claims.

Our intent is not to offer a comprehensive survey of historical and contemporary philosophical views on the semantics and pragmatics of names. Rather, we concentrate solely on the particular issues connected to these three theses. Here we make no claim that our interpretations of the views of historical authors are sanctioned by current exegesis of these authors' works. Nevertheless, we believe that they are accurate enough for the dialectical ends of the chapters.

The discussion of this chapter proceeds as follows. §2 defines DR and reviews three well-known objections to it. §§3–8 discuss three different strategies, proposed by Braun (1998, 2001), Salmon (1986, 1989), and Soames (2002, 2004, 2007), for rescuing DR from these objections. A summary of each proposal is followed by a discussion of its shortcomings. This chapter does not present the whole case against DR; in the next chapter, certain considerations will be discussed which may be viewed as additional reasons for rejecting DR.

2. Definition of Direct Reference[42] and Three Objections

For the purposes of this chapter, a Direct Reference (DR) theory of names may be defined as any one of the following three claims:[43]

Direct Reference

DR1. The sole semantic function of a name is to refer to its bearer.

DR2. A name does not have a meaning beyond the object to which it refers.

DR3. The truth-conditional value or contribution of a name is its bearer.

DR may be traced back to (Mill, 1867/2001). In that work Mill described proper names such as "John" as expressions having "denotation" but no "connotation." For Mill, the denotation of an expression is the thing or things to which the expression "is capable of being truly affirmed," i.e. the entities to which the expression applies. For example, the noun "man" denotes ("is capable of being truly affirmed of") an individual man. Mill takes "man" to be a general name for the set of all men, and he holds that the term may

[42] Direct Reference is also known by other appellations: "Millianism," Russellianism," "Neo-Russellianism," "Naïve Russellianism," "The New Theory of Reference," "The Naïve Theory," and "The 'Fido'-Fido Theory."
[43] Braun (1993, pp. 449-50) identifies and endorses these three claims.

be truly affirmed of each member of the set. (p. 266) The *connotation* of an expression is characterized by Mill as information conveyed by the expression that regards the attributes possessed by the thing to which the expression applies. An expression with connotation, says Mill, "implies, or involves, or indicates, or…*connotes*" the attributes of something to which the expression may be truly applied. (p. 268, emphasis in the original) So "man," for example, connotes an attribute that all men have in common; one such attribute, suggests Mill, is rationality. (p. 268)

Mill may be interpreted as assimilating the notion of connotation to the notion of meaning. He says:

> …whenever the names given to objects convey any information—that is, whenever they have properly any meaning—the meaning resides not in what they *denote*, but in what they *connote*. (p. 270, emphasis in the original)

According to Mill not all expressions connote: there are connotative expressions and non-connotative expressions. Proper names are paradigmatic examples of non-connotative expressions:

> Proper names are not connotative; they denote the individuals that are called by them, but they do not indicate or imply any attributes as belonging to these individuals. (p. 269)

> [T]he only names which connote nothing are *proper* names; and these have, strictly speaking, no signification.... (p. 270, emphasis in the original)

A proper name, then, simply stands for, or "signifies directly" (p. 268), an individual. The reason it does is not that the individual possesses certain attributes connoted by the name—attributes through which we could determine that the name may be "truly affirmed" of it—but simply because the individual is called by that name.

This, we shall assume, is the basic idea of DR as articulated by Mill. Theses D1–DR3 may then be taken to give concise expression to this idea. Contemporary defenders of DR include Kaplan (1989), Salmon (1986, 1989), Soames (2002, 2004, 2007), and Braun (1993, 1998, 2001). DR is also often attributed to Kripke (1971, 1970/1980).

Among the considerations motivating the contemporary versions of DR are the following. First, there is the feeling that a name, as Marcus (1961/1993, p. 12) put it, is "just a tag." Unlike nouns such as "lamp," there are no criteria which must be fulfilled in order for an object to be named "N;" "N" is often said just "to stand for" that object. Second, and related to this point, there is the fact that people introduce new names all the time. Seemingly, naming is an act over which people have complete control. By comparison, we have little control over the

meaning and application of most other words. One cannot simply choose to use "lamp" to mean "cat," for example, and hope to be understood by others. One can, however, name one's cat "Lamp," if one so wishes.[44] Third, there is the assumption, made by many philosophers, that many aspects of natural language may be fruitfully elucidated through the application of formal tools. The implication in this case is that natural language names may be understood on the model of logical constants, which simply represent objects in a universe of discourse, and have no further meaning.[45] Finally, many philosophers have been impressed by Kripke's (1970/1980) critique of Descriptivism, and have been led to the conclusion that names do not encode *any* descriptive information, and hence, to DR. Some of Kripke's claims will be examined in Chapter 4.

DR faces three well-known difficulties. First, there is Frege's Puzzle concerning the informativeness of identity statements. DR2 says that a name does not have a meaning beyond the object to which it refers. So, if two names "N_1" and "N_2" refer to the same object, a sentence of the form "$N_1 = N_2$" would seem to predicate the identity of the referent with itself.[46] The puzzle is how, despite seeming

[44] This point is made by Récanati (1993, p. 136).
[45] Salmon (1989) explicitly endorses the names-as-constants view.
[46] By "referent" is meant an individual bearing the name.

to express this triviality, the sentence can nevertheless be informative to someone.

For example, consider (1).

(1) George Orwell is Eric Blair.

On DR, (1) would seem to predicate the identity of the object denoted by these two names with itself. Which is what (2) also predicates:

(2) Eric Blair is Eric Blair.

(1) and (2) are both true. But intuitively, (1) seems more informative than (2). Someone might affirm (2) but deny (1), for example.

Although in (1892a/1997a) Frege originally formulated the informativeness puzzle in terms of identities such as (1) and (2), the puzzle can be straightforwardly adapted to ordinary predications, as Dummett (1981, pp. 125-6) points out. Compare (3) and (4):

(3) George Orwell was a writer.

(4) Eric Blair was a writer.

For someone—say, A—who knows that George Orwell was a writer, but does not know that George Orwell and Eric Blair are (or were) the same individual, (4) would represent new information. A might affirm (3) but deny (4), for example. But how (3) and (4) may convey different and potentially new information is at first sight a mystery on DR, since

according to DR, (3) and (4) would seem to say the same thing; they ascribe the same property to the same individual.

The second problem facing DR is Frege's other puzzle, also presented in (1892a/1997a), which concerns the failure of substitutivity of co-referential terms within non-extensional contexts.[47] If a name's sole truth-conditional contribution is its referent, as DR3 says, then co-referential names should be substitutable for each other salva veritate.[48] But substitution of co-referential names within the "that-clauses" that typically follow propositional attitude verbs[49] often fails to preserve the truth value of the sentence. For example, (5) may be true but (6) false:

(5) A believes that Eric Blair is Eric Blair.

---

[47] A linguistic context is non-extensional if the expressions in it do not allow for truth-value-preserving substitution of co-referring terms. A linguistic context whose expressions allow such substitution is extensional.

[48] This supposition is often justified by an appeal to Leibniz's principle of substitutivity: "*eadem sunt quorum unum in alterius locum substitui potest, salva veritate*." ("Two things are the same that can be substituted for one another, while preserving truth.") Englebretsen (1997, p. 38) notes that this formulation of the principle does not specify whether the things to be substituted are linguistic expressions or some other kind of thing. Englebretsen presents the following quotation from Leibniz, which makes it clear that Leibniz had linguistic expressions in mind: "Those terms are 'the same' or 'coincident' of which either can be substituted for the other wherever we please without loss of truth—for example, 'triangle' and 'trilateral.'" (Both passages from Leibniz are quoted in [Englebretsen, 1997, p. 38]).

[49] Propositional attitude verbs are verbs that, when followed by a sentence, create a non-extensional context. Examples are "to believe," "to know," "to fear," "to feel, "to think," "to say," "to assert," etc.

(6) A believes that Eric Blair is George Orwell.

Specifically, (5) is true if A believes that Eric Blair is self-identical and (6) false if A doesn't know or believe that Eric Blair is George Orwell. It appears to follow from DR3, however, that if (5) is true, so must (6), since "Eric Blair" and "George Orwell" are co-referential. In both cases, DR seems to imply that the sentences following "that" predicate the identity of the referent to himself. But simply from the fact that "Eric Blair" and "George Orwell" are co-referential it doesn't follow that (6) is true, or that the belief that Eric Blair is George Orwell can be attributed to A.

The third problem is the problem of so-called empty names.[50] An empty name is a name that has no referent. Three consequences follow from DR2 and the fact that some names are empty. First, DR2 would compel us to say that empty names such as "Winston Smith," i.e. the name of the fictional character, are completely meaningless. Second, a sentence containing an empty name, such as (7),

(7) Winston Smith doesn't exist,

would also be meaningless, since one of its component expressions is meaningless. Third, a sentence containing

[50] A more specific variant of the empty names problem, the problem of empty names in "negative existential sentences" was first presented in (Russell, 1905).

an empty name, such as (7), would appear to be neither true nor false.

But, our intuitions are that "Winston Smith" is not entirely devoid of meaning; that (7) is perfectly meaningful; and that (7) has a determinate truth value. ([7] seems true.)

3. Braun's Psychological Strategy

In (1998) and (2001), Braun put forth a proposal for defeating these objections to DR.[51] The basic idea of Braun's proposal is that there is a distinction to be made between the proposition semantically expressed by a sentence containing a name and the ways a person may understand, entertain, or believe that proposition.[52] A person may believe one and the same proposition in different ways; as a result, different beliefs, represented as different belief reports expressing the same proposition, may or may not be attributable to someone (depending on whether or not the person believes in the proposition in ways that correspond to the reports), and moreover, a person's believing a proposition in different

[51] Braun's proposal is similar to the account presented in (Perry & Crimmins, 1989). This chapter will not consider Perry and Crimmins' account. It assumes with Schiffer (2006, pp. 365-6) that Braun's version represents the best version of the psychological strategy.
[52] Braun articulates his account in terms of belief; this summary of his account does so as well.

ways would explain the informativeness, to that person, of multiple sentences expressing the same proposition. There are thus two essential components to Braun's solution to the informativeness and substitution problems: the notion of a singular proposition and the notion of a way of believing a proposition. Braun's solution to the empty names problem requires an additional ingredient, which we will introduce and discuss once we have reviewed Braun's understanding of propositions.

Braun, like other DR advocates,[53] endorses Russellianism about propositions.[54] According to Russellianism, well-formed sentences express propositions; a proposition is the truth-conditional content of a sentence. A Russellian proposition is a complex, structured entity: the constituents of a proposition might include objects, properties, relations, and other propositions. If a sentence contains a singular term such as a name, the sentence is said to express a "singular proposition;" and one of the constituents of the proposition will be a single object. Thus, a singular proposition may be represented as an ordered pair <o, F>, where o is the object referred to by the name and F is some property predicated of it by the

[53] See, e.g., Kaplan (1989, pp. 494, 496), Soames (2002, pp. 55-6), and Récanati (1993, pp. 26-44).
[54] The "Russellian" view of propositions that is described here is based on Kaplan's (1989) interpretation of Russell.

expressions making up the rest of the sentence. For example, the proposition expressed by (3) is represented as <Eric Blair, being a writer>. And this proposition is also the proposition expressed by (4).

Braun explains his notion of a way of believing via the notion of a belief state. A belief state is a psychological state, "a state of the brain or soul." (1998, p. 573) Braun doesn't characterize exactly the relation between a belief state and a believed proposition, but he assumes that it's plausible to distinguish between the two. Once this distinction has been made, he says, it is also plausible to think that one and the same proposition may be related to different belief states. (pp. 573-4) He justifies the plausibility of these assumptions by suggesting that belief states can be characterized independently of the propositions to which they are related, through the different causal roles such belief states might play. For instance, A's being in the belief state of *knowing* that George Orwell is Eric Blair may dispose him to accept the assessment that Eric Blair is a good writer. And A's being in the belief state of *wondering* whether George Orwell is Eric Blair may cause him to begin looking for information that would decide the question. Differences in causal roles point to differences

in ways of believing propositions. On the basis of these two assumptions about belief states, says Braun, "we could then plausibly say that each of the belief states is a *distinct way* to believe [a] proposition P." (2001, p. 6)

Braun suggests that this explanation could be fleshed out further by appealing to the notion of a mental sentence and Schiffer's (1981) metaphor of a "belief box." Braun, like other philosophers, such as Fodor (1975), believes it's plausible to assume that thinking occurs in a kind of mental language, a "language of thought." Mental sentences, he supposes, are similar to natural language sentences in that they possess syntax and certain semantic properties.[55] (1998, p. 574) In particular, like natural language sentences, mental sentences express propositions. For a subject to be in a mental state involving[56] a mental sentence is for her to have the sentence in her belief box. To say that a person has a mental sentence in her belief box is supposed to be a colorful way of saying that she has a belief (understood as a mental sentence), whether the belief be consciously held or not.

A way of believing a proposition can then be characterized as a mental sentence expressing that

[55] More on Braun's conception of mental sentences in §4 below.
[56] As noted, Braun doesn't characterize exactly the relation between a mental sentence and the mental state which "involves" it.

proposition which one has in one's belief box. Just as different natural language sentences may express the same proposition, suggests Braun, different mental sentences may also express the same proposition:

> [T]here might be distinct mental sentences that have the same propositional content, just as there are distinct natural language sentences that have the same propositional content. (2001, p. 7)

So, to believe the same proposition in different ways is to have different mental sentences expressing the same proposition in one's belief box.

The explanation is supposed to work as follows. Braun would say, for instance, that in virtue of having different mental sentences in his belief box, A may believe the proposition expressed by (1) and (2), namely, <Eric Blair, Eric Blair, Identity> in different ways. In the same fashion, proposes Braun, we can explain how A may rationally believe this proposition and its negation. For example, this would be the case if A had in his belief box the mental equivalent of (2) and the equivalent of the negation of (1).

Braun's psychological account answers the informativeness objection to DR by showing how one and the same singular proposition may be believed in different ways by a subject. It aims to translate the loose, intuitive notion of

informativeness into the more precise notion of a way of believing. Thus, someone may find (1) more informative than (2), in the sense that he may believe them in different ways. However, Braun insists, the two sentences express the same singular proposition.

Essentially the same account explains the substitution problem. For example, suppose B hears the reports (8) and (9).

(8) A thinks that George Orwell is a writer.

(9) A thinks that Eric Blair is a writer.

She may believe the first and not the second, even though the two reports express the same proposition, since she believes this proposition in two different ways.

In believing that (8) is true and that (9) is false B would not be guilty of irrationality, according to Braun. She merely believes the proposition expressed by the two sentences in different ways. She is in different belief states involving the same proposition. As in the identity case, such belief states may be further characterized as mental sentences expressing the same proposition which figure in B's belief box.

To account for the empty names problem, Braun first modifies the standard Russellian picture of propositions. Elaborating upon an idea due to Kaplan (1989, p. 496, fn.

23), Braun (1993) introduces the notion of a “gappy” proposition, which he describes as a structure possessing an unfilled position. For example, the gappy proposition expressed by (10),

(10) Winston Smith is a clerk,

consists of the following structure: <_, being a clerk >. “_” represents the slot normally occupied by the referent of a name. “Winston Smith” has no referent so in this case the slot is unfilled.

Braun’s gappy proposition account, coupled with his notion of a way of believing a proposition, would explain our intuitions regarding the meaningfulness of empty names and the meaningfulness and truth-evaluability of sentences containing empty names. (10) would seem meaningful and truth-evaluable for someone, for example, were this person to have a mental equivalent of the sentence in her belief box. The mental sentence, like the English sentence (10), expresses a gappy proposition. Such a combined account would also explain how someone may rationally believe (11) but deny (12), for example.

(11) Zeus lives on Mt. Olympus.

(12) Jupiter lives on Mt. Olympus.

Someone who believes (11) but denies (12) believes the same gappy proposition in different ways.

4. Problems with Braun's Psychological Strategy

There are two sorts of problems with Braun's account, those having to do with the notion of a singular proposition, on the one hand, and those having to do with the notion of a way of believing, on the other. Here the problems having to do with the notion of a singular proposition will be set aside, since they involve thorny metaphysical issues that would take us too far afield of our main concerns.[57] This section discusses problems relating to Braun's notion of a way of believing.

The main problem with Braun's notion is its dubious explanatory value. On the face of it, his explanation in terms of belief states (or mental sentences in a subject's belief box) seems weak, since the explanandum—the phenomena of informativeness, failure of substitutivity of co-referential names, and meaningful though empty names—seems intuitively clearer than the explanans he proposes. (As a side point, DR is a semantic theory of names, so one would expect that its adequacy should be entirely assessable by empirical linguistic methods, and not by a consideration of controversial psychological and metaphysical matters.)

[57] We refer the reader to (Plantinga, 1983) for a critique of the notion of a Russellian proposition.

Most importantly, however, Braun's account of the three objections in terms of ways of believing seems merely to push the objections back one (or perhaps more) level(s); the account doesn't offer a definitive answer to them.

What follows concentrates on Braun's explanation of the informativeness objection. The substitution objection can be safely ignored, since his explanation of this issue is not fundamentally different from his explanation of the informativeness objection. The empty names issue can also be set aside, since the notion of a way of believing a proposition is independent of the question of which sorts of propositions one may believe. The notion is supposed to hold for both gappy and regular propositions; as we saw, a subject may also believe the same gappy proposition in different ways.

One difficulty with Braun's notion of a way of believing a proposition is that it depends on a particular, and disputable, view of mental sentences.[58] Namely, it explicitly presupposes a "public language" view of mental sentences. The public language view holds that thinking occurs in a language identical or very similar to one's

[58] Braun does say that one does not have to accept the mental sentence picture of ways of believing in order to accept his psychological explanation (1998, p. 576); however, since Braun nevertheless *relies* on this picture to articulate his account, it is fair to point out some of its defects.

public language, i.e. a natural language such as English.[59] Braun writes:

> To simplify matters, let's assume that the mental sentences in a person's belief box are sentences in that person's natural language. (1998, p. 575)

Bound up with this assumption are two additional ones: 1) that the syntax of a mental sentence mirrors that of an English sentence (since the mental sentence *is* an English sentence) and 2) that mental sentences possess the same semantic properties as English sentences. In particular, mental sentences express propositions. Also, mental names refer to objects. This is in effect what Braun holds:

> These mental sentences express propositions because of their structures and because their constituents refer to individuals and express properties and relations. (The constituents might do this because they stand in appropriate causal or historical relations to individuals and instances of properties and relations.) (p. 574)

The conflict between the public language view assumed by Braun and the alternative view, represented by Chomsky's "I-language," is simply too broad to be considered here.[60] This dissertation takes for granted Chomsky's hypothesis according to which underlying our production and

[59] For a clear statement of the public language view, and comparison with other views on the nature of thinking, see (Devitt & Sterelny, 1999, Ch. 7).

[60] All sorts of obvious questions arise concerning the idea that thought is carried out in a public language. For example, what does "to refer" mean for a "name" of this mental language? Is it the same as for natural language names? (Braun seems to think so.) Are then mental names *directly* referential expressions? If so, Braun's account would be in danger of being viciously circular.

understanding of language is an I-language substantially *different* from the language in which we speak and write.[61] This assumption plays no part, however, in the following critical comments on the explanatory role Braun assigns to mental sentences.

The trouble with assuming that mental sentences are in a public language, e.g. English, is that the informativeness, substitution, and empty names problems would seemingly apply to these mental sentences as well. To see this, consider Braun's explanation of the fact that (1) appears informative to A but (2) does not. According to Braun, (1) appears informative to A but (2) does not because A believes the proposition expressed by them in different ways; that is, he has the different mental sentences (1m) and (2m) in his belief box.

(1m) George Orwell is Eric Blair.

(2m) Eric Blair is Eric Blair.

(1m) and (2m) include the public language (English) names "George Orwell" and "Eric Blair." Like the names in (1) and (2), they refer to the same individual. Suppose now that A thinks about (1m). Is (1m) informative to A? If it is (and it should be, since its English counterpart is

[61] See (Chomsky, 1986, §2.3; 2000, Chs. 2 and 6) for an explanation of the notion of an I-language.

informative to him), it must be because there are two other, distinct mental sentences (1m*) and (2m*) also in his belief box (or perhaps in a second belief box).

(1m*) George Orwell is Eric Blair.

(2m*) Eric Blair is Eric Blair.

Suppose A reflects on these. Is (1m*) informative to him? If it is, then this is because there are two other mental sentences (1m**) and (2m**) also in his belief box (or perhaps in a third belief box).

(1m**) George Orwell is Eric Blair.

(2m**) Eric Blair is Eric Blair.

It seems that this process could go on forever, so long as a subject may reflect on the mental sentences in his belief box.[62] (And why shouldn't this be allowed, since they are supposed to be English sentences?) Now, Braun could avoid this infinite regress scenario simply by giving up his richer characterization of a way of believing in terms of mental sentences plus a belief box. Presumably, then, he would fall back on the coarser notion of a belief state and the many-to-one relation he assumes exists between belief states and propositions.[63] But, it would

[62] By definition, a subject's "belief box" consists of her consciously *accessible* beliefs; the beliefs need not be occurrently and consciously held at any given moment, but they *could be*.

[63] Alternatively, he could adopt a different, non-public view of mental sentences. But then he'd have to explain in some detail how sentences

seem that the notion of a belief state stands in need of greater precision for it to accomplish the theoretical work Braun hopes to derive from it. Aside from the general conjecture that belief states may plausibly be distinguished on account of their differing causal roles, Braun has very little to say regarding the nature and structure of belief states. Not enough, it would seem, to serve as the basis for a comprehensive and convincing answer to the three objections presented in §2.

5. Salmon's Gricean Strategy

A different strategy for resolving the informativeness and substitution problems has been proposed by Salmon (1986, 1989).[64] Salmon's approach is based on a distinction between the information "semantically encoded" by a sentence and the information "pragmatically imparted" by an utterance of the sentence in a context. According to Salmon, the information semantically encoded by a sentence,

---

in this non-public mental language express propositions. Chomsky, for example, explicitly denies that internal linguistic representations possess the semantic properties that have historically interested philosophers, in particular, reference. See, e.g., (Chomsky, 2000, p. 42). Other theorists who work within, or accept the basic assumptions of the generativist tradition concur that internal linguistic representations are not fully propositional, and are not the type of thing that may be evaluated for truth or falsity. See, e.g., (Levinson, 2000, Introduction), (Sperber & Wilson, 1986/1995, p. 193), and (Neale, 2004a and 2004b).

[64] In (1986) and (1989) Salmon largely ignores the issue of empty names, so this section will do so as well.

the sentence's "information content" (1986, p. 13) is what the sentence literally expresses.[65] A sentence's information content is a function of the semantic contributions of the syntactic components of the sentence, together with the way that these components are arranged. Salmon calls the contribution made by a syntactic component to the sentence's information content the "information value" of the expression. (p. 15)

An expression's information value is not the same thing as its linguistic meaning. Salmon holds, for example, that the information value of a name is just the name's bearer. As an advocate of DR, Salmon maintains that a name has no meaning over and above the object that is its information value. Nor does a name contribute anything but this object (its information value) to the proposition or information content expressed by the sentence containing it. For instance, Salmon would say that the information value of "Eric Blair" is the individual Eric Blair, who is also the information value of "George Orwell." Hence on Salmon's view sentences (1) and (2) semantically encode the same information (have the same information content/express the same proposition/have the same truth conditions). Namely,

[65] Salmon employs "proposition" as another term for "information content."

both (1) and (2) encode the information that Eric Blair is self-identical.

Even though (1) and (2) semantically encode the same information, according to Salmon utterances of the two sentences may pragmatically impart different information. The same is true of belief reports as well: utterances of belief reports such as (5) and (6) (in which [1] and [2] are embedded) may also pragmatically impart propositions distinct from the information contents semantically encoded by the sentences. Salmon's solution of the informativeness and substitution puzzles thus consists of denying, against the intuitions cited in §2, that (1) and (2) and (5) and (6) express different propositions and that these propositions disagree in truth value. Strictly and literally speaking, he argues, these two pairs of sentences encode the same information content and have the same truth value. The intuitions to the contrary are due to the frequent confusion by "the general masses and most philosophers" of semantic and pragmatic information. Using the sentences "Hesperus is Phosphorus" and "Hesperus is Hesperus" as examples, Salmon explains:

> Since the form of words 'Hesperus is Phosphorus' is considerably richer in pragmatic impartations than other expressions having the same semantic information content (e.g. 'Hesperus is Hesperus'), if one is not careful one cannot help but mistake the 'that'-clause

> as referring to this somewhat richer information—information which *A* may not believe…Utterances of the locution ^*a* believes that *S*^ may even typically involve a Gricean implicature to the effect that the person referred to by a believes the information that is typically pragmatically imparted by utterances of *S*. Even so, that is not part of the literal content of the belief attribution. The general masses, and most philosophers, are not sufficiently aware of the effect that an implicature of this kind would have on ordinary usage. It is no embarrassment to the modified naïve theory [as Salmon calls his version of Direct Reference] that ordinary speakers typically deny literally true belief attributions (and other propositional-attitude attributions) when these attributions involve a 'that'-clause whose utterance typically pragmatically imparts information which the speaker recognizes not to be among the beliefs or other propositional attitudes of the subject of the attribution. In fact, it would be an embarrassment to the modified naïve theory if speakers did not do this. With widespread ignorance of the significance of the distinction between semantically encoded and pragmatically imparted information, such violation of the rules of the language is entirely to be expected. (1986, p. 85)

As can be seen in this passage, Salmon suggests that the idea of "pragmatic imparting" could be explicated in terms of the Gricean notion of an implicature.[66] In the next section, we identify some problems with this suggestion.

6. Problems with Salmon's Gricean Strategy

The first question raised by Salmon's suggestion concerns the type of implicature needed to spell out the notion of pragmatic imparting. In (1989) he suggests two

[66] See also (1989, p. 252).

possibilities, conventional implicature and generalized conversational implicature:

> It is doubtful whether… the situation I have in mind in connection with propositional-attitude attributions, qualify as cases of what Grice calls *particularized conversational implicature*…in a number of respects, these cases better fit one or the other of Grice's two contrasting notions of *generalized conversational implicature* and *conventional (non-conversational) implicature*. Surely a great many speakers may be confused by the conventional or generalized conversational implicature of a sentence into thinking that the sentence literally says (in part) what it in fact only implicates. (pp. 275-6, fn. 11, emphasis in the original)

Two imprecisions should be noted in Salmon's remarks. The first is that according to Grice, implicating is something a speaker does. Sentences do not "have" implicatures.[67] Second, even though they are not conversational implicatures, conventional implicatures are not *non-conversational* implicatures, as Salmon appears to suppose. Grice divides implicatures into two main categories: conventional and non-conventional. Non-conventional implicatures split further into the conversational and non-conversational. Conversational implicatures may then be either particularized or generalized.

As noted by Récanati (1993, p. 346, fn. 12) and Green (1998, p. 68-9), it is unlikely that Salmon intends his

[67] For helpful discussion on this point, see (Neale, 1992).

pragmatic proposal to be understood in terms of non-conversational implicature. Non-conversational implicatures are not governed by Grice's Cooperative Principle or Maxims of Conversation, nor are they triggered by the presence of particular expressions in a sentence, as in the case of conventional and generalized conversational implicatures. In fact, they may be generated by someone failing to say anything at all, by refraining to use language. For example, suppose A asks B, "Who moved my books?" and B, whom A knows would know the answer, responds by nodding her head in the direction of C. B has implicated that C moved the books but has done so without uttering a word.[68] Since Salmon is concerned with information implicated when someone says something—in particular, when someone says something of the form "N_1 = N_2" or "A believes that S"—we will assume that he is not appealing to non-conventional, non-conversational implicature.

Let us then consider the first possibility contemplated by Salmon, that a speaker who utters a sentence of the form "N_1 = N_2" or "A believes that S" conventionally implicates information that may have a different truth value from the proposition semantically encoded by the sentence. Grice

[68] This example is based on an example in (Green, 1998, p. 68-9)

characterizes conventional implicature as a type of implicature that is prompted by the linguistic meaning of certain expressions. Unlike *conversational* implicatures, conventional implicatures are not inferences derived from special features of the context together with the Cooperative Principle and Maxims of Conversation. Grice states that although the implicated information depends on the linguistic meaning of the words uttered, it is not part of "what is said" by the speaker. (1989, p. 37) That is, the implicated information is not part of the truth-conditional content expressed by the utterance. Nevertheless, by uttering a sentence containing a conventional implicature-triggering expression, a speaker manages to communicate information that goes beyond the sentence's truth-conditional content. For example, Grice would say that (13) and (14) express the same proposition and have the same truth conditions (the speaker *says*, in Grice's technical sense of "saying," the same thing by uttering them):

(13) Sam is big but agile;

(14) Sam is big and agile,

since he takes the truth-conditional contribution of "and" and "but" to be the same (truth-functional conjunction). Yet someone who utters (13) implicates that there is some

sort of contrast or contradiction between being big and being agile. This contrast is prompted by the word “but,” whose meaning suggests a contrast between the expressions preceding and following it. Salmon, then, would be proposing that ordinary names such as “Sam” and the construction “believes that” similarly trigger specific conventional implicatures. Thus, for example, a speaker who utters (15),

(15) Ralph believes that Eric Blair was writer,

expresses a true proposition (since—let us suppose—Ralph believes that George Orwell was a writer and George Orwell and Eric Blair are the same individual), but may conventionally implicate a false proposition on account of the meaning of “believes that.”

Salmon’s proposal, construed in terms of conventional implicature, presents the following three difficulties.

First, it seems unlikely that proper names and the construction “believes that” could be classified as conventional implicature-triggering expressions. For, the conventional implicatures that would be triggered by these two expressions appear to be “cancelable.” But according to Grice, conventional implicatures are non-cancelable. Grice characterizes cancelability as a property of conversational implicatures. Cancelability is presented as

a sort of test: if a putative implicature can be canceled, then it may be viewed as conversationally implicated. Grice writes:

> [A] putative conversational implicature that *p* is explicitly cancellable if, to the form of words the utterance of which putatively implicates that *p*, it is admissible to add, *but not p*, or *I do not mean to imply that p*, and it is contextually cancellable if one can find situations in which the utterance of the form of words would simply not carry the implicature. (1978/1989, p. 44)

For example, suppose a manager has been asked to give his opinion regarding the professional qualities of one of his employees, Jones, to the head of a personnel search committee at another company. The manager says "Jones has the most beautiful handwriting." Since this reply does not address the professional qualities of Jones, the manager may be taken to have implicated that those professional qualities are wanting. However, the manager may cancel this implicature by adding, "But don't get me wrong; Jones is the best analyst in the history of the company. He is extremely meticulous, sharp, and competent. It's just that his handwriting is really something to see!"

Conventional implicatures cannot be canceled in this fashion. The manager cannot say, for example, "Jones is an analyst but he is smart," without implying that analysts are not smart. To add, "I didn't mean to imply that

analysts are not smart," for example, would not succeed in defeating the implication; at most it would seem a clumsy attempt to correct a blunder.

If, as Salmon proposes, names and the expression "believes that" do carry particular information in the same way that "but" suggests an unexpected contrast, then it would seem that this information can be disavowed or canceled. Consider names first. Suppose for the sake of argument that "George Orwell" induces a conventional implication involving the description "the author of *1984*." The name conveys this information just as "but" conveys an unexpected contrast. Suppose now that A says to B, "It's time to feed George Orwell." At B's puzzled look, A says: "I did not mean to imply that it's time to feed the author of *1984*, but my dog, George Orwell." The conventional implicature hypothetically carried by "George Orwell" is thus cancelable.

Now consider "believes that." Presumably Salmon is to be interpreted as proposing that this expression conveys the information that the individual to whom the belief is ascribed believes it just as it has been reported.[69] For example, someone who utters (15) would conventionally

[69] We follow Récanati (1993, pp. 332-3) in making this presumption. Salmon nowhere specifies the exact information supposedly conveyed by "believes that," nor does he offer examples of conventional implicatures arising from reports containing this expression.

implicate that Ralph has the belief that Eric Blair is a writer and that he has this belief precisely as it has been articulated in (15)—i.e. including the name "Eric Blair." The trouble with this idea is that, as Barwise and Perry point out (1981, p. 394), belief reports such as (15) do not imply, but at most suggest, that the believer would endorse the report as stated.[70] It is perfectly appropriate to say (adapting Barwise and Perry's example [p. 394] to the present case): "Ralph believes that Eric Blair was a writer, but only knows to call him 'George Orwell,'" for instance. In short, the suggestion that the believer would endorse the report as stated is cancelable.

The second difficulty with Salmon's proposal is a difficulty given Salmon's Directly Referential view of names. Suppose once again that the name "George Orwell" conveys the information expressed as "the author of *1984*." Whenever the name is used, it conveys this information, just as "but" always conveys an unexpected contrast. The problem for Salmon is then that this information would have to be considered as part of the name's meaning—since that is how Grice defines "conventional implicature"—but according to Salmon, names have *no* (descriptive) meaning.

[70] Ducrot (1984, Ch. 8) argues that neither should it be automatically assumed that the report reflects the point of view of the reporter, even when she is speaking seriously.

In other words, the hypothesis that a name may prompt a certain conventional implicature seems incompatible with a DR view of names.

Finally, it should be mentioned that the very notion of conventional implicature is open to question, and has been rejected by several well-known theorists.[71] If these theorists are correct, Salmon would be appealing to an empty notion.

Let us look now at the second possibility, generalized conversational implicature. According to Grice, a generalized conversational implicature is a conversational implicature that may be inferred without appealing to special features of the context. (p. 37) Like a conventional implicature, a generalized conversational implicature is triggered by the presence of a particular expression in the sentence uttered. (p. 37) Grice says that the implicature may "normally" (p. 37) be derived—it may be inferred in almost all contexts—whenever someone makes an utterance containing the expression. Unlike a conventional implicature, however, a generalized conversational implicature is cancelable. (p. 39)

As illustration, consider the following sentences:

[71] E.g., Récanati (1993), Bach (1999), Levinson (2000), and Carston (2002). Neale (2001) examines the problems surrounding the notion of conventional implicature, but does not reject the notion outright.

(16) She married John and she got pregnant.

(17) She got pregnant and she married John.

Grice would say that (16) and (17) have the same truth conditions and express the same proposition (a speaker who utters [16] and [17] "says" the same thing). As noted above, Grice holds that the semantic value of "and" is simply the familiar truth-functional one. But "and" may trigger different implicatures, which would be automatically picked up by the interlocutors regardless of the specific context. For (16) the implicature would be that the woman first married John *and then* got pregnant; i.e. that the two events happened in a certain temporal order. In the case of (17) "and" triggers two distinct implicatures: 1) the temporal one, to the effect that the woman first became pregnant and subsequently married John, and also 2) a causal implicature, to the effect that she married John *because* she got pregnant. Yet these implicatures are cancelable; one can say, for instance, "She got pregnant and married John, but not in that particular order."

Understood in terms of the concept of a generalized conversational implicature, Salmon's notion of pragmatic imparting would seem to escape the first objection mentioned above to a construal in terms of conventional

implicature. Consider names first. Suppose once more that "George Orwell" conveys certain information, e.g. "the author of *1984*." The presence of "George Orwell" in a sentence would thus trigger an implicature involving this information, in the same way that "and" may trigger an implicature involving the information that the events described in the clauses preceding and following it occurred in a certain temporal order. This time, however, the implicature triggered by "George Orwell" is a generalized conversational implicature. It may be canceled. Thus, for example, A may say, "George Orwell sheds a lot—but I don't mean that the author of *1984* sheds a lot; it's my dog George Orwell that sheds a lot." Similarly, the putative implicature triggered by "believes that" would be cancelable. Assume once again that "believes that" conveys the information that the person to whom a belief is ascribed has the belief in the way it has been ascribed. To defeat the implicature carried by (15), someone may then utter the sentence we adapted from (Barwise & Perry, 1981): "Ralph believes that Eric Blair was a writer, but only knows to call him 'George Orwell.'"

But does this alternative construal avoid the second objection mentioned above? It would seem not. Generalized conversational implicatures, while not part of the

implicature-triggering expression's conventional (i.e. linguistic) meaning, are characterized by Grice as being nonetheless closely associated with it; they are "non-detachable" from it. Because of this relationship of non-detachability, "…it will [generally] not be possible to find another way of saying the same thing, which simply lacks the implicature in question…" (p. 39). Grice goes on to say that: "one may expect a generalized conversational implicature that is carried by a familiar, non-special locution to have a high degree of non-detachability." (p. 39)

Salmon's proposal that names give rise to generalized conversational implicatures would thus have to presuppose that names have some sort of meaning[72] from which the implicature may not be detached. But the only sort of "meaning" Salmon allows names to have is their bearers; the only "information value" a name possesses, according to him, is an object.

[72] Recall that generalized conversational implicatures, like all implicatures, are supposed to be derived on the basis of what is said, and that Grice defines "what is said" as "closely related"—one could say identical, with the exception of indexicals—to the linguistic meaning and syntactic structure of the sentence uttered. The expression Grice offers as an example of a generalized conversational implicature-triggering expression—"a/an X"—has linguistic meaning, a fact that is recognized by Grice's discussion in p. 38. So presumably it is this meaning that goes into triggering the generalized conversational implicatures carried by uses of this expression. The problem is that, as discussed above, proper names for Salmon do not have linguistic meaning; they only have "information values"—objects.

The main problem, however, with understanding pragmatic imparting in terms of generalized conversational implicature has to do with the notion of a generalized conversational implicature itself. As in the case of conventional implicature, various theorists have expressed doubts about the cogency of this notion.[73] They argue that the information conveyed by a generalized conversational implicature-triggering expression can be shown to constitute part of the truth-conditional content expressed by the sentence containing it, contrary to what Grice holds. If the information conveyed by the expression affects the truth conditions of a statement that contains it—i.e. if it is part of what is said—then it can't be part of what is implicated.

For instance, if, as Grice believes, (16) and (17) really express the same proposition and have the same truth conditions (a speaker says the same thing by uttering them), then so should (18) and (19), which result from embedding (16) and (17) in a larger conditional sentence.[74]

---

[73] E.g., Récanati (1993) and Carston (2002).

[74] Due originally to Cohen (1971), this type of "scope argument" aims to show, for any expression which according to a Gricean analysis triggers a specific kind of implicature, that the (supposedly pragmatic, implicated) information associated with the expression is actually part what is said, i.e. part of the truth-conditional content expressed by the sentence containing it, by placing this sentence within the scope of logical operators such as "if…then," "or," and "not." Many writers have availed themselves of this argument to question Gricean analyses of specific expressions or more generally the mainstream idea that an

(18) If she married John and she got pregnant, then grandma will be happy.

(19) If she got pregnant and she married John, then grandma will be happy.

However, the intuition is that (18) and (19) express quite different propositions and do *not* have the same truth conditions. It appears, rather, that the temporal (in [18]) and causal (in [19]) relations associated with "and" are indeed relevant to the truth-conditional content of the antecedents of these conditionals. Therefore, the Gricean analysis, which construes these relations as something that goes beyond what is said, and not as part of the truth-conditional contribution of "and," must be mistaken.

Some of these critics[75] of the Gricean notion of a generalized conversational implicature view the temporal and causal information that uses of "and" may convey as pragmatic (not linguistically mandated, but understood in a given context) and at the same time truth-conditional (part of what is said, not of an implicature). They view the non-truth-functional information conveyed by other generalized conversational implicature-triggering

---

utterance's truth-conditional content ("what is said" by the utterance) is determined by its meaning and syntax; see, e.g. (Anscombre & Ducrot, 1983, Ch. 2), (Sperber & Wilson, 1986/1995), (Carston, 1988, 2002), (Récanati, 1993, Appendix to Part II), (Bezuidenhout, 1997a), and (Green, 1998).

[75] E.g. Carston (2002, Ch. 3) and Récanati (1993, Appendix to Part II).

expressions, such as “some” (“some but not all”), “a/an NP” (“a/an NP not closely associated to the speaker”), and “if” (“if and only if”) in the same way: as pragmatic but truth-conditional. Others[76] regard the temporal and causal connotations of “and” as aspects of the word’s linguistic meaning and hence as a straightforwardly semantic matter. If these critics are correct, then the concept of a generalized conversational implicature is of little aid in fleshing out Salmon’s notion of pragmatic imparting.

What about *particularized* conversational implicatures? Perhaps Salmon’s proposal could be made sense of in terms of this remaining type of implicature. Here we shall not explore this alternative; concurring with Récanati (1993, p. 344) and Green (1998, p. 75), we take it that Schiffer (1987) has shown it to be untenable.[77]

7. Soames’s Non-Gricean Pragmatic Strategy

In (2002, 2004, 2007), Soames has proposed a strategy for handling the three objections that is similar to Salmon’s in that it is based on a distinction between the information “semantically expressed” by a sentence and the

[76] E.g. Bar-Lev and Palacas (1980).
[77] Salmon himself is skeptical of the idea, as seen in the passage quoted at the beginning of this section.

information "pragmatically asserted"[78] by an utterance of the sentence in a context. Unlike Salmon's account, however, Soames's does not appeal to Grice's notion of implicature.

According to Soames, the information semantically expressed by a sentence S, or the "semantic content" of S, is the information that would be expressed in all normal contexts in which S is used by competent speakers with its literal meaning (without irony, sarcasm, or defeating conversational implicatures), and once all ambiguous and indexical elements in S have been disambiguated and assigned referents. (2002, pp. 105–6 and 2004, p. 360) However, in many cases, argues Soames, a speaker who utters S asserts more than just the semantic content of S:

> [N]ormally the semantic content of the sentence uttered, relative to the context of utterance, is *one* of the propositions asserted by an agent's assertive utterance. (2002, p. 131, emphasis added)
>
> [A]n assertive utterance of a sentence … by an agent in a context c often results in the assertion not only of the semantic content of s with respect to c, but also of other, sometimes stronger, propositions as well." (2002, p. 133).

[78] Soames's somewhat peculiar use of "to assert" and "assertion" will be interpreted as referring not to the performance of a speech act, but to the content of the speech act: a proposition evaluable for truth or falsity. In alternative but equivalent vocabulary, Soames's "assertion" or "asserted proposition" is *what is said* by the utterance of a sentence in a context.

Soames notes that the semantic content of S might be quite poor; indeed, S might even fail to express a complete proposition:

> Semantic contents of grammatically complete sentences (relative to contexts) are not always complete propositions; sometimes, they are incomplete propositional matrices, together with partial constraints on how to contextual information may be used to complete them. (2004, p. 357)

Soames argues that in actual contexts of utterance the proposition or propositional matrix expressed by S is "pragmatically enriched" with various sorts of information.[79] As a result, one and the same sentence may be used to assert different propositions from one context to the next, even though the sentence will express one and the same semantic content in all contexts.[80] According to Soames, it is the diversity of information that utterances typically convey in ordinary communicative situations that is the source of our divergent intuitions about the infomativeness, substitution and empty names cases.

Soames presents various examples to illustrate how sentences having the same semantic content may in different contexts be used to assert different, pragmatically

[79] Bach (1994b, 1994c, 2006) has defended a very similar view. What Soames calls a "propositional matrix" Bach calls a "propositional radical."
[80] Cf. (Neale, 2001).

enriched, propositions.[81] Here we will go over just one of them, which appears in all three works ("The Party," 2002, pp. 75–7; 2004, pp. 360–5; 2007, pp. 7–9), and which specifically concerns the question of the informativeness of sentences containing different co-referential names.

The setting is the Princeton Philosophy Department party held at the beginning of the academic year. A professor, Paul, asks a new graduate student, Mary, "Have you been introduced to Peter Hempel?" Mary then asks, "Who is Peter Hempel?" Paul responds by gesturing in Carl Hempel's direction and saying:

(20) Peter Hempel is Carl Hempel.

Later, Mary reports to another student that Paul said (asserted) the following:

(21) The man, Peter Hempel, who is standing over there, is the famous philosopher of science Carl Hempel.

Soames argues that Mary's report is correct, in virtue of the following facts and assumptions manifest in the conversation between Paul and Mary: (a) Mary could see that Paul was gesturing at a man and referring to him as "Peter Hempel;" (b) Paul and Mary both associate the name "Carl Hempel" with the information that its bearer is a famous philosopher of science; (c) Paul knows that Mary

[81] See, e.g., (2002, pp. 74–86).

knows that Carl Hempel is a famous philosopher of science (and Paul knows that Mary knows that he knows this); and (d) that Mary will, as a result, associate that specific bit of descriptive information with the name.

Now, such descriptive information is not part of the semantic content of the name. For Soames, a sentence of the form "$N_1 = N_1$" and a true sentence of the form "$N_1 = N_2$" express the same proposition; they have the same semantic content. Thus (20) has the same semantic content as (22),

(22) Carl Hempel is Carl Hempel.

There is no unnoticed semantic fact that would create a difference between (20) and (22). However, the information conveyed by them in different conversational situations can vary widely. According to Soames, it is this information we are sensitive to when we are asked whether (20) is more informative than (22) or have the same meaning.

In claiming that an utterance pragmatically conveys information distinct from the information semantically expressed by the sentence, Soames is not arguing that the richer information is implicated by the utterance of the sentence. Unlike Salmon, Soames does not propose that names or propositional attitude verbs such as "to say" give rise, by virtue of their meaning, to specific conventional or generalized conversational implicatures that may diverge

in truth value from the truth value of the proposition literally and semantically expressed by the sentence containing them. Nor is the richer information conveyed as a *particularized* conversational implicature: Soames does not construe "pragmatically asserted" information as a separate proposition derived through an application of Grice's Cooperative Principle and Maxims of Conversation to the utterance of a sentence expressing a different proposition. Rather, it is the *very* proposition or propositional matrix, the semantic content, of the uttered sentence that is enriched in the context:

> [T]he proposition asserted arises from the semantic content of the sentence uttered by adding contextually determined pragmatic content to one of its constituents. (2007, p. 12)

Soames, unlike Salmon, claims that *what is said*, the proposition semantically expressed by a sentence, can be pragmatically enriched in a context, which is something that a Gricean cannot admit.

8. Problems with Soames's Non-Gricean Pragmatic Strategy

Soames's strategy appears more promising than either of the proposals of Braun and Salmon. This section considers two objections to Soames's view of assertion, in particular to his idea that a sentence may be used in a context to

assert propositions “stronger” than the semantic content expressed by the sentence.

The first objection, due to McKinsey (2005), looks at Soames’s examples purporting to show that a sentence may be used to assert a proposition(s) stronger than its semantic content, and denies that this is in fact what occurs. With respect to the example reproduced above, McKinsey says the following:

> Soames’s claim [that in addition to the proposition semantically expressed by the sentence uttered, the speaker would have asserted and said something else] is false. Had Carl Hempel not really been a philosopher of science, or had the man standing over there at whom Paul gestured not really been Peter Hempel (perhaps he was wearing a Peter Hempel mask), Soames’s claim implies that Paul would have said something false in uttering the true sentence [20]. But this is wrong. Given the context, Paul may have *conveyed* or *implicated* something false by uttering [20], since he may have conveyed or implicated the (hypothetically) false information expressed by [21]. But even so, Paul would have neither asserted nor said anything false in uttering the true sentence [20]. (2005, p. 156, emphasis in the original)

It would appear, however, that Soames’s claim can be defended. Suppose that McKinsey’s counterfactual hypothesis is right: suppose that Carl Hempel was never a philosopher of science and that he had cleverly fooled people into thinking that he was all of these years. Suppose further that Paul was the only person in on the secret. (Perhaps he wrote the works published under

Hempel's name.) Let everything else about the party situation be as Soames described it. In particular, Paul knows that Mary thinks that Carl Hempel is a famous philosopher of science, and that she knows that Paul knows that she believes this. Then, if Paul utters (20) in this situation, does he say something true or false?

It would seem that only the strictest interpreter—one who, moreover, would already be persuaded by DR—would judge that Paul said something true, simply in virtue of the alleged fact[82] that (20) literally expresses the necessarily true proposition that Carl Hempel is self-identical. Most unprejudiced people would probably say that Paul deliberately misled Mary, if not flat-out lied to her. Suppose Mary later found out the truth and confronted Paul about the false information he "asserted" at the party. It would be a sorry defense indeed if Paul were to claim in response that all he literally said was that Carl Hempel is self-identical.

But McKinsey and many others are likely to remain unconvinced. For there are alternative interpretations of the party example that would appear to be better grounded than Soames's. McKinsey suggests, for instance, that the example illustrates the pervasive phenomenon that Bach

[82] Frege (1892a/1997a), for example, would dispute this "fact."

(1987/1994a, pp. 77–85) has labeled "standardization."[83] As Bach has amply documented, many expressions (words, phrases, whole sentences) are frequently (perhaps most frequently) used non-literally. For example, phrases of the form "can you…?" are standardly used to perform requests, even if the request is literally expressed only by the VP that follows the construction (e.g. "pass the salt"). McKinsey proposes (p. 158, fn. 6) that identity sentences containing names, such as (20), are standardly used non-literally, to assert "partly quotational" propositions, such as *Peter Hempel is [named] "Carl Hempel."*

Standardization cannot be disputed, but intuitively it would seem that the two cases, the case of requests performed by utterances of "can you VP?" and that of identity sentences containing names, differ. But this is only an impression. Against McKinsey's "standardization argument" it may be replied that much more needs to be ascertained regarding the literal and non-literal uses of names in order to evaluate the merits of the argument. Without further empirical evidence, McKinsey's argument risks begging the question, since, presumably, McKinsey's reason for concluding that identity sentences containing

[83] See also (Bach, 1998).

names are standardly used not-literally to express metalinguistic information is the belief that the *literal* use of a name is to refer to an object. But this is precisely the idea that DR advocates need to prove and that is being questioned in this and the following chapter.

McKinsey questions Soames's notion of "pragmatic enrichment." And indeed, one may ask, how exactly is it, according to Soames, that a proposition may be pragmatically enriched in a context? After all, (20) contains no expressions traditionally regarded as context-sensitive. How, for example, does the information that Carl Hempel is a famous philosopher of science wind up in the proposition Paul supposedly asserted in uttering (20)?

This is the second and main problem presented by Soames's pragmatic strategy: Soames provides no detailed account of how pragmatic enrichment is supposed to work. Soames himself admits that he has no detailed account to offer (2007, p. 6), and only glosses the matter over as follows:

> In many contexts, the semantic content of S—whether it is a complete proposition or not—interacts with an expanded conception of pragmatics to generate a *pragmatically enriched proposition* that it is the speaker's primary intention to assert. (2007, p. 5, emphasis in the original)

There are available, however, a number of "expanded conceptions of pragmatics" ("expanded," meaning,

presumably, “non-Gricean”) which possess the theoretical resources to explain the phenomenon of pragmatic enrichment Soames has illustrated; one such conception is the Relevance-Theoretic conception that has been adopted in this dissertation. Until a detailed and principled account of pragmatic enrichment is provided, critics of Soames’s view of assertion and his account of the examples are likely either to reject the view and the account as insufficiently motivated or to seek a more traditional explanation of the examples.

In sum, the main problem with Soames’s pragmatic strategy is not the idea that the truth-conditional content of sentences (in this case, sentences containing names) may be pragmatically enriched with contextual information, or that a sentence may be used to assert propositions other and stronger than the information “semantically expressed” by the sentence, but that Soames fails to spell out the crucial notion of pragmatic enrichment. In the next chapter, an account will be presented of what is said by the utterance of an identity sentence containing names that in many respects agrees with Soames’s. In particular, it is based on the idea, which is vital to Soames’s strategy, that propositions in general may be pragmatically enriched with contextual information, in ways that are not traceable

to indexical or ambiguous elements in a sentence. An important difference, however, is that the notion of pragmatic enrichment is characterized more fully, by drawing on Relevance-Theoretic concepts.

Chapter 3

The Truth-Conditional Values of Proper Names

1. Introduction

The mainstream view on the semantics of proper names is that they have objects as truth-conditional values; let us call this the Referentialist view. The standard Referentialist analysis of sentences containing names may be illustrated via this simple example. Consider (1).

(1) Aristotle is wise.

The first step in the analysis is to observe that the sentence seems to allow a natural partition: (1) may be split into a subject, composed of an NP ("Aristotle"), and a predicate, composed of a VP ("is wise"). Next, the Referentialist assumption is made that the NP has an object, Aristotle, as truth-conditional value. That is, Aristotle, the man, is the contribution the name makes to the truth conditions of the sentence. Finally, it is assumed that the VP predicates something, a property (or, alternatively, a set), of Aristotle—*being wise*. Thus, (1) is true just in case Aristotle has the property of being wise. That is, the conditions for the truth of (1) are that there be an Aristotle and that he be wise. In propositional parlance, one might say that the thought or proposition expressed by (1) is about a certain individual,

the “referent” of “Aristotle,” viz. Aristotle, and that this thought is true if in fact Aristotle is wise.

In the previous chapter, we were concerned with the most extreme version of Referentialism, Direct Reference. In this chapter, our target is Referentialism in its general form. The chapter seeks to challenge Referentialism by establishing the following three claims in order:

(A) Proper names encode descriptive information;

(B) Proper names may be literally interpreted as having the descriptive information they encode as their truth-conditional value (instead of, or in addition to, an object);

(C) The type of truth-conditional value a proper name may be interpreted as having in a given occasion of use is pragmatically and contextually determined by relevance-guided processes of interpretation.[84]

The discussion proceeds as follows. §2 presents linguistic evidence that would support (A). §3 further explains the dialectical role of some of the evidence presented in §2. §4 gives reasons for concluding (B). §5 sketches a Relevance-Theoretic explanation of names’ varying truth-conditional values. §6 characterizes the

[84] See Chapter 1, §6, for an overview of key Relevance-Theoretic concepts.

relationship between thesis (A) and Descriptivism, which, along with Direct Reference, is one of the two main philosophical approaches to the semantics of names.

Some of the consequences of this discussion can be stated at the outset: If (A) and (B) are true, then Direct Reference, as defined in the previous chapter, is false. If (B) is true, then the standard semantic view according to which names are referring expressions (= expressions having objects as truth-conditional values), purely in virtue of their meaning and syntactic type, is false. (B) and (C) together undermine the very idea of semantic reference.

2. Evidence for Descriptionality

Some reasons for thinking that names encode descriptive information (thesis [A]) are the following. (The reasons go from least to most compelling.)

First, as Horwich (1998, pp. 124–5) notes, there is the intuitive observation that we commonly talk of *understanding* or failing to understand a name, and of *knowing* or being ignorant of the "meaning" of a name. Such locutions are perfectly ordinary. Furthermore, knowing or failing to know the "meaning" of a name would seem to be a matter distinct from knowing or failing to know *who* (or

what) the bearer of a name is. For instance, one may be able to accurately specify some of the information encoded by the name "Gareth Evans," and have a good idea of how the name functions in the language, without knowing who bears the name.

Second, many names can be translated from one language to another, just like other types of expression (e.g. "Aristotle" is "*Aristóteles*" in Spanish).[85, 86] Even empty names are capable of being translated (e.g. the English "Santa Claus" is "*Papa Noël*" in French). Intuitively, as Horwich suggests (1998, p. 124), what is thought of as translated from one language to another is an expression's linguistic meaning.

Third, there is syntactic evidence. In (1973) Burge pointed out that names may be "modified." In particular, Burge noticed that (a) like nouns, names can be pluralized, e.g.,

---

[85] By "can be" is meant here that a name is translatable in principle; in practice, of course, many names are not translated, but are borrowed from one language into another. Names of well-known people or places are often translated: e.g., the Spanish "*Aristóteles*" names both the Greek philosopher and the late Greek shipping magnate ("*Aristóteles Onassis*"). However, Abbott (2005, p. 12) notes that as a rule, "borrowed" names conform to the phonotactic constraints of the borrowing language. So from a phonotactic point of view names are almost always "translated." For instance, "Hsieh Ho" (to use Ziff's example [1960, p. 85]) is pronounced in English according to the phonotactic rules of English, not Chinese.

[86] Admittedly, it is a bit strange to say that "*Aristóteles*" is the *translation* in Spanish of "Aristotle," but neither is it a transliteration, of course. Perhaps, following a suggestion by Neale (2004a), it is better to say that this is the name's *rendering* in Spanish.

(2) There are many John Smiths in London;

(b) like nouns, they can combine with the definite and indefinite articles,[87] e.g.,

(3) A John Smith joined the department;

(4) The John Smith who joined the department is an old friend;

and (c), like nouns, they can combine with determiner-quantifiers,[88] e.g.,

(5) Some Johns are clever; some are dull.

To Burge's four examples, the following should be added:

(6) That John is funny (combines with the distal demonstrative);

(7) This John is a bore (combines with the proximal demonstrative);

(8) Our John is talented (combines with possessive determiners);

(9) Which Bush did you mean? [George Hebert or George Walker?] (combines with an interrogative determiner);

(10) The new Bob likes coffee (combines with adjectives).[89]

---

[87] In other languages (e.g. Romance languages, German, and Modern Greek) names combine with the definite article much more frequently than in English.

[88] In von Fintel's (1994, p. 2) terminology, determiner-quantifiers are determiners that have traditionally been called quantifiers; they combine with noun phrases. Examples of determiner-quantifiers include "all," "every," "each," "some," and "many."

Syntactically, the name in the preceding examples behaves just as a noun, a type of expression standardly taken to have a property or a set as truth-conditional value. (To see this comparison more clearly, substitute "John" in the examples with a noun that would harmonize with the rest of the components of the sentences, such as "medievalist.") Of course, according to current syntactic theory, names are NPs.

The resemblance of names to nouns can be characterized further. In (2)-(10) "John" seems to behave as a count noun. Count nouns are nouns that can combine with numerals, the definite and indefinite articles, determiner-quantifiers, and have both a singular and a plural form. They apply to things that may be counted. In contrast, *mass nouns* apply to things that are not counted, such as water. Mass nouns present characteristics opposed to those of count nouns: they often do not combine with the definite and indefinite articles; they do not usually pluralize by adding "s," and they combine with only certain determiner-quantifiers, such as "some" and "much." Quine (1960, p. 91) observed that many count nouns can "double"

---

[89] Other examples are *Ugly Betty* (the title of the TV show), and adjuncts such as "junior" and "senior" (e.g. "Jr." in "John F. Kennedy, Jr." conveys that John F. Kennedy, Jr. is younger than John F. Kennedy, Sr.).

as mass nouns. For example, the count noun “rabbit” has a mass reading in (11),

(11) John served rabbit and garlic mashed potatoes for dinner.

The same appears to be the case with names. The following examples are from (Allers, 2006, p. 12).

(12) Hannibal Lector served Jason and garlic mashed potatoes for dinner.

(13) I can only handle so much Jason at this hour in the morning.

(14) Jason in the morning is bad for your health.

Another piece of syntactic evidence that suggests that names encode descriptive information is that they commonly take noun suffixes such as “ic,” “ist,” “ian/an,” and “esque” (e.g. “Platonic,” “Marxist,” “Cartesian,” “Kafkaesque”), which serve to adjectivize a name. Adjectives, like nouns, are taken to have properties or sets as truth-conditional values.

Fourth, names encode cultural and linguistic information. The name “John,” for example, is an English name. Just like any other English word, it conveys the information that it is an English word. Moreover, the name suggests that its bearer is an English speaker and that this individual is in some way or other (perhaps through genetic

heritage or assimilation) connected to Anglo-Saxon culture. That names carry such cultural and linguistic information is evidenced by the common occurrence of questions seeking to confirm it. Two typical situations in which such questions arise are, for instance, when one is traveling in foreign countries (e.g. "Your name is 'John'…Are you American?") or when one is introduced to someone with a foreign name (e.g. "'François'…Is that French?" "Are you French?" "Do you speak French?"). Evidently there are many individuals who do not conform to the expectations raised by the cultural and linguistic connotations associated with his or her name. However, this does not negate the fact that names encode such information; if anything, such cases would confirm this fact by presenting an unexpected contrast.

Fifth, names encode natural gender. An expression that possesses natural gender (as opposed to "grammatical" gender[90]) encodes the information that the things to which the expression applies are sexed. "John," for instance, encodes the information that its bearer is male, whereas "Jane" encodes the information that its bearer is female.

[90] Grammatical gender is a morphological feature of words of certain languages. A word's having grammatical gender is unrelated to the question whether the word applies to sexed things.

In the great majority of cases, the gender of name and bearer agree.

To some, these last two points may seem trivial; it might be thought that it is a purely adventitious feature of natural language names that they encode cultural/linguistic information and natural gender. As we shall see in §4, it is far from trivial that names do this. For these two types of descriptive information may constitute the contribution of a name to the truth-conditional content of an utterance containing it. That is, the name will be interpreted to have general, descriptive information as its truth-conditional value, instead of an object, contrary to what the standard semantic picture of names implies.

Summing up, the reasons for thinking that names encode various sorts of descriptive information (thesis [A]) are the following:

1. As with other expressions, it is perfectly natural to say that one knows or doesn't know the meaning of a name.

2. Names may be translated or "rendered."

3. Names behave syntactically like nouns (syntactically, they are NPs), which are expressions that encode descriptive information and whose truth-conditional values are properties or sets.

4. Names encode cultural and linguistic information.

5. Names possess natural gender.

3. What the Syntax of Names Implies about their Semantics

The purpose of this section is to state in greater detail the implications of the third reason given above, that names behave syntactically like nouns. The examples presented above make it clear that names do in fact behave syntactically like nouns. Sentences (2)-(14) are not just well formed, they are also completely commonplace.

From a formal semantic point of view, nouns have properties or sets as truth-conditional values.[91] Thus examples (2)-(14) would warrant the following argument: Names sometimes are nouns. Nouns have sets as truth-conditional values. Therefore names sometimes have sets as truth-conditional values. The problem is that according to standard truth-conditional semantics, names have objects as truth-conditional values, not sets. Following Larson and Segal (1995), a semantic axiom schema that gives the truth-conditional value[92] of a name "N" may be put as follows:

(O) Val (x, *N*) iff x = N,

where x is an object. If "N" is instantiated by "John," for example, this yields the following axiom:

Val (x, *John*) iff x = John.

---

[91] The following discussion is in terms of sets rather than properties. Nothing hangs on this, though; the points made here can be expressed in terms of properties as well.

[92] Though Larson and Segal, in formulating a Davidsonian truth-theoretic account, would not necessarily call this a *truth-conditional* value.

But if the argument just given is sound, the truth-conditional value of a name may also be specified by the axiom schema for nouns. Again using Larson and Segal's notation, the schema for a noun "P" would be represented thus:

(P) Val (X, *P*) iff X = {x : x is a P},

where X is a set. For example, the axiom for "book" is the following:

Val (X, *book*) iff X = {x : x is a book}.

So, the axiom for "John," treated as a noun, is:

Val (X, *John*) iff X = {x : x is a John}.

So which of these two axiom schemata, (O) or (P), correctly gives the semantics of names? (O), it might be argued, adequately characterizes the semantics of singular "unmodified" names in subject position, as in (15):

(15) John is tall.

Yet (P) seems required by the syntactic behavior of names in sentences like (2)-(14).

To our knowledge, there are two main ways of reconciling the examples of noun-like behavior of names with the mainstream view of their semantics captured by axiom (O). The two approaches, which are similar, aim to provide a unified semantic theory of names. The first is Burge's (1973) view that the singular, unmodified, subject-position

use of names is to be explained by positing a phonologically null demonstrative element associated to the name. For example, (15) would have the deeper syntactic form "THAT John is tall," where THAT refers to some contextually salient John. The second is Larson and Segal's (1995) account according to which it is not a demonstrative that is covert in the name's structure but the definite determiner "the." For example, (15) would have the deeper syntactic form "THE John is tall," where THE singles out a contextually salient John. Larson and Segal's approach has recently been developed further by Elbourne (2005).

It is beyond the scope of this chapter to examine these theories in detail, or to mention their flaws. Here we do not need to endorse either theory or indeed any semantic theory of names at all. The aim of this chapter is not to explain from a semantic point of view how it is that names can have sets as truth-conditional values. (Our claim, remember, is that the truth-conditional value of a name is the value it is interpreted as having by interlocutors tacitly following relevance-guided principles of interpretation. That is, the truth-conditional value the name may be interpreted as having is a pragmatic and not a semantic matter.) Nevertheless, we shall take examples

(2)-(14) at face value. The idea that the syntactic similarity with nouns is illusory or philosophically unimportant is to be rejected. Taken at face value, then, the data culled in the previous section seem sufficient to justify claim (A), which had been our goal.

4. Non-Standard Truth-Conditional Values of Names

The following two points support the conclusion that names may be literally interpreted as having the descriptive rather than objectual truth-conditional values (thesis [B]).

First, there is the syntactic argument discussed above, which bears repeating: names sometimes are (syntactically) nouns; nouns have properties or sets as truth-conditional values; therefore names sometimes have properties or sets as truth-conditional values. Independently of how to make sense of the first premise of the argument from a semantic point of view, the fact remains that this argument provides a strong prima facie ground for thinking that names may have truth-conditional values other than objects. Syntactic evidence is always to be taken seriously.

Second, as Martinich (1977), Devitt (1981, §5.6), and Bezuidenhout (1997b, p. 385, fn. 6) have suggested, names, like definite descriptions and other expressions, admit of

the referential/attributive distinction.[93, 94, 95] An attributively interpreted expression contributes descriptive information to the truth-conditional content of an utterance that contains it. The utterance expresses a "general" or "object-independent" proposition.[96]

So, when interpreted attributively, a name would contribute descriptive information—not an object—to the proposition expressed by the utterance that contains it. An utterance containing an attributively interpreted name would thus (be taken to) express an object-independent

---

[93] The referential/attributive distinction will be characterized more fully in Chapter 5, §5.

[94] Martinich (1977) presents the following example of the attributive use of a name:

> Suppose that a drawing has just taken place for the grand prize in some raffle for charity and the chairman of the raffle committee announces the winner by saying: "Jane Smith has won the grand prize" …[S]uppose that the chairmain does not know anything more about Jane Smith…the chairman is using "Jane Smith" to state that whoever is Jane Smith won the grand prize. It does not matter whether Jane Smith completed the winning ticket, knew of the entry in her name or would have approved of it…Jane Smith, whoever she is, has won the grand prize. (1977, p. 161)

Martinich, however, expresses doubts as to the clarity and import of the notion of "attributive use" as characterized by Donnellan in (Donnellan, 1966).

[95] Devitt, however, considers "attributive names" to be a different "type of name" from "referential names."

[96] Following Neale (1990, pp. 49–50, fn. 1; 2008 forthcoming, p. 139), by an object-dependent proposition we mean a proposition (thought, truth-conditional content) whose existence (whether conceived abstractly or as a mental entity, we take no position on this) essentially depends upon the existence of a particular object. Put in terms of utterances and truth conditions, an utterance has object-dependent truth conditions if they include a particular object. By an object-*in*dependent or general proposition we then understand a purely qualitative or descriptive proposition, a proposition whose existence is not dependent on the existence of any particular object; and correspondingly, an utterance has object-independent truth conditions if no specific individual is included among these truth conditions.

proposition. If it is true that names have an attributive interpretation, this would be sufficient to prove thesis (B).

The following examples show how names may be interpreted attributively.

First example, demonstrating the existence of the attributive interpretation. Suppose Albert is checking the messages recorded on his answering machine. He listens to a message left by someone who identifies himself as "John," and who says is calling to invite Albert to dinner that night. However, Albert is unable to determine who John is. Even so, Albert later explains to his wife that he might not be home for dinner since:

(16) John, whoever he is, invited me to dinner.[97]

In uttering (16), Albert is not using "John" referentially, since he does not know who John is—i.e. since he has no particular individual in mind—and as the clause "whoever he is" suggests. Albert's use of "John" would be interpreted as expressing certain general information, and his utterance of (16) a general proposition, whose truth

[97] Some might object that "John" here is really "short for" a longer, perhaps quotational expression such as "A person called 'John,'" for example. See §3 of the next chapter for a reply to this objection.

conditions may be partially represented in first-order logic as follows:[98]

∃x (Jx & Ix) [Something or other has the property of being a John[99] and the property of having invited me to dinner.]

Second example, illustrating that the natural gender information encoded in a name may constitute the name's truth-conditional contribution. Suppose now Albert returns home from his dinner with John. His wife notices that the collar of his shirt is stained with the imprint of a woman's lips in bright red lipstick. Recognizing which particular brand of lipstick it is, she says:

(17) Funny that John should prefer Laura Mercier [the lipstick's brand].

As in (16), the name "John" in (17) seemingly does not denote a specific individual, but rather descriptive information. And (17) would likewise be taken to express a

[98] Normally a natural language name would be rendered in logical notation as a constant. However, this would not be accurate for an attributively interpreted name, which is interpreted as predicating certain properties of a random object. This random object is better represented by a variable.

[99] As indicated in Chapter 1, this dissertation is open and neutral as to what sorts of things may be considered properties. From our perspective, the property of *being a John* is a perfectly legitimate property, and on a par with the property (also encoded by "John") of *being male*, for example.

general proposition, whose truth conditions may be partially represented in first-order logic as follows:

<<∃x (Jx & Mx & Px)>, F> [Someone or other has the properties of being a John, of being male, and of preferring Laura Mercier lipstick, and this proposition is funny.]

Note that, as in (16), it would have been natural to insert in (17), following "John," the phrase "whoever he is," which serves as a test for the attributive use. Note, too, that the various implicatures that Albert's wife may be making in this context by uttering (17)—e.g. *you* [Albert] *did not have dinner with a man called 'John'* or *you* [Albert] *lied to me*—depend on the gender information encoded by "John."

Third example, illustrating that the cultural/linguistic information encoded in a name may be interpreted as constituting the name's truth-conditional contribution. Suppose two students are looking for help in translating Baudelaire's "L'Invitation au Voyage." They go to the school's Language Center, where pasted on the door is a list of tutors' names. Seeing the name "François" on the list, one student says to the other:

(18) Judging by the name, I think that François, whoever he is, can probably help us out.

What are the truth conditions of this utterance? The speaker has the belief that a tutor named "François" should be able to help them translate the poem. And she is explicitly appealing to the cultural/linguistic information encoded by the name to support this belief. Clearly in this case she isn't using "François" referentially, since the students have no idea who François is. (18) would seem to be true, then, just in case the speaker believes that the tutor, whoever he is, lives up to the information encoded by François, i.e. possesses one of certain properties—*being French* or *being a French speaker*—and, *because* he possesses such a property, would thus be able to help them.

Once again, we may partially represent the seemingly object-independent truth conditions of (18) in logical notation as follows:

<BEL{A, <∃x (Fx & Gx & Hx)>}> [A (the student) believes (stands in the binary belief relation BEL to the proposition that) someone or other has the property of

being a François, the property of being French, and the property of being able to help.][100, 101]

Now, it might be objected that while it is true that names encode gender and cultural/linguistic information, it is too strong to claim that this information is truth-conditional.[102] After all, one can make a true utterance containing "François," where the name is interpreted attributively, and this utterance would not be falsified if "François" turned out not to be French. For example, an utterance of (19) (in the situation described above) would appear to be true, even if François turned out to be German and not French.

(19) François, whoever he is, works at the Language Center.

Hence, if *being French* (or *being a French speaker*) were a property with truth-conditional import, (19) could not be true if François were German, but it appears that (19) *can* be true if he is German.

---

[100] As suggested in §2, natural language names convey the information that their bearers belong to the linguistic and cultural group with which the name is associated. While this information may of course turn out to be false (e.g. François may not be French), this does not negate the fact that the *name* conveys this information.

[101] We are employing here a standard relational analysis of belief purely for convenience's sake, without a commitment to the adequacy of such an analysis.

[102] This objection was made by Neale (conversation).

This objection might be correct as far as (19) goes, but it doesn't apply to (18). If evaluated in the circumstances in which it was uttered, which were described above, (18) would indeed be falsified if François were German. The reason the speaker is proposing that François could help them translate the poem is because the name "François" suggests that its bearer is French or knows French. That is why she *thinks* that François would be able to help them. Remember, she is going by the name alone; there would have been no point to her suggestion if she had had any reason to believe that François were not in fact French (or a speaker of French) but German, for example.

In other words, the objection fails to engage the claim that was actually made above. Our claim is not that "François" (or any name, for that matter), when interpreted attributively, *must* have as its truth-conditional contribution the cultural/linguistic information it encodes, but that the name *may* have this information as its truth-conditional contribution. (Note that "François" encodes natural gender just like "he" does, but this information is truth-conditionally inert in the example.) The example tries to bring this out, without presuming that such descriptive information is always, in all contexts, truth-conditional. The aim of course is to refute the

Referentialist who claims that the only truth-conditional contribution a name can have is an object.

The attributive interpretation of names, as illustrated by these examples, seems literal. From an intuitive standpoint, uncorroborated by any empirical research, the attributive interpretation would appear rather frequent, or at least not rare.[103] Be that as it may, no great weight will be placed on the notion of "literal" here.[104] Of more interest to us is the following: that the examples of the attributive interpretation seem to show clearly that the descriptive information encoded by a name may enter into the truth-conditional content of the sentence containing the name. In other words, thesis (B) is shown correct.

Recapitulating the section, thesis (B) is warranted on the following grounds:

1. Names behave syntactically like nouns (syntactically, they are NPs), which are expressions that encode descriptive information and whose truth-conditional values are properties or sets. (This is reason 3 of §2.)

2. Names admit of the referential/attributive distinction. Attributively interpreted expressions contribute general or descriptive information, not objects, to the truth conditions of the utterances that contain them.

---

[103] Frequency, however, is not a sufficient condition for literality, as shown by the phenomenon labeled "standardization" by Bach (1987/1994a, 1998), which was mentioned in the previous chapter.

[104] The attributive interpretation appears literal by the lights of the notion of "literal interpretation" defined in footnote 1 in Chapter 1.

5. A Relevance-Theoretic Explanation of Names' Varying Truth-Conditional Values

The preceding sections offered support for theses (A) and (B) by citing positive evidence for them. This section concerns thesis (C). Our aim is to show that names' varying truth-conditional values can be straightforwardly explained through Relevance-Theoretic concepts. The following example illustrates how a name may have an object as truth-conditional value. Suppose Andy and Brady are talking about the food preferences of a mutual friend, Chris. Andy utters (20),

(20) Chris hates anchovies.

Brady interprets Andy's utterance of (20) as expressing the proposition *Chris strongly dislikes anchovies*. The proposition understood by Brady concerns a specific individual, the Chris known to Andy and Brady. This proposition is an "object-dependent" proposition. Brady would say that Andy's utterance is true just in case the Chris she knows hates anchovies.

Brady's interpretation process may be outlined as follows. According to the Communicative Principle of Relevance,[105] Andy's utterance should be optimally relevant to Brady, that is, it must have positive cognitive effects

[105] See Chapter 1, §6.3.

for Brady while at the same time requiring little processing effort. Certain information is available from the context and is manifest to both Andy and Brady, for instance:

(a) Andy and Brady have been talking about their mutual friend Chris.

Following an interpretive path of least effort, from contextual fact (a) Brady interprets "Chris" to refer to the Chris known to them, the person they've been talking about. Brady's knowledge of the linguistic meanings of "to hate" and "anchovies" allows her to determine the rest of the proposition.[106]

Such an interpretation of Andy's utterance has positive cognitive effects for Brady. For example, it leads her to correct her mistaken impression that Chris loves pizza topped with anchovies.

Let us now look at the more controversial case, that of a name having a descriptive rather than an objectual truth-conditional value. In the previous section we saw that a name can have a descriptive truth-conditional value, and not an objectual one, when it is interpreted attributively. It is not immediately clear how the standard Referentialist

[106] As noted in Chapter 1, Sperber and Wilson conceive of the comprehension process as occurring quickly and largely unconsciously.

account of the semantics of names would explain the attributive interpretation, which, it was surmised above, seems literal.

This is how Relevance Theory would explain the attributive interpretation of “François” in the third example presented in §4. The example describes the following situation. Two students, A and B, go to their school’s Language Center to look for a tutor to help them translate a French poem. Seeing the name “François” on the Center’s list, A says to B:

(21) François [whoever he is] can help us.

The following contextual facts are manifest in the situation:

(a) A and B are talking to each other.

(b) A and B know that “François” is a French name, which carries the information that its bearer is French/speaks French.

(c) The students do not know who the François who works at the Language Center is.

Following a path of least effort, this is how B’s relevance-guided interpretation process would go:

1. B interprets “us” to refer to A and B. [From contextual fact (a).]

2. B interprets “help” to mean “to assist.” [From her knowledge of the linguistic meaning of “to help.”]

3. B interprets “can” to mean “is able to.” [From her knowledge of the linguistic meaning of “can.”]

4. B interprets “François” attributively, to mean something like “someone or other who is French and a François.” [From the information encoded by “François,” information which (b) says she is able to decode.]

Thus the information encoded in “François” allows B to interpret A’s utterance as expressing the proposition that *someone who is French and a François is capable of assisting them [A and B]*. The proposition she recovers is a general and object-independent one. In §4, it was represented as follows:

∃x (Fx & Gx & Hx) [Someone or other has the property of being a François, the property of being French, and the property of being able to help.]

The variable x of course reflects the fact that the object possessing these three properties is not known and is arbitrarily posited.

That “François” is interpreted as denoting (in Russell’s sense) a random someone possessing certain properties and

not a particular person known to A and B is compelled by contextual fact (c); that is, (c) blocks the referential interpretation of "François." In this situation, it would be completely unwarranted for B to think that A is using "François" of a particular individual known to her.

The reasoning employed in these two examples may also be applied to Soames's "Carl Hempel" case, which was discussed in the previous chapter.[107] In Soames's example, Mary, an entering graduate student, and Paul, a professor, are conversing during a Philosophy Department party at Princeton. Paul asks Mary if she's been introduced to Peter Hempel and points to a man standing nearby. When Mary replies that she hasn't, and doesn't know who Peter Hempel is, Paul says:

(22) Peter Hempel is Carl Hempel.

Mary later reports to another student that Paul asserted the following:

(23) The man, Peter Hempel, who is standing over there, is the famous philosopher of science Carl Hempel.

In the previous chapter we expressed our agreement with Soames's view that Paul's utterance of (22) conveyed to Mary the information that she later articulated explicitly as (23). We also concurred with Soames's assessment that

[107] Chapter 2, § 7.

this information was not conveyed as a conversational implicature, but resulted from the enrichment of the linguistic components of (22) with contextual material. We faulted Soames, however, for not providing a detailed account of such "pragmatic enrichment."

Here's how RT would describe Mary's comprehension process of Paul's utterance of (22). In the situation described by Soames, the following unstated facts and assumptions are available to Mary and Paul:

(a) Mary sees that Paul is gesturing at a man and referring to him as "Peter Hempel;"

(b) Mary and Paul both associate the name "Carl Hempel" with the information that its bearer is a famous philosopher of science.

Mary interprets "Peter Hempel" to refer to the man to whom Paul is pointing. This interpretation is licensed by contextual fact (a); it is the easiest and most obvious interpretation in the situation. On the basis of contextual assumption (b), she takes "Carl Hempel" to mean "the famous philosopher of science." She then interprets "is" to signify identity. So Mary understands Paul's utterance as expressing the proposition that the man standing close to them, at whom Paul is gesturing, is

identical to the person satisfying the description they both associate with the name “Carl Hempel.”

This interpretation meets the “least processing effort” requirement of the Communicative Principle of Relevance. But does it have positive cognitive effects for Mary? Since Soames doesn’t formulate the example in these terms, the question really is moot. However, with little trouble we may build on Soames’s example to suggest how such an interpretation would have cognitive effects for Mary. Suppose Mary was under the impression that Hempel had retired from all social and academic activities at Princeton. Then Paul’s utterance, interpreted as above, would lead her to revise this belief. It would improve her representation of reality by removing a mistaken assumption.

The Relevance-Theoretic account just sketched is predicated on two fundamental assumptions: 1) that Sperber and Wilson’s Cognitive Principle of Relevance is true of human cognition and 2) that expressions may only be ascribed truth-conditional values when the sentences that contain them are used in contexts. As indicated at the end of Chapter 1, this dissertation does not offer a defense of the main concepts and assumptions of RT; they are taken for granted here. As for the second assumption, general

reasons were given in Chapter 1 (see §§3–5) in support of it.

From a Relevance-Theoretic point of view, the meaning and syntactic type of a proper name constrains, but does not determine, its truth-conditional value. Truth-conditional values are not mechanically spit out by an expression's meaning and syntax. That is to say, they are not encoded by language. What determines a name's truth-conditional value is the relevance-guided comprehension process. The name's contribution has "to make sense" in relation to the goal of producing an interpretation of the speaker's utterance that enhances the hearer's knowledge of the world.

In Chapter 5, Relevance-Theoretic notions will be appealed to again to explain the different truth-conditional values that phrases of the form "every F" may make. We will defer a consideration of alternatives to the RT explanation presented here until that discussion.

6. Thesis (A) and Descriptivism

In §2 it was argued that names encode descriptive information (thesis [A]). What is the connection between this claim and the traditional Descriptivist theories of names? Let us first see what traditional Descriptivism is.

Descriptivism is one of two main philosophical approaches to the semantics of names; the other is Direct Reference. The approach originates in the work of Frege, who argued in (1892a/1997a) that names, in addition to having referents, describe or present them as being a certain way. This additional function of a name he called the "sense" of a name. A name's sense, he claimed, serves as a "mode of presentation" of its referent. His view was that the sense of a name determines or uniquely identifies its referent.[108] According to Frege (1892a/1997a, p. 153, fn. B), senses vary from speaker to speaker.[109] He gives the following example: one speaker may associate the sense "the teacher of Alexander the Great" with the name "Aristotle" while another might associate the different sense "the pupil of Plato" with the name. As in this example, Frege represented the senses of names using definite descriptions (expressions of the syntactic form "the F"). Frege called an expression that "has as its referent (*Bedeutung*) a definite object" a "proper name." (p. 153) In this

[108] Dummett (1973, p. 110) suggests that the sense of an expression for Frege offers a "criterion" by which the expression's referent may be uniquely identified.

[109] In (1914/1997, p. 321), however, Frege says that "a common store of thought, a common science would be impossible" if sense were in this way subjective. Sense objectivity is necessary for the transmission of knowledge. Thanks to Neale (conversation) for pointing this out.

category he included not only ordinary names, but also definite descriptions and whole sentences.[110]

Descriptivism is also associated with Russell (1910, 1919), who, like Frege, believed that the information connected to a name could be expressed as a definite description. However, he went further than Frege in claiming that most names are "disguised" or "abbreviated" definite descriptions.[111] Russell also agreed with Frege that a name is associated with different definite descriptions for different people or even for the same speaker at different times. (1910, p. 114) Unlike Frege, however, Russell did not think that definite descriptions denote[112] "definite objects." In (1905), Russell argued that definite descriptions are to be analyzed as complex

---

[110] This is not to say that for Frege names are semantically equivalent to definite descriptions. As Dummett (1973, p. 110) points out, it is doubtful that the thesis that the sense of a name is always the same as that of a definite description may be attributed to Frege, even as tacitly assumed by him (since, according to Dummett, Frege nowhere explicitly states such a view). But, as Dummett (p. 110) also admits, whenever Frege gave examples of the sense of a name, he articulated this sense by means of a definite description. This is all that is being claimed above.

[111] The phrase "most names" is meant to acknowledge that Russell held that there are "logically proper names" that are non-descriptive and solely refer to objects. However, the examples he gave of logically proper names are not names at all, but the demonstratives "this" and "that:" "We may even go so far as to say that, in all such knowledge as can be expressed in words—with the exception of 'this' and 'that' and a few other words of which the meaning varies on different occasions—no names, in the strict sense, occur, but what seem like names are really descriptions." (1919, p. 174).

[112] Russell calls definite descriptions "denoting expressions" and the object satisfying the properties expressed by a definite description the description's "denotation." Throughout the dissertation, this object is being called the "satisfier" of the description.

quantificational structures. The surface grammatical form of a definite description might give the impression that it has an object as truth-conditional value. For example, "the current President of the U.S." might be taken to have the current President of the U.S., George W. Bush, as its value. However, Russell's analysis purports to show that definite descriptions have a deep "logical form" representable in logical notation by quantifiers, predicate-letters, variables, and logical connectives.[113] This structure does not contain an element that would single out any particular individual.[114] Consequently, names, insofar as they are abbreviated definite descriptions, are not really referring expressions on Russell's account. Their truth-conditional values are not individuals, but whatever set is expressed by the NP they abbreviate, together with the second-order set expressed by "the."[115]

Russell's view that names abbreviate definite descriptions was modified in an important way by Wittgenstein (1953, §79) and Searle (1958). These philosophers argued that a name is associated not with *one*

[113] Strictly speaking it is the sentence containing the definite description that has such a "logical form."
[114] Russell's analysis of descriptions is briefly summarized in §4 of Chapter 5.
[115] Russell is being interpreted here as claiming that "the" expresses the (higher order) set that the set F is uniquely instantiated.

definite description, but with a cluster of them. To qualify as the name's satisfier,[116] an object must satisfy an indeterminate number ("a weighted logical sum or inclusive disjunction," according to Searle [1958, p. 160]) of the definite descriptions associated with the name.

A newly popular form of Descriptivism, which may be labeled "Metalinguistic Descriptivism," claims that the meaning of a name consists of a definite description that mentions the very name. Thus, a name "N" is taken to mean "the bearer of 'N'" or "the individual called 'N.'" Metalinguistic Descriptivism was first seriously advocated by Kneale (1962, p. 630), though Russell (1919, p. 171) also alludes in passing to the possibility of interpreting names in this way. Authors who have more recently argued for or explored the view include Loar (1976), Bach (1987/1994a, Ch. 7; 2002), Katz (1990, 2001), Récanati (1993, Chs. 8–9), Geurts (1997), and Justice (2001).

Historically, Descriptivist theories have been motivated by the informativeness, substitution, and empty names/empty names-in-existence-sentences problems discussed in the previous chapter. Frege's Descriptivism provides relatively straightforward solutions to these problems.

[116] As on Russell's view, names on the cluster view abbreviate a set of definite descriptions; hence a name may be *satisfied by*, but does not refer to, an object.

The informativeness problem is solved by noting that identity statements composed of different but co-referential names may appear informative to someone who associates different senses or definite descriptions with the two names. Similarly, the substitution problem is explained if it is assumed that what falls under the scope of the propositional attitude verb is the name's sense and not its referent. For then there is no substitution of identicals for identicals, but of two non-identical things, two distinct senses. It is to be expected that the truth value of the sentence will change after a substitution of one sense for another. And the empty names problem disappears if we consider that while a name may lack a referent, it still has a sense. In virtue of the name's sense, the sentence containing the name may be considered meaningful and truth-evaluable.

The claim that names encode descriptive information, thesis (A), may be viewed as a species of Descriptivism about names, insofar as it affirms, as all traditional Descriptivist theories claim and contrary to what Direct Reference holds, that names do express certain information in addition to having the role of standing for the individuals that bear the names. That is, (A) agrees with this core insight of Descriptivism: that the semantics of

names includes a level of information that is distinct from the individuals that may bear the names.

However, (A) is to be considered Descriptivist only to that limited extent. In particular, (A) should not be interpreted as being equivalent to, or as presupposing or entailing, other claims that may characterize traditional forms of Descriptivism, such as the following, for example:

Descriptivism

D1: Names are/are semantically equivalent to definite descriptions.

D2: Names have "senses" and also referents.

D3: The sense of a name determines its referent.

Names encode descriptive information, but in most cases (e.g. "John," "François") this information seems to be of a very general and schematic sort. For our purposes, it is not necessary to express this information in the form of a definite description. (It is debatable even whether the descriptive information encoded by most names is adequately captured via a definite description, as compared to, for example, an *indefinite* description.) The view we have defended in this chapter should therefore not be assimilated to traditional Descriptivism.

Chapter 4

Are Proper Names Rigid Designators?

1. Introduction

No, or so this chapter argues. In particular, it is argued that proper names are not rigid, if rigidity is taken to be a semantic property of name types. The argument against rigidity is simple and may be put as a dilemma: rigidity is either a property of name types or name tokens;[117] if understood as a property of types, then there arises a well-known problem, the multiple bearers problem, and the responses to it (it will be shown) are unsatisfactory; if taken as a property of tokens, rigidity becomes uninteresting from a semantic point of view, since rigidity would then have to be viewed as a (pragmatic) property of a *use* of a name by a speaker in a context.

The popular view that names are rigid designators needs to be challenged, since this view does not square with our claim, articulated in the previous chapter, that the truth-conditional value of a name is determined (= identified) pragmatically and in context, and not semantically, that is, simply in virtue of a name's linguistic meaning[118] and syntactic type. For if reference (= having an object as

[117] The type/token distinction is explained in footnote 13 of Chapter 1.
[118] As defined in §3 of Chapter 1.

truth-conditional value) is not a semantic matter, then neither can rigidity (= constant, linguistically mandated reference to one and the same object across possible worlds) be a semantic matter. Semantic rigidity implies semantic reference.[119]

Our aim in this chapter is not to offer new reasons against the semantic notion of rigidity, but to present in a streamlined fashion those that already exist (most of them are extracted from [Katz, 2001] and [Bach, 1987/1994; 2002]). The next section reviews Kripke's (1970/1980) concept of rigidity. The dilemma is then laid out in §3.

2. Rigidity

Under the influence of Kripke's forceful argumentation in (1970/1980),[120] many philosophers have come to believe that natural language proper names are "rigid designators." This has been considered an "empirical discovery," and the empirical status of the rigidity thesis has long been recognized. (See, e.g., (Evans, 1973), (Lockwood, 1975), (Peacocke, 1975), and (Stanley, 1997, p. 555).

[119] The converse is not true, of course: reference does not imply rigidity.
[120] See especially the Preface, pp. 48-9, and pp. 76-8.

Rigidity is taken to be an inherent, semantic property of names ("rigidity de jure").[121, 122] That names are rigid de jure, or as a matter of the "semantic rules of the language," is indeed the thesis that Kripke aimed to defend in (1970/1980): "Clearly my thesis about names is that they are rigid de jure…" (p. 21, fn. 21)

A rigid designator is defined by Kripke as an expression that refers to one and the same object in every possible world where that object exists. (Kripke, 1970/1980, p. 48) By a "possible world" is meant an actual or counterfactual situation which concerns that individual in some stipulated way or other. (p. 15) "Names are always rigid designators," Kripke states (p. 58), but other expressions are not, notably most definite descriptions, expressions that had been taken by earlier philosophers, in particular Frege (1892/1997), Russell (1910, 1919), Wittgenstein (1953, §79), and Searle (1958), to give the meaning, or to be synonymous with, proper names.[123] Most[124] definite

[121] For helpful discussion on this point, see, (Récanati, 1993, §§1.2, 1.3).
[122] The contrast is with "rigidity de facto": an expression is rigid de facto if it just happens to be the case, metaphysically, that the expression cannot but denote a particular object. One of Kripke's examples is the definite description "the smallest prime." This expression is held to be rigid de facto, since, metaphysically speaking, no other entity but the actual smallest prime—the number two—can be the smallest prime number and thus satisfy that description.
[123] See §6 of the previous chapter for a brief discussion of traditional Descriptivism.

descriptions, seemingly unlike names, may be satisfied by different individuals in different possible worlds.

One of Kripke's examples is the following. Kripke says that the definite description "the President of the U.S. in 1970" is non-rigid, since in the actual world it is (was) satisfied by Nixon, but in another possible world it would be satisfied by someone else. For instance, in the possible world in which Humphrey defeated Nixon in the 1968 elections, the description is satisfied by Humphrey. Since what is (was) actually true of Nixon may not have been so—the President of the U.S. in 1970 might have been someone else—"Nixon" and "the President of the U.S. in 1970" cannot be synonymous.

"Nixon," on the other hand, is rigid, since "no one other than Nixon might have been Nixon." (p. 48)

That is, sentence (1) below seems true (since it would be true in a counterfactual situation), but (2) seems false (since it appears to be false in every possible world):

(1) The President of the U.S. in 1970 might not have been Nixon.

(2) Nixon might not have been Nixon.

---

[124] As suggested in footnote 122, Kripke allows that some definite descriptions, like "the smallest prime" and "the positive square root of four," are rigid. According to him, though, these are rigid de facto and not rigid de jure. It is a metaphysical and not a linguistic fact that supposedly makes them rigid.

Thus, one piece of evidence in favor of the thesis that names are rigid comes from modal discourse: names appear to behave differently from definite descriptions in modal environments—i.e. in sentences containing moral operators such as "might have been."

However, the main evidence for rigidity is supposed to be furnished by our intuitions: in (1970/1980) Kripke presents rigidity as a notion that is justifiable and graspable intuitively.[125] In the Preface, for instance, he argues that our intuitions corroborate that names are rigid in "simple" sentences, too—i.e. in sentences lacking modal operators such as "might have been."

He presents the following example. Suppose, he says, that "Aristotle" means "the last great philosopher of antiquity." Then consider the following pair of sentences:

(3) Aristotle was fond of dogs.

(4) The last great philosopher of antiquity was fond of dogs.

Regarding (3) and (4), Kripke says:

> [3] and [4] are 'simple' sentences…Yet the issue of rigidity makes sense as applied to both. My view is that 'Aristotle' in [3] is rigid, but 'the last great philosopher of antiquity' in [4] is not. No

[125] For instance, he characterizes (p. 48) the rigidity thesis as one of the "intuitive theses" he wants to establish in (1970/1980); he says that he will argue for the thesis "intuitively" (p. 49); and he describes the example given above concerning "Nixon" as an "intuitive test" for rigidity. (p. 48)

> hypothesis about scope conventions expresses this view; it is a doctrine about the truth conditions, with respect to counterfactual situations, of (the propositions expressed by) *all* sentences, including *simple* sentences. (1970/1980, pp. 11–12, emphasis in the original)

To support his view on the rigidity of "Aristotle" and the non-rigidity of "the last great philosopher of antiquity," Kripke appeals to our intuitions concerning (3) and (4). Specifically, "the intuition is about the truth conditions, in counterfactual situations of (the proposition expressed by) a *simple* sentence." (1970/1980, p. 12, emphasis in the original) If we compare (3) and (4) with respect to the actual world and other possible worlds, Kripke argues, it would seem that the truth values of (3) and (4) are the same in the actual world, but diverge in other possible worlds. (3) is true in the actual world or in other possible worlds just in case there is a particular individual, namely Aristotle, who was fond of dogs. In contrast, (4) is true, and is so for the same reasons as (3) in the actual world, but not in other possible worlds. There is no one individual such that (4) is true in every possible world just in case that individual is fond of dogs. For instance, in the counterfactual situation in which Aristotle died in his infancy, (4) would seem to be true just in case *Plato*, not Aristotle, was fond of dogs.

In this case “the last great philosopher of antiquity” would be satisfied by Plato, not Aristotle. Imagine now the counterfactual situation in which Alexander, instead of becoming a conqueror, opted to follow in his teacher’s footsteps and become a philosopher; and, what’s more, he succeeded in equaling or surpassing Aristotle in terms of the quality of his work. In such a situation, (4) would be true just in case Alexander, not Aristotle, was fond of dogs. In this case “the last great philosopher of antiquity” would be satisfied by Alexander, not Aristotle. Thus, according to Kripke, the intuitive difference in the actual and counterfactual truth conditions of (3) and (4) reveals that names such as “Aristotle” are rigid, whereas definite descriptions such as “the last great philosopher of antiquity” are non-rigid.

In sum, the rigidity thesis ultimately rests on our intuitions regarding the truth conditions of any and all sentences containing names, as Kripke insists throughout the Preface and the first two lectures of (1970/1980). An important consequence of the rigidity thesis, pointed out by Kripke, is that Names behave differently from definite descriptions in modal environments; the satisfier of a

definite description varies in modal environments whereas a name always refers to the individual bearing the name.[126]

3. Problems with Rigidity

Kripke's repeated assertions that rigidity can be justified and understood intuitively are problematic, but the topic of philosophical intuitions is a delicate one and is too broad to be addressed here.[127] In this section we lay out the dilemma mentioned at the beginning: rigidity, if it exists, must be a property of either name types or name tokens, but there are problems with both alternatives. Let's look at what these problems are, starting with the first possibility, rigidity as a property of name types. Kripke's remarks would suggest that this first interpretation of rigidity is the appropriate one, since, recall, Kripke explicitly says that he is arguing for rigidity de jure, or rigidity as a matter of the semantic rules of the language.[128] His claim is about "language as we have it" (p. 12), and not about particular uses of language by particular speakers in particular contexts.

---

[126] I am grateful to Neale for clarifying to me that this is a *consequence* of the rigidity thesis, and not an argument for it.
[127] It might be argued, for example, that what Chomsky says of "denote," "refer," and "true of" seems apt in the case of "rigidity" as well: "there can be no intuitions about these [technical] notions, just as there can be none about 'angular velocity' and 'protein.'" (2000, p. 130)
[128] On p. 21, quoted above.

(Or so it would seem at first sight; the discussion below suggests that another interpretation is possible.) Understood as a claim about name types, then, Kripke's rigidity thesis is that ordinary English name types are rigid designators.

One reason to doubt that names, considered abstractly as types, are rigid, is that typically names have more than one bearer. Take the name "Aristotle," for example. Many different individuals have been and currently are called "Aristotle." The name, considered in the abstract of a context—simply as an item in the language, synchronically construed—does not pick any one of them out. The problem, in other words, is this: if "Aristotle," abstractly considered, does not refer to any one of its bearers to the exclusion of the others, then neither can it *rigidly* refer to one of its bearers to the exclusion of the others. This is the familiar multiple bearers problem.

Kripke deals with this problem by denying that names may have more than one bearer. Instead, he proposes that names may be individuated by bearer: "[D]istinctness of referents will be a sufficient condition for the uniqueness

of the names." (1970/1980, p. 8, fn. 9) Kripke further explains his proposal as follows:[129]

> Some have thought that the simple fact that two people can have the same name refutes the rigidity thesis. It is true that…I spoke for simplicity as if each name had a unique bearer. I do not in fact think, as far as the issue of rigidity is concerned, that this is a major oversimplification…For language as we have it, we could speak of names as having a unique referent if we adopted a terminology, analogous to the practice of calling homonyms distinct "words," according to which uses of phonetically the same sounds to name distinct objects count as distinct names. This terminology certainly does not agree with the most common usage, but I think it may have a great deal to recommend it for theoretical purposes. (1970/1980, pp. 7–8)

The analogy Kripke is suggesting here is the following. Just as one might consider "bank" (= financial institution) to be a different word from "bank" (= side of a river)—instead of one word with two meanings—so might one consider "Aristotle" (referring to Aristotle, the philosopher) to be a different word from "Aristotle" (referring to Aristotle Onassis, the late shipping magnate), for example. The two names "Aristotle" (call them "Aristotle_1" and "Aristotle_2"), just like the two nouns "bank" ("bank_1" and "bank_2"), sound the same, are spelled the same, and belong to the same

[129] Notice that the reason given by Kripke for the rigidity of "Nixon"—viz. "no one other than Nixon might have been Nixon" (p. 48)—needs to be supplemented with the (tacit) premise—"and 'Nixon' has no other bearer than Richard Nixon, the U.S. president"—in order for the conclusion to go through; otherwise it seems a non-sequitur: the (pretty uncontroversial) metaphysical fact that no other entity than Nixon could have been Nixon seems irrelevant to whether an English expression has a certain posited property (rigidity) or not.

syntactic category, but are better classified as two distinct words.[130]

However, there are various problems with this proposal. On an intuitive level, Katz (2001, pp. 150–2) identified the following implausible consequences of Kripke's "homonymy" view: (a) natural-sounding statements such as "The fifty John Smiths living in New York are Yankees fans" and "I know one John Smith, but there must be more in America," would be "semantically deviant" (in Katz's words) and seemingly unanalyzable; (b) if a speaker named "John" said to someone else also named "John," "We have the same name," then what the speaker asserted would be literally false; (c) by the same token, if one John were to say to the other, "We have different names," he'd be saying the literal truth; (d) the term "namesake" would be rendered meaningless; (e) it would be redundant for a son named after his father to append "Jr." to his name; (f) there would be no family names, in the sense that the members of the Smith family are said to have the same name "Smith;" and (g) if, for example, Brenda Starr married Kenneth Starr, then, literally speaking, she would change her last name by taking his.

[130] Kaplan (1990) has also advocated this one bearer/one name view of names, which he calls the "common currency" view, and explicitly draws the analogy with "bank" to justify it.

In addition, Katz suggested that Kripke's proposal appears based on a misunderstanding of homonymy. As Katz explained, true homonyms, such as "bear" and "bare," "differ in etymology, (typically) orthography, syntax, and meaning." In contrast, "Kripke's distinct names, such as 'Smith' used to name one person and 'Smith' used to name another, do not differ in etymology, orthography, syntax, or meaning (since both names 'Smith,' being names, are supposed to have zero meaning[131])." (2001, pp. 148-9)

To Katz's criticisms, the following should be added: other expressions (e.g. indexicals such as pronouns and demonstratives) may have multiple referents too, but in their case there is no inclination to individuate expression by referent. We wouldn't want to say, for example, that there are as many (homonymous) pronouns "I" as there are individuals referred to by this word (potentially the totality of English speakers). Rather, it seems more plausible to say that there is a single English expression, the personal pronoun "I," whose tokens may

[131] Katz seems here to be attributing a Direct Reference view of names to Kripke. There is some textual evidence for such an attribution: e.g., in (1979) Kripke states that "the linguistic function of a proper name is completely exhausted by the fact that it names its bearer" (p. 240), which was identified as claim DR1 in Chapter 2, and in (1971) he says, "the function of a name is simply to refer, and not to describe the objects so named by such properties as 'being the inventor of bifocals.'" (p. 57) However, we have it on extra-textual grounds that Kripke's view is not as close to Direct Reference as these quotes would make it seem.

refer to different individuals in different situations. Why then should names be any different?

Well, a Kripkean (though perhaps not Kripke) might reply, names are different because in their case there is no "character rule[132]," no statable linguistic meaning, whereas in the case of indexicals there is. That is to say, the one bearer/one expression approach is unmotivated in the case of indexicals, since indexicals possess a linguistic meaning which may readily be distinguished from the things tokens of the expressions may be used to refer to. "I," for example, has the context-invariant linguistic meaning "the person currently speaking or writing," but proper names appear to lack meanings expressible in this manner. What would be the linguistic meaning of "Aristotle," for instance?[133]

But this possible Kripkean reply may be countered, and this rejoinder can again be put as a dilemma. Names (types) either have or don't have linguistic meanings that are non-identical with the individuals that may bear the names. If taken at its word, the Kripkean reply assumes

---

[132] In Kaplan's (1989) sense.

[133] Of course, in (1970/1980) Kripke offers a suite of arguments against the idea that names are synonymous with definite descriptions specifying properties possessed by the bearers of the names. In this chapter we are not concerned with Kripke's many-sided critique of traditional Descriptivism, but only with his rigidity thesis, a positive and stand-alone claim about proper names.

that names don't have linguistic meanings, so that it might fairly be interpreted as involving a commitment to a Directly Referential view of names.[134] But it is known that Kripke himself does not explicitly endorse Direct Reference, despite the attributions of some (e.g. Katz, quoted above, see footnote 131) to the contrary. And, regardless, Direct Reference is not a satisfactory theory, for reasons given in Chapters 2 and 3. So, pending some clarification as to Kripke's actual views on the linguistic meaning of names, the reply, interpreted this way, must be seen as unsatisfactory.

On the other hand, if it is allowed that names do have some sort of linguistic meaning over and above the individuals that may bear them (and, to escape Kripke's objections to traditional Descriptivism, they would have to be meanings not synonymous with one or a cluster of definite descriptions specifying the properties possessed by the bearer of the name), then the parallel with indexicals becomes an irresistible reason to reject the one bearer/one name view. So the reply is unsuccessful in this case as well.

Indeed, the parallel with indexicals has in recent years led a number of philosophers to view proper names as a type

[134] Direct Reference was defined in §2 of Chapter 2.

of indexical with a semantics similar to that of indexicals: e.g. Récanati (1993, Chs. 8–9), Pelczar and Rainsbury (1998), and Carston (2002, pp. 37–8). These philosophers argue that names, qua expression types, have linguistic meanings of a metalinguistic sort (i.e. a meaning mentioning the name itself), and that this meaning helps interlocutors identify the referent of a token of a name used in a particular context, just as the linguistic meaning or character of an indexical expression is supposed to.[135] So, if this "indexical" view of names is correct, then the alleged difference between names and indexicals vanishes altogether: both types of expression would have meanings distinguishable from the individuals of whom we may use the names.[136, 137]

To be sure, Kripke does admit that the one bearer/one name view he proposes "does not agree with the most common usage," yet he justifies the view on account of certain

---

[135] See also (Bach, 2002).

[136] In (1970/1980) Kripke also objects to the idea that names possess metalinguistic meanings; and the authors mentioned above have elaborated different responses to these objections. It is unnecessary to weigh in on the merits of these arguments here: the remarks made above should not be viewed as an endorsement of the indexical view; rather, we are merely citing the parallel with indexicals as a reason to contest Kripke's one bearer/one name view of names.

[137] As Carston (2008 forthcoming, fn. 16) notes, Szabó (2006, p. 36) has shown that proper names pass all the tests that Cappelen and Lepore (2005) have devised to determine whether an expression is context-sensitive and must be assigned a referent in context, based on factors such as the speaker's intention to refer to a particular individual by her use of the expression, for example.

"theoretical purposes." (p. 8) Kripke does not specify which theoretical purposes he has in mind, but presumably they are related to his project in modal logic. The problem with such a presumed justification, however, is that the usefulness of a device or notion for logic does not license us to postulate it as a feature of a natural language. While individual constants in a modal system may perhaps be fruitfully regarded as "rigid," this does not constitute a sufficient ground for holding that ordinary names are rigid. (This point is also stressed by both Katz [2001, p. 149] and Bach [2002, pp. 89-90]).

But let us take a different tack. Let us *accept* Kripke's proposal; let us grant him that names may be individuated by bearer. Every name, then, is "unique;" and there are as many names as there are bearers. Even so, rigidity would not be plausibly seen as a property of a unique name type.[138] For as we saw in the previous chapter, names (tokens) may not only be interpreted referentially, they also admit of a literal attributive interpretation.

Consider the following example. Suppose two students are trying to assemble the Philosophy Department's team for an all-male inter-departmental soccer league. They are

[138] The one bearer/one name strategy does not erase the type/token distinction, of course. In the sentence "Aristotle$_1$ liked dogs, but dogs didn't like Aristotle$_1$," for example, there are two tokens of the "unique" name "Aristotle$_1$," which is individuated by one of its bearers.

missing one player. Checking the department roster, A says to B:

(5) Adam Smith, whoever he is, could play for us.

A and B, both Kripkeans, interpret "Adam Smith" as having a tacit subscript, i.e. as "Adam Smith_n." As Kripkeans, they would deny that the student has the "same" name as the famous British political economist, for example. But in this case they are not interpreting the name referentially, since neither of them knows this person, Adam Smith_n. They have no particular student in mind, a fact that is signaled by the inclusion of the clause "whoever he is," which serves as a test for the attributive interpretation. They are interpreting the name attributively, as expressing certain information encoded by it, natural gender. (5) is true, they would say, if and only if this person, whoever he is, is male. Only then could he play for them.

In other words, since the tokens of "Adam Smith_n" can be literally interpreted either referentially or attributively,[139] it follows that "Adam Smith_n," qua name type, is not inherently and exclusively referential, even

[139] And if one is willing to leave the realm of the "literal," it becomes apparent that names may be taken to express all sorts of things of a non-objectual nature. Consider the following example, due to Clark (1992), in which a name is used as a verb to express a relation derived from a property associated with its bearer: "He Houdini'd his way out of a locked closet." Here the speaker is alluding to Houdini's legendary abilities to escape from confinement and ascribing similar abilities to the person he is referring to.

when “unique” in Kripke’s sense. And if the name is not inherently and exclusively referential (that is, if it is not referential as a matter of its semantics), then neither can it be rigid. As mentioned above (§1), semantic rigidity (rigidity de jure) implies or presupposes semantic reference.

Now, it might be objected that “Adam Smith” here is really “short for” a longer, perhaps quotational expression, and that it could be paraphrased as “A person called ‘Adam Smith,’” for example. This revisionist impulse should be resisted. First, observe that (5) is perfectly grammatical as it stands. Second, notice that utterances of sentences like (5) seem to be quite common. Third, it seems completely appropriate and felicitous to utter (5) in the situation described above. Finally, and most importantly, a question arises as to how exactly the name should be paraphrased. For instance, which would be the better paraphrase in this case? A) “A person called ‘Adam Smith’”; B) “The student whose name is ‘Adam Smith’”; C) “This student listed here as ‘Adam Smith’”; D) “Some guy called ‘Adam Smith,’ who is new to the department”; or E), “The student called “Adam Smith” whose name appears on the department roster and is advised by Professor Jones”? There seems to be no fact of the matter

as to how such paraphrasing should be done. Moreover, it is doubtful that if asked, the speaker would say that he had a specific paraphrase in mind when he uttered (5). (The problem here is analogous to that of trying to complete an "incomplete" definite description; see [Wettstein, 1981]).

Let us turn now to the other horn of the dilemma: rigidity construed as a property of name tokens. Admittedly, this is probably not Kripke's intended construal, yet there is at least some textual justification for interpreting the rigidity thesis as being about tokens, since Kripke sometimes speaks of "particular uses of sentences" and admits that what an expression means is determined in context. For instance, he writes:

> When I spoke of the "truth conditions" of [3], I perforce assumed a particular reading of [3]…In practice it is usual to suppose that what is meant in a particular use of a sentence is understood from the context. In the present instance, that context made it clear that it was the conventional use of "Aristotle" for the great philosopher that was in question. Then, *given* this fixed understanding of [3], the question of rigidity is: Is the correctness of [3], *thus understood*, determined with respect to each counterfactual situation by whether a certain single person would have liked dogs (had that situation obtained)? I answer that question affirmatively. (1970/1980, pp. 8–9, emphasis in the original)

If the correct understanding of an expression is determined in context, then what we are talking about is a

token of the expression, a concrete, datable utterance or inscription of the expression. The main problem with understanding rigidity to be a property of name tokens is that it would trivialize the rigidity thesis, as Bach (2002, p. 89) has argued. It would do so by relativizing rigidity to *uses* of names in contexts, as he points out. As Bach also notes (p. 89), from a semantic point of view, it is uninteresting to be told that a name (e.g. "Aristotle"), as it is being used in a given context, rigidly refers to an individual (e.g. Aristotle the philosopher) because the speaker is using the name to refer to that individual (e.g. Aristotle the philosopher) in that context.

Put more generally, the trouble with assuming that what a name refers to is given by the context is that rigidity could no longer be considered a semantic property of a name (rigidity de jure), but only a presumably pragmatic property of a use of a name in the context. As Bach emphasizes (1987/1994a, p. 155), from the fact that an expression may be used in a certain way in a given context nothing in general follows about the semantics of the expression.[140] To recall the example of Chapter 2, phrases of the form "can you…?" are regularly used to perform

[140] See also (Bach, 1999).

requests, even if the request is literally expressed only by the VP that follows the construction (e.g. “pass the salt”). Thus, even if it made sense to speak of someone as using a name “to refer rigidly” to an individual made salient by the context—or even if we were to assume that a whole community regularly used the name to refer rigidly to this individual in most contexts—it does not follow from this fact alone that the name itself rigidly refers to this individual.

In discussing his understanding of the notion of a possible world, Kripke states that “in describing that world, we use *English* with *our* meanings and *our* references.” (1980, p. 77, emphasis in the original) Interpreted loosely, Kripke’s remark may serve as a helpful pointer for grasping his idea of a possible world. However, upon a closer reading, it becomes apparent that Kripke here is *stipulating* that an English name has a certain referent. This practice, while permissible in logic, is simply not the way natural language works.[141]

So why not just say that rigidity is a property of tokens, a property supervening on the intentions and uses of proper names of speakers in particular contexts? As far

[141] Again, even if we were to grant Kripke his controversial one bearer/one name view, names still couldn’t be viewed as inherently referential expressions.

as this dissertation is concerned, we see no problem with this option. In fact, Neale has recently developed an account of rigidity from a thoroughly pragmatic, intention-based perspective. (See [Neale, 2008 forthcoming, Ch. 7].) Neale's pragmatic rendering of rigidity was accepted by Kripke after Neale's talk on this topic at the gala opening of the CUNY Saul Kripke Center on May 22, 2008. Nevertheless, we very much doubt that such a pragmatic construal is the idea of rigidity that Kripke and his followers have defended during the last thirty-odd years.

In conclusion, the thesis that proper names are rigid designators, as a matter of the semantic rules of the language, should be rejected for the following reasons:

1. Rigidity cannot plausibly be construed as a property of name types, even when Kripke's problematic one bearer/one name view is granted.

2. To construe rigidity instead as a property of tokens is to make rigidity a pragmatic matter, if it makes sense to speak of rigidity at all. (Specifically, rigidity would be a matter of the speaker's intention to refer rigidly, by her use of a name, to an individual who is salient to speaker and audience in the context.) Rigidity would not be rigidity de jure.

Chapter 5

"Every F"

1. Introduction

In Chapter 3 we saw that proper names, which are widely viewed as prototypical referring expressions, may be literally interpreted as having properties instead of objects as truth-conditional values, contrary to what the standard Referentialist picture implies. This chapter argues that noun phrases of the form "every F," which are widely viewed as prototypical *non-referring* expressions, may be literally interpreted as having multiple objects as their truth-conditional value, instead of the truth-conditional value standardly ascribed to them. The thesis to be defended, then, is that "every F" admits of a truth-conditionally significant referential interpretation that is not countenanced by the standard semantic account of this expression.

"Every F" is a quantificational noun phrase. By "quantificational noun phrase" or "QP" we mean expressions of the syntactic form **determiner-quantifier + noun phrase**. The term "determiner-quantifier" is borrowed from von Fintel (1994); it identifies a subset of determiners that includes "every," "all," "some," and "no," among others. These determiners, as well as the QPs they form, have often

been called “quantifiers” in philosophy.[142] This word has also been applied to English expressions belonging to other grammatical categories, such as pronouns (e.g. “all” in “dear all” or “with liberty and justice for all,” “some” in “some like it hot”), adverbs (e.g. “always,” “sometimes,” “never,” “usually,” “seldom”), and adjectives (e.g. “much,” “sufficient,” “several,” “ten thousand”). Here “quantifier” will be used as an umbrella term to refer either to a determiner-quantifier or to a quantificational noun phrase, when the distinction between the two is unimportant for the purposes of the discussion. Pronouns, adverbs and adjectives will not be considered in this chapter. Even though it is tempting to try to extend our thesis about “every F” to other QPs—and it is reasonable to suppose that other QPs may also be interpreted referentially—this chapter focuses solely on “every F” for the sake of clarity.

The chapter is organized as follows. §2 briefly reviews the standard semantic account of “every F.” §3 gives reasons for concluding that “every F” may have objects as truth-conditional values. §4 further articulates this claim. §§5-6 defend this conclusion against two possible objections. §7 offers a Relevance-Theoretic account of the

[142] Russell (1905) called QPs like “every F” “denoting phrases.”

differing truth-conditional values of "every F." §8 considers an alternative to this account.

2. The Standard Semantics of "Every F"

Over the last thirty years, quantification in natural language has been extensively investigated and one theory has emerged as the standard theory of quantifiers, namely, Generalized Quantifier Theory (GQT). The origins of GQT are in (Mostowski, 1957) and (Lindström, 1966). The theory was developed in important ways by Montague (1974), Barwise and Cooper (1981), Higginbotham and May (1981), Westerståhl (1985), and Keenan and Stavi (1986). This section discusses in an informal manner the main ideas of GQT, with the aim of providing a background for the discussions of subsequent sections. It is beyond the scope of this chapter to offer a detailed presentation of the theory. There are a number of good expositions of GQT; see, for example, (Westerståhl, 1989), (Keenan & Westerståhl, 1997), and (Glanzberg, 2006).

The core claim of GQT is that the truth-conditional value of a determiner-quantifier such as "every" is a set of sets (or, alternatively, a property of properties).[143] More

[143] For ease of exposition, the following discussion is in terms of sets rather than properties.

precisely, the determiner-quantifier is said to express[144] a *relation* between sets. The relation expressed by the determiner-quantifier is said to hold between the set expressed by the nominal element with which it is associated in the subject of a sentence ("F") and the set expressed by the VP of the sentence. For example, consider (1).

(1) Every bottle is empty.

GQT offers the following analysis of (1). The VP "is empty" expresses the set of empty things. The subject QP "every bottle" is broken down into its component parts, the determiner-quantifier "every" and the common noun "bottle." "Bottle" expresses the set of bottles. "Every" then relates the two sets in a certain way. Specifically, "every" is taken to express the subset relation. Thus, the sentence "Every bottle is empty" is taken to express that the set of bottles is a subset of the set of empty things.

The very same analysis applies to the closely related determiner-quantifiers "all" and "each" as well. For instance, the sentence "All renates are cordates" is

[144] In this chapter, we will use "to express" to mean "to have as truth-conditional value," where (as in the case of the theory discussed in this section) the property of having a truth-conditional value is viewed as a purely linguistic and semantic matter. This is done simply for expository ease; it should be clear that our position is that the ascription of truth-conditional values to expressions is a pragmatic and contextual matter, and not something encoded in language.

interpreted as expressing that the set of renates is a subset of the set of cordates. "Each student presents a paper" is taken as expressing that the set of students is a subset of the set of things presenting a paper. Other determiner-quantifiers are treated in a similar fashion. In the case of "some," for example, the determiner-quantifier is taken to express the non-disjointness relation. Thus, the sentence "Some flowers are green" is taken to express that the set of flowers overlaps the set of green things.

The relational view of determiner-quantifiers may be traced back to (Frege, 1980, §47). The essence of the view is clearly stated by Frege in (1892b/1997b):

> [T]he words "all," "every," "no," "some" are prefixed to concept words [predicates]. In universal and particular affirmative and negative utterances we express relations between concepts and indicate the specific nature of these relations by means of these words. (1997b, p. 187)

One of the advantages of such a relational analysis of determiner-quantifiers is that these natural language expressions may be given explicit definitions in terms of well-understood mathematical notions. For instance, "every" ("all," "each") and "some" would be defined as follows:

Every_U (X, Y) iff $X \subseteq Y$.

[Sets X and Y, which belong to a universe U, satisfy the binary relation “Every” if and only if set X is a subset of set Y.]

Some_U (X, Y) iff $|X \cap Y| \geq 1$.

[Sets X and Y, which belong to a universe U, satisfy the binary relation “Some” if and only if the intersection of X and Y has a cardinality equal to or greater than 1.]

Definitions in terms of sets and relations can be given for a variety of other natural language determiner-quantifiers. (For details, see any of the references mentioned above.)

A consequence of the relational and second-order “meaning” of determiner-quantifiers is that both determiner-quantifiers and the QPs they form are barred from having objects as their truth-conditional value. Determiner-quantifiers such as “every” cannot express objects, since according to GQT their truth-conditional value is a relation between sets. QPs such as “every bottle” cannot express objects either, since the components of a QP do not express objects. In the case of “every bottle,” for example, the phrase is composed of two elements, a determiner-quantifier and a noun. Neither expression has objects as its truth-conditional value:

nouns express sets and determiner-quantifiers express relations between sets.

According to GQT, therefore, quantifiers do not belong to the semantic category of expressions that have objects as their truth-conditional value (a category that is usually held to include proper names, pronouns, and demonstratives). The "semantic category" of an expression may be defined by the kind of truth-conditional value it has; and as we've just seen, according to GQT quantifiers express relational rather than objectual truth-conditional values.

What evidence supports this standard semantic account of quantifiers? A common argument involves a comparison with names. The argument goes as follows.[145] We are asked to consider a pair of sentences, one of which contains a name and the other a quantifier of extreme generality, such as "everything." For instance:

(2) Sam weighs 300 lbs.

(3) Everything weighs 300 lbs.

It is noted that on the surface the two sentences appear similar, in that they both seem to predicate the property of weighing 300 lbs. to whatever the subject noun phrase expresses. It is assumed that "Sam" has an object as

[145] In rehearsing the argument we follow (Glanzberg, 2006).

truth-conditional value, Sam, and it is asked whether "everything" could similarly have an object as truth-conditional value. Then it is wondered what sort of object everything could be. On (reasonable) metaphysical grounds it is rejected that everything could be an object; for if it were, it would be at the same time a particular object and all objects, i.e. everything there is. It is concluded that in general, unlike names, quantifiers do not express particular objects. Then a second hypothesis is put forward: perhaps "everything," like a predicate, expresses a *set* of objects. So (3) would now be interpreted as saying that the set of all objects is identical or coextensive with the set of objects that weigh 300 lbs. Plausibly enough, the hypothesis predicts that (3) is false (not everything in the world weighs 300 lbs.). However, it is observed that the same analysis returns the wrong results for (4):

(4) Nothing weighs 300 lbs.

"Nothing" (on the foregoing reasoning) would express the empty set. But then, every member of the empty set also belongs to the set of things that weigh 300 lbs.: trivially so. Hence (4), contrary to what we would expect, is true. On the basis of this absurd result, the hypothesis that a quantifier might express a set of objects

is rejected. Therefore, it is concluded that quantifiers express neither objects nor sets of objects; they must express something else. Taking a cue from Frege's interpretation of determiner-quantifiers as second-order "concepts," it is alleged that for many natural language determiner-quantifiers, including "every," to construe them as second-order relations explains much of the behavior we intuitively attribute to them, while at the same time avoiding wrong results. So it is concluded that determiner-quantifiers express relations between sets, viz. the set expressed by the nominal element in the subject and the set expressed by the sentence's VP. From this conclusion, plus the assumptions that nouns express sets and that the truth-conditional value of a complex expression is built out of the truth-conditional values of its parts, it is further concluded that QPs such as "every F" do not express objects, but a set and a second-order relation.

3. Evidence Conflicting with the Standard Account of "Every F"

It cannot be denied that GQT is a rich and well-developed theory that succeeds in explaining a great deal of the behavior of natural language quantifiers. But one might

wonder whether the disanalogy with names constitutes an adequate evidential basis for the theory. No matter how mathematically elegant and productive a semantic theory for a natural language expression may be, it still has to square with the facts of usage. The GQT claim that determiner-quantifiers express second-order relations between sets is at bottom an empirical claim. The Fregean-derived relational view of determiner-quantifiers, with its consequence that QPs cannot expresses objects, has not always enjoyed its present hegemonic status. In (1952), Strawson argued that QPs could be used to refer:

> The curious assumption implicit in regarding the quantificational logic as adequate for the analysis of ordinary categorical utterances [utterances beginning with a quantifier such as "all," "some," "no," etc.] in general is, roughly speaking, the assumption that the only subject-predicate utterances [utterances "in which we use one part of our sentence to play the referring role and the remainder to play the ascriptive or classificatory role"] are utterances in the singular; that all other categorical utterances are positively or negatively existential. But of course we can *refer to*, and *talk about*, the members of a group collectively in just the same way that we can refer to and talk about an individual member of a group. Compare a utterance made in the words 'Miss Robinson has gone home' with a utterance made in the words 'All the members of the Robinson family have gone home.' The grammatical subject of each sentence is used to refer; in the one case, to an individual, in the other, to the members of a group. (p. 182, emphasis in the original)

Strawson's view that QPs may be used to refer to multiple individuals seems to have been mostly overlooked in recent

philosophical discussions on the semantics of quantifiers.[146] To our knowledge, only Neale (1990) and Bach (1987/1994a, Ch. 12) and have admitted that QPs may be used to refer. Neale's views will be addressed in more detail in §6. At present, a few more examples should be provided to cement the notion that a QP may be used to refer. Consider the following, which is modeled on an example of Bach's (1987/1994a, p. 244).

(5) A few examples will show how QPs may be used to refer.

In (5) the QP "a few examples" is being used to refer to the examples to follow. As Bach points out, such a referential use could be made explicit by adding a clause that specifies the examples in question, as in (6):

(6) A few examples, namely the ones presented here, will show how QPs may be used to refer.

[146] Linguists have paid attention to Strawson's examples, but they have focused on the question of presupposition (i.e. on the fact that the examples seem to show that certain English QPs carry a presupposition of existence of objects) and not on the issue of reference proper. See, e.g., (de Jong & Verkuyl, 1985), (Partee, 1987), (Diesing, 1992), and (Heim & Kratzer, 1998, pp. 153-63). Strawson's examples of the referential use—"All John's children are asleep" (p. 176) and "All the members of the Robinson family have gone home" (quoted above)—could be reasonably cited as cases where QPs seem to presuppose the existence of objects, viz. John's children in the first case and the members of the Robinson family in the second. But this is relatively uncontroversial, since these QPs contain an expression widely admitted to carry existential presuppositions, namely "the." ("All John's children" is usually viewed as a surface form of "all of the children of John.")

For Bach, if a "namely…" clause may be inserted to make explicit the reference, then one may be reasonably sure that the QP is being used to refer to multiple individuals. Consider another example:

(7) Some chapters are about proper names.

As a response to a question about the contents of this dissertation, "some chapters" in (7) may be used to refer to Chapters 2, 3, and 4, which concern proper names. This reference can be made explicit using Bach's "namely…" formula:

(8) Some chapters, namely Chapters 2, 3, and 4, are about proper names.

It would seem that the referential *use* of QPs can hardly be debated. The question, then, is whether such a use is of truth-conditional significance. That is to say, the issue worth investigating would seem to be whether utterances containing QPs may have different truth conditions, and diverging truth values, depending on whether the QP is interpreted referentially (as having multiple objects as truth-conditional value) or according to its standard GQT semantics.

In what follows it will be argued that the referential use of "every F" is indeed truth-conditionally significant. Specifically, it will be argued that if "every F" is

interpreted referentially, the utterance containing it will have truth conditions involving particular objects, and not the set of objects that are F plus the subset relation, as ordained by GQT. In other words, if "every F" is interpreted referentially, the utterance will express an object-dependent proposition and not an object-independent proposition.[147]

Bach denies that the referential use of a QP is truth-conditionally significant in this sense. He holds that from a strictly semantic point of view, QPs are non-referring expressions:

> Semantically, quantifier phrases do not denote[148] and are not referring expressions (strictly literal uses of quantifier sentences express general propositions), but our examples have shown that they also have referring uses. (1987/1994a, p. 248)

As we shall see in §6, Neale also rejects that the referential interpretation of a QP is truth-conditionally significant. For both Bach and Neale, sentences containing QPs semantically express object-independent propositions; in no case do their truth conditions include particular objects. What is Strawson's position? Strawson, after

[147] See footnote 96 of Chapter 3 for an explanation of what is meant by "object-dependent proposition" and "object-independent proposition."
[148] Bach uses "to denote" to mean "to refer," in the traditional philosophical sense of this term: "having an object as its truth-conditional value."

all, appears not only to be the first[149] to recognize the referential use of a QP, but also the tension between the "official," logical construal of these expressions and their everyday referential uses. To extract Strawson's view on the truth-conditional significance of the referential use of a QP is a tricky matter, however, since he is writing at a time prior to the development of GQT, Chomskyan linguistics, Tarski-Davidson truth-theoretic semantics, and the current debate on semantics-pragmatics distinction. Here we will not venture a guess as to what Strawson's view on all the relevant issues would be.

Let us now turn to the following example, which serves to demonstrate the truth-conditional significance of the referential reading of "every F." Consider (1) again.

(1) Every bottle is empty.

Suppose two people, A and B, are talking about four bottles sitting on a counter in front of them. A utters (1), intending to refer to the bottles by using "every

[149] Well, Strawson was perhaps the first *philosopher of language* to recognize the referential use of a QP, but probably the first recorded referential use is one that appears in Homer's *Odyssey*. The Cyclops has been told by Odysseus that Odysseus's name is "Nobody" (or "No man"). Battling Odysseus, he desperately cries out to his friends, "Nobody is killing me!", but he fails to get them to come to his aid. Here the Cyclops is interpreting "nobody" referentially, with unpleasant consequences for him partly as a result, while his friends, understanding the phrase in a non-referential manner, figure he's doing just fine. Lewis Carroll's *Alice in Wonderland* and *Through the Looking Glass* also contain various examples of referential uses of QPs. Thanks to Neale (conversation) for the *Odyssey* example and reminder of the examples in Carroll.

bottle." B, correctly recognizing this intention, interprets A as using the subject QP of (1) as a whole to refer to the four bottles on the counter. That is, B understands A to be asserting something about these four bottles and not about bottles in general or the set comprised of these four bottles: viz. that the bottles are empty.

Our claim, then, is that the truth of what is expressed by (1) in this situation depends on the four bottles' being empty. It is only if *these* particular bottles are empty that what A said, the proposition expressed by his utterance of (1) in this context, is true.

Once again, the question is not whether "every bottle" can be used to refer to particular bottles; Strawson, Bach and Neale have already argued that many quantifiers may be used to refer. The issue rather is whether the referential use results in truth conditions different from those predicted by the GQT analysis. We have seen in the previous section what the GQT analysis for (1) is. According to this analysis (1) expresses the general proposition that the set of bottles is a subset of the set of empty things; (1) is true just in case the set of bottles *is* a subset of the set of empty things. So if this analysis is correct, the interlocutors' interpretation of

"every bottle" as referring to the four bottles on the counter will likely have to be explained extra-semantically; a hearer's referential interpretation does not endow the phrase with an objectual truth-conditional value in this context.

The following reasons support the opposing view that in this context, "every bottle" expressess the four bottles and (1) expresses an object-dependent proposition.

First, in this case "every bottle" seems replaceable without a marked loss of truth-conditional content by four names for the four bottles or a plural demonstrative phrase containing a quantifier of greater specificity. Names and demonstratives are considered paradigmatic referring expressions.[150] For instance, "every bottle" could be replaced as follows:

(9) Bottle 1, Bottle 2, Bottle 3, and Bottle 4 are empty.

(10) These/those four bottles are empty.

In (9) the conjunction of names specifies (in an artificial way, of course) the four bottles the speaker is referring to and in (10) the demonstrative phrase "these/those four bottles" does the same. Intuitively, (1)

[150] This is not to deny that demonstratives and demonstrative phrases have non-referential uses. For an extensive discussion of non-referential uses of demonstrative phrases, see (King, 2001). (King's discussion, though extensive, may still be viewed as incomplete—cf. [Neale, 2007b].)

(as uttered in the context described above), (9) and (10) appear to express the same proposition and to have the same truth conditions. In each of the three cases the utterance is true if and only if the particular bottles are empty.

Just to clarify, it is not being claimed here that in *all* cases "every F" may be replaced without remainder by a conjunction of names or a plural demonstrative phrase. In other cases, including perhaps also referential uses in contexts different from the one described above, the quantifier imparts the additional information that the bottles in question are all the bottles which possess the property of being empty; i.e. that those are all the empty bottles there are. But in many other, perfectly quotidian cases, such substitutions seem acceptable.

Second, the objectual truth-conditional value of "every bottle" and the object-dependence of (1) in this context would seem to be confirmed by a report of what A said that contained a plural pronoun referring to the bottles. Like names and demonstratives, pronouns are also considered prototypical referring expressions.[151] For example, it seems that (11) is true:

(11) A said that they [the four bottles] are empty.

[151] This is not to deny that pronouns have non-referential uses. Bezuidenhout (1997b), citing an example in (Schiffer, 1995), argues that pronouns may be used attributively.

But if (11) is true (and it intuitively seems to be), then this can only be because the embedded sentence expresses the same proposition as (1). In other words, the truth of (11) piggybacks on the truth of (1), where "every bottle" is interpreted to expresses objects.[152]

[152] It might be argued that this is an example of unbound cross-sentential anaphora, and that there are analyses that would treat "they" in (11), not as a referring expression, but as having the quantificational truth-conditional value expressed by its antecedent in (1), "every bottle." One such analysis would be Neale's (1990, Ch. 5). Neale's analysis may be better illustrated by altering the example. Suppose A utters (1*):

(1*) Every bottle is empty. I'll put them in the recycling bin.

The pronoun "them" in the second sentence is anaphoric on the QP "every bottle" in the first; by "anaphoric" is meant that the pronoun is supposed to have the same truth-conditional value as the previously mentioned phrase, which serves as its "antecedent." Neale's claim is then that "them" in the second sentence does not refer to particular objects, the individual bottles, but has a relational truth-conditional value similar to the one standardly ascribed to "every bottle." How does the pronoun carry the relational truth-conditional value of its antecedent (or a truth-conditional value that is very similar)? According to Neale, unbound, cross-sententially anaphoric pronouns such as "them" in (1*) have the semantics of definite descriptions. For example, "them" in (1*) would have the same semantics as the plural definite description "the bottles." For Neale, definite descriptions have a quantificational (relational) semantics, just as QPs of the form "every F" do. So in (1*), the pronoun (qua definite description) and the antecedent QP both expresses a set and what he calls a "maximal" relation—a relation "exhausting the range of the predicate" with which it combines. (1990, p. 180) In (1990) Neale supports this analysis by showing how, for a variety of sentence sequences containing different types of QP, treating the anaphoric pronoun in the sequence as a definite description yields the truth conditions that we would intuitively ascribe to them.

While Neale's analysis seemingly applies to (1*), and hence appears to undermine the "pronoun co-reference" argument in favor of the referential reading of "every F" that is being made above, it is unclear that it applies to (11). For (11) can be discursively disconnected from (1) and still be true. For one thing, in our example it is not the speaker of (1), A, but a different person (say, B) who utters (11). For another thing, B can utter (11) hours, days, weeks, etc., after A's utterance of (1), and what she expresses still appears true. Suppose B utters (11) several days after A uttered (1), and that she points to the bottles as she makes her utterance. In this case B's use of "they" is unequivocally referential. "They" expressess the

Third, expressions for "every" in other languages encode natural gender. As explained in the Chapter 3 (§2), an expression that possesses natural gender encodes the information that the things to which the expression applies are sexed. Interestingly, the majority of referring expressions in English and other languages encodes natural gender. To be used correctly, the referring expression's gender must correspond to the referent's sex. One example of a language in which the word for "every" encodes natural gender is German. To be used properly to refer, the German counterpart of "every F" ("*jede/r/s* F")[153] must correspond to the sex of the things referred to, if they are sexed.[154]

Now, from the fact that "*jede/r/s*" possesses natural gender it doesn't follow that it is a referring expression. (Though the fact that the majority of referring expressions in English, German, and other languages possesses natural gender is certainly very suggestive.) What does follow is that the quantifier encodes descriptive information. That

---

bottles B is pointing to. The problem is that though the pronoun is referential in this case, what is expressed by (11) still appears true. In sum, Neale's analysis does not seem adequate when further contextual details are filled in.

[153] "*Jede*" is the nominative feminine form, e.g. "*jede Frau*" ("every woman"); "*jeder*" is the nominative masculine form, e.g. "*jeder Mann*" ("every man"); "*jedes*" is the nominative neuter form, e.g. "*jedes Auto*" ("every car").

[154] A couple of notorious exceptions are "*Weib*" ("woman") and "*Mädchen*" ("girl"), which have feminine natural gender, but neuter grammatical gender. When combined with these words, "*jede/r/s*" must take neuter gender.

is to say, the expression *describes* the entity to which it applies; it specifies one of its attributes. In the case of "*jede/r/s*," the quantifier specifically encodes the information that the entity is gendered.[155] Why is this relevant to the semantics of "every F"? Because this information seems incompatible with the truth-conditional value ascribed to "every F" (or rather, "*jede/r/s* F") by GQT. For it would be absurd to suggest, for instance, that the subset relation supposedly expressed "every" is a sexed entity. In contrast, the gender information is entirely compatible with the claim that "every F" may have objects as truth-conditional values, since objects may be gendered.

In conclusion, the reasons to think that "every F" admits of a truth-conditionally significant referential interpretation are the following:

1. In many cases (though not all), it seems that "every F" can be replaced by a conjunction of names of objects without loss of truth-conditional content.

2. In many cases (though not all), it seems that "every F" can be replaced by a plural demonstrative phrase, e.g. "these/those Fs," without loss of truth-conditional content.

3. The truth of a report of what was said by the utterance of a sentence containing "every F" will piggyback on the truth of a sentence in which the

[155] It would be reasonable to assume that a quantifier encodes further information besides natural gender. However, it is unnecessary to guess as to the nature of this further information, if any. For our purposes, the important thing is that quantifiers *do* encode descriptive information.

quantifier is interpreted referentially. The report will contain a plural referring term that is co-referential with the quantifier in the original sentence.

4. In other languages, quantifiers possess natural gender, a fact seemingly incompatible with their GQT semantics but compatible with a referential interpretation positing an objectual truth-conditional value.

4. "Every F" and the Referential/Attributive Distinction

One way of characterizing the two interpretations of "every F," the relational GQT interpretation and the referential one presented in the previous section, is to say that "every F" admits of the referential/attributive distinction. The referential/attributive distinction is usually associated with (Donnellan, 1966), though Neale (1990, pp. xii-xiii) traces it to (Rundle, 1965), (Mitchell, 1962), (Geach, 1962), (Marcus, 1961/1993), (Hampshire, 1959), and the theological writings of Antoine Arnauld (Neale refers to [Dominicy, 1984]).

In (1966) Donnellan argued that the referential /attributive distinction applies to definite descriptions. The distinction has since been claimed to hold for other types of expression as well: (Chastain, 1972) proposes that it holds in the case of indefinite descriptions;

(Bezuidenhout, 1997) in the case of pronouns;[156] and (Martinich, 1977), (Devitt, 1981), and (Bezuidenhout, 1997) in the case of proper names. In suggesting that "every F" also admits of the referential/attributive distinction, we draw on an interpretation of the distinction that has its origins in (Hornsby, 1977) and (Kaplan, 1978) and has been elucidated further in (Neale, 1990, Ch. 2) and (Récanati, 1993, pp. 278-9).[157]

According to this interpretation, the referential/attributive distinction is a distinction between two types of proposition (seemingly) literally expressed by a sentence containing an expression to which the distinction is said to apply. For instance, the referential /attributive distinction for definite descriptions is a distinction between two types of proposition (seemingly) literally expressed by a sentence containing a definite description in subject position, i.e. a sentence of the form "The F is G." When used referentially, the sentence expresses an object-dependent proposition. The sentence expresses a truth if and only if the intended referent in the situation is G. When used attributively, on the other hand, the sentence expresses an

[156] As mentioned in footnote 152, Neale (1990) also argues that not all uses of pronouns are referential.
[157] See also (Devitt, 2004) and (Bach, 2008 forthcoming).

object-independent proposition. The sentence expresses a truth if and only if there exists an object and that object (and no other) is both F and G.

When used attributively, in other words, the sentence will have the kind of truth conditions specified by Russell in (1905). According to Russell's classic analysis, sentences of the grammatical form "The F is G" have the underlying "logical form" $\exists x\ [\forall y\ (Fy \leftrightarrow y = x)\ \&\ Gx]$, which, informally put, says that there is one unique thing that is both F and G. Here the logical notation allows us to see clearly how the attributive use of the (singular) definite description would result in the expression of a general proposition. We see that the quantificational (relational) meaning imputed to "the" relates two properties, F and G. Specifically, it ascribes to F the higher-order properties of being uniquely instantiated and being non-disjoint from G. Although on Russell's account "The F is G" is true if and only if there is an object that uniquely instantiates F, the logical notation makes it clear that no particular object has been singled out; the existential quantifier ranges over arbitrary objects represented by variables. In short, the sentence expresses general information; it expresses a relation between two properties.

It is worthwhile to review briefly the first of the examples in (Donnellan, 1966) in order to better bring out the contrast between the referential and attributive uses and their truth-conditional consequences. Donnellan asks us to consider (12):

(12) Smith's murderer (i.e., the murderer of Smith) is insane.

Suppose A encounters the body of Smith, which shows signs of having been murdered in a gruesome way. In uttering (12) in these circumstances, A may be understood as using the description "Smith's murderer" attributively. That is, A is saying that there is some unique individual who is the murderer of Smith and whoever this individual is, he or she is insane. The claim is a general one: someone or other murdered Smith and this person is insane.

Let us now suppose that A has not seen the body, has no direct knowledge of the circumstances of Smith's murder, and is sitting at the trial of Jones, who has been arrested and charged with the murder. Jones is behaving in a very strange way. Suppose further that A is convinced of the prosecution's case against Jones. In uttering (12) in this second situation, A may be understood as using "Smith's murderer" referentially, to point to the individual before her, at whom she is looking. The proposition expressed by

(12) in this situation is a proposition about Jones; and it is true if and only if Jones is insane. Donnellan stresses that A succeeds in referring to Jones by means of the description "Smith's murderer" even if in fact it was not Jones who murdered Smith, but someone else, say Robinson. The description may be interpreted as referring to Jones even if Robinson murdered Smith, or even if Smith wasn't murdered at all.

This example, like the others that Donnellan discusses in (1966), exploits the possibility of a misdescription, of a situation in which the satisfier of the description and the object referred to in the situation are not identical. Presumably, the reason he lays so much emphasis on the possibility of misdescription is to highlight the contrast between the two kinds of use in a vivid way. Because of Donnellan's focus on the misdescription cases, one might be led to think that such misdescription cases effectively furnish the contrast between the two uses, or that misdescription cases demonstrate the truth-conditional significance or even the downright semantic ambiguity presented by the referential use of definite descriptions. However, as Récanati (1993), Reimer (1998), and Devitt (2004) point out, this would be a mistake. The semantic ambiguity theorist must rest his case not on the

misdescription examples, but on the cases in which the description's satisfier and the intended referent in the situation are identical.

Be that as it may, if the referential/attributive distinction is understood in the manner described above, it would seem to apply to "every F." As we said above, the usual interpretation cashes out the distinction in terms of the types of proposition expressed by a sentence containing an expression admitting of the distinction. The referential interpretation yields an object-dependent proposition and the attributive interpretation an object-independent proposition. Now, it is in these precise terms that the referential and relational GQT interpretations of "every F" were explained in the previous section. We said that when (1) is interpreted relationally according to the GQT analysis, the sentence expresses an object-independent proposition. The sentence expresses a truth just in case the set of bottles is a subset of the set of empty things. In this case "every," like "the" on the Russellian construal, expressess a second-order relation between sets of arbitrary objects. Thus, we may say that the relational interpretation of "every F" corresponds to the attributive reading. On the other hand, when (1) is interpreted referentially as in the example presented in §3, the

sentence may be taken to put across an object-dependent proposition; (1), as uttered in the context described above, would be true just in case certain individual bottles are empty. Thus, like definite descriptions, "every F" also admits of a referential reading.

As in the case of definite descriptions, the referential and "attributive" interpretations of "every F" are truth-conditionally significant, in the sense that a utterance containing the expression will have different sets of truth conditions (will express different types of proposition), depending on whether it is interpreted attributively or referentially.

Now, not all philosophers believe that the referential/attributive distinction is of theoretical importance. No opinion on this matter is proffered here. The point of this section was merely to show that our claim that sentences containing "every F" may have two different sorts of truth conditions, object-dependent and object-independent, could be expressed in terms of this distinction.

5. The Narrowing Objection

The next two sections consider and respond to two objections that might be made against the idea that "every

F" may be interpreted as having objects as truth-conditional values. First, it might be objected that this idea confuses reference with the restriction (or "narrowing") of the quantifier to a specific set by the context.[158] To make the objection concrete, let us focus once again on (1).

(1) Every bottle is empty.

The objection then is that in the situation presented above, where, we argued, "every bottle" expressess four bottles on the counter, there is narrowing but no referring.

The objection observes that, in general, QPs used in real-life conversations are contextually restricted.[159] They do not have universal values. In the situation presented above, "bottle" should not be construed as expressing[160] the set of all bottles, but rather the contextually and implicitly restricted set of the four bottles on the counter. Accordingly, "every" should not be interpreted as ranging over the set of all bottles either,

[158] This objection was made by Robert Fiengo via electronic correspondence.
[159] We set aside the issue of how exactly such contextual restriction is supposed to work. A prominent proposal in philosophy on quantifier domain restriction is Stanley's (Stanley, 2000; Stanley and Szabó, 2000; and Stanley, 2002a). This proposal, however, has many difficulties. We cannot go into these here.
[160] In the sense specified in §2.

but over the set comprised by the four bottles on the counter.

The problem, the objection continues, is that the restricted reading of the QP and the alleged referential interpretation are indistinguishable, in the sense that (1) will have the same truth value, regardless of whether "every bottle" is interpreted referentially or according to a restricted relational reading. The worry is that if the referential interpretation cannot be distinguished from the restricted relational reading, the referential interpretation might be illusory. The objector, in other words, challenges us to produce an instance in which (1), where "every bottle" is interpreted referentially, yields a different truth value from (1), where "every bottle" is interpreted as contextually restricted.

This objection fails for the following reasons. First of all, it is plausible to suppose that the essential differences that exist between particular objects, on the one hand, and sets of objects, on the other, would be at least tacitly clear to the interlocutors in the example concerning (1). Certain things may be truly predicated of the set of the four bottles on the counter that may not be truly predicated of the four individual bottles, and vice-versa. For example, (13) is true but (14) seems false:

(13) The set of these four bottles on the counter is such that it has a lower cardinality than the set of natural numbers.

(14) These four individual bottles on the counter (or: Bottle 1, Bottle 2, Bottle 3, and Bottle 4) are such that they (each individual one) have a lower cardinality than the set of natural numbers.

(13) is true but (14) seems false, since the property of having a certain cardinality is a property that can only be properly ascribed to a set of things. Here we are talking about the individual bottles themselves, and not the unit set made up of each bottle.

Likewise, (15) seems true but (16) false:

(15) These four individual bottles on the counter (or: Bottle 1, Bottle 2, Bottle 3, and Bottle 4) are made of very thin and breakable glass.

(16) The set of these four bottles on the counter is such that it is made of very thin and breakable glass.

What (15) says might be true of the bottles, but clearly (16) is false, since sets, as abstract entities, are not made of very thin and breakable glass.

But even if the differences between a set and a particular object are not implicitly assumed or understood by the interlocutors, it *is* possible to produce an example

showing that (1) may have diverging truth values, depending on whether "every bottle" is interpreted referentially or relationally. This possibility would decisively refute the objection being entertained. That (1) may yield diverging truth values, depending on whether it is interpreted referentially or relationally, can be demonstrated by formulating a mistake case analogous to the misdescription scenarios Donnellan (1966) devised to illustrate the referential/attributive distinction.

For instance, suppose the speaker isn't wearing his glasses and is looking at the counter from some distance away. What he mistakes for bottles are really four bottle-shaped vases. The vases are empty; they have no liquid in them. By uttering (1), has he said something true or false? On the GQT interpretation of "every bottle" in uttering (1) he certainly hasn't uttered a truth, since (1) in this case is analyzed as saying that the set of four bottles on the counter is a subset of the set of empty things. The sentence is true just in case the set of four bottles on the counter *is* a subset of the set empty things. But there is no set of four bottles on the counter, so the sentence cannot be true.[161] However, there does seem to be

[161] As von Fintel (1994, p. 7) notes, while the first-order formula $\forall x (P(x) \rightarrow Q(x))$ is true even if $P(x)$ is true of nothing (in virtue of the

a sense in which what the speaker said is true. The speaker is thinking of those four objects at which he is looking. He is claiming that these objects are devoid of contents and he is right about this. His charitable audience, who recognizes his intention and knows which objects he is referring to, would judge that his utterance of (1) is true (or true enough). Therefore, to say that "every bottle" has the four bottles on the counter as its truth-conditional value is not the same thing as saying that this quantifier is contextually restricted to the set of four bottles on the counter.

6. The Gricean Objection

The second objection accepts that a quantifier may on occasion be interpreted as referring to particular objects, but denies that such an interpretation is truth-conditionally significant. The second objection thus appraises the consequences of the referential interpretation of "every F" differently than we do, and proposes an alternative explanation of it.

The second objection would explain the referential interpretation of "every F" by means of Grice's pragmatics,

---

definition of material implication), natural language determiner-quantifiers such as "every" do not allow for vacuous truth.

as follows. In a given context, someone may use a sentence containing a quantifier to implicate a proposition concerning certain contextually salient objects. However, the fact that a speaker may so use a sentence containing a quantifier implies nothing about the semantics of the expression. The quantifier's semantics are still given by GQT. The sentence containing it should not be taken to express an object-dependent proposition as a result of the referential use. The sentence literally expresses a relation between sets. The object-dependent proposition understood in the context is an implicature; it is not *what is said* by the utterance of the sentence.

This Gricean line is explicitly taken by Neale (1990) in the course of his well-known rebuttal of the incompleteness objection to Russell's Theory of Descriptions.[162] The objection, originally raised by Strawson (1950, pp. 332-3), notes that a great many definite descriptions used in ordinary speech are "incomplete," in the sense that they do not explicitly provide conditions that are specific enough to pick out a unique object satisfying them. Yet in

[162] Neale's stance on the truth-conditional significance of referential uses of QPs and definite descriptions has since changed. (Neale, 2004a) represents his current thinking on this issue in connection to definite descriptions. (Neale, 2004b) and (Neale, 2008 forthcoming) present his overall view on referring. Here we will nevertheless continue to use (Neale, 1990), since the position expressed in that work may be seen as the best available statement of the Gricean Objection.

uttering sentences containing incomplete definite descriptions, communicators nevertheless often succeed in expressing true propositions. Since, to take Strawson's example, the sentence "The table is covered with books" may be used to make a true utterance despite the existence of many tables, Russell's analysis, according to which the proposition expressed by the sentence is true just in case there is one table in the universe, cannot be correct.

Neale counters that definite descriptions, like other quantifiers,[163] can be similarly incomplete, but we wouldn't want to say that these other quantifiers should be interpreted referentially. As an illustration of Neale's point, consider (17).

(17) All renates are cordates.

The QP in (17) is incomplete: "all renates" doesn't explicitly delimit the domain over which it ranges. The speaker could be saying that everything in the universe that has a kidney is also something that has a heart, or that everything currently alive on planet Earth that has a kidney has a heart. Or suppose the speaker is a biologist studying newly discovered life forms on Mars. After much fieldwork and analysis, she reports to her colleagues back

[163] As mentioned above, Neale believes that "the" is a determiner-quantifier with a relational semantics like "every," "all," and "some."

on Earth, “All renates are cordates.” In uttering (17), she is saying that all renates *on Mars* are also cordates. In this case, the domain of quantification is implicitly and contextually restricted to the class of renates on Mars.

Hence the question of whether the domain of quantification is implicit or not, and if it is, how to identify it, is independent of the question of whether the QP may be taken to have objects as truth-conditional values. Neale’s argument is that just because a QP is often incomplete doesn’t mean that it needs to be ascribed a referential interpretation; by analogy, the incompleteness of many definite descriptions (such as “the table”) doesn’t justify ascribing a second, referential semantics to definite descriptions.

Neale is correct that incompleteness doesn’t entail reference, but his assumption that quantifiers do not admit of a truth-conditionally significant referential interpretation is contestable. Neale allows that depending on the context, certain utterances could give the impression that a quantifier may have objects as values. However, he unequivocally rejects the possibility that this may be more than mere appearance. He presents the following example, concerning “everyone:”

> Suppose it is common knowledge that Smith is the only person taking Jones' seminar. One evening, Jones throws a party and Smith is the only person who turns up. A despondent Jones, when asked the next morning whether his party was well attended, says,
>
> [18] Well, everyone taking my seminar turned up
>
> fully intending to inform me that only Smith attended. The possibility of such a scenario, would not lead us to complicate the semantics of "every" with an ambiguity; i.e., *it would not lead us to posit semantically distinct quantificational and referential interpretations* of "everyone taking my seminar." (1990, pp. 87–8, emphasis added)

This passage makes it plain that Neale rejects a truth-conditionally significant referential interpretation for quantifiers. It also makes it plain that his argument against the semantic ambiguity thesis for definite descriptions crucially turns on the assumption that quantifiers are relational and do not admit of a referential interpretation. Here, though, the issue is not incompleteness so much as the appearance of reference. But the structure of the argument is the same: the appearance of reference created by certain uses of QPs does not sanction a semantic ambiguity for these expressions; by analogy, the appearance of reference created by certain uses of definite descriptions should not lead us to posit a semantic ambiguity in their case, either. According to Neale, the referential use of "everyone taking my seminar" to say something about Smith is to be explained by recourse

to Gricean pragmatics: the proposition about Smith seemingly expressed by the sentence is not literally expressed, but only implicated in the context. And this is also his overall strategy against the challenge posed by referential uses of definite descriptions to Russell's Theory of Descriptions. For him, the semantics of definite descriptions is given by Russell's theory. In a given context, however, communicators may derive through Gricean procedures an object-dependent proposition from the utterance of a sentence containing a definite description.[164]

While Neale is perhaps right to maintain that Gricean procedures adequately explain the understanding of a proposition about Smith from an utterance of (18) in the context he describes, there are reasons to doubt that such a Gricean explanation (i.e. one that involves deriving an implicature from the utterance of a sentence literally expressing a different proposition) will work for all cases. In particular, one may doubt whether such a Gricean explanation works for (1) in the context we described above. We argued above that in this case it is (1) itself, as uttered by A in the context we described, which

[164] Neale provides a full derivation for an utterance containing the referentially used definite description, "the Chairman of the Flat Earth society," at p. 89.

expresses an object-dependent proposition about the four bottles on the counter.

One way to show that the objectual truth-conditional value of “every bottle” is part of what is said by A’s utterance of (1), and not part of something she implicates, is by demonstrating that the objectual truth-conditional value figures in the truth conditions of a larger sentence in which (1) is embedded. If the four particular bottles that are the values of “every bottle” in this case are included in the truth conditions of a utterance that contains (1)—i.e. if these four objects are part of what is said by this larger utterance—then they can’t be part of an implicature generated by an utterance of (1). For Grice, information is either part of what is said or part of what is implicated.

Take once again the situation described in §3. Suppose B replies to A:

(19) If every bottle is empty, then you can put them in the recycling bin.

Or, suppose that B walks over to the bottles and, looking at them more closely, corrects A by saying:

(20) Not [it is not the case that] every bottle is empty; look, this one is half full.

The truth of (19) would seem to depend on whether the particular bottles are empty, and not on whether the set of bottles-on-the-counter is a subset of the set of empty things. For the consequent contains a pronoun ("them") that is co-referential with "every bottle." And in (20), the sentence following the semi-colon, which contains the demonstrative phrase "this one," can only be properly interpreted if the QP in the first sentence is also interpreted referentially.[165] As (20) itself implies by the inclusion of "look," the situation is one in which the communicators have the bottles in front of them, so that B may rely on perceptual evidence in order to determine A's informative intention to refer to those specific bottles.

Let us recapitulate the conclusions of the last two sections. We entertained two possible objections to the referential interpretation of "every F." The first, which might be called the Narrowing Objection, claimed that the referential interpretation of "every F" and the narrowing

[165] Once again, it may be objected that both the anaphoric pronoun "them" and the anaphoric demonstrative phrase "this one" may be ascribed relational truth-conditional values, and hence that the presence of these expressions in (19) and (20) does not by itself show that the antecedent QP has an objectual truth-conditional value. A proper reply to this objection would require going into the semantics of pronouns, definite descriptions and demonstrative phrases in considerable detail. This we cannot do here. The remarks made above must thus stay at an intuitive level. But though intuitive, they are nevertheless plausible. Footnote 152 sketches a possible reply to this objection: to interpret the expression in question according to one of the proposed relational analyses might turn out not to be convincing, once the context is described in greater detail.

of the quantifier are being run together in the example presented in §3, which purported to illustrate the existence of such an interpretation. Why are the referential interpretation and the restricted relational interpretation not the same thing?

1. It is plausible to suppose that people tacitly know the difference between objects and sets.

2. It is possible to concoct "mistake cases" analogous to misdescription examples, which reveal the difference.

The second, Gricean Objection, acknowledges the possibility that quantifiers may be interpreted referentially, but denies that this interpretation is truth-conditionally significant. Instead, the object-dependent proposition seemingly expressed when a quantifier is interpreted referentially is merely pragmatically conveyed, as an implicature, in the context. Against this objection it was argued:

> The object-dependent proposition can be shown to be what is said by embedding the sentence containing the quantifier in a larger sentence. The larger sentence will also have truth conditions involving particular objects, and not sets and relations.

But perhaps the most effective way of answering the Gricean Objection is by offering an entirely different yet compelling explanation of the literality and truth-conditional significance of the referential interpretation. The next section offers such an explanation.

7. A Relevance-Theoretic Account of the Referential Interpretation

The explanation presented in this section is based on Relevance Theory. It follows the same pattern as the account concerning names that was given in §5 of Chapter 3. For simplicity, our explanation shall continue to focus on (1) and the situation described in §3. What needs to be explained is then the fact that B has interpreted A's utterance of the phrase "every bottle" referentially. B identifies particular objects, the four bottles on the counter, as the truth-conditional contribution of this phrase to the content expressed by A's utterance of (1).

Let us flesh out the example just a bit further. Suppose the following dialogue took place:

B: (0) Are those bottles on the counter empty or full?

A: (1) Every bottle is empty.

B's relevance-guided interpretation process may be outlined as follows. According to Sperber and Wilson's Communicative Principle of Relevance,[166] A's utterance carries the expectation that it is optimally relevant to B. More than any other stimulus in the environment, the utterance commands B's attention and promises specific

[166] Stated in Chapter 1, §6.2.

cognitive effects for B without demanding excessive processing effort on her part. This guarantee of relevance validates an interpretive path of least effort.

The following facts obtain in the situation:

(a) There are four (and only four) bottles on the counter.

(b) The bottles are empty.

From contextual fact (b), B interprets "empty" to mean "containing nothing." From (a), B interprets "every bottle" to refer to the four individual bottles on the counter. She then interprets "is" predicatively: each one of the four bottles has the property of being empty. Thus she interprets (1) as expressing an object-dependent proposition concerning the four individual bottles.

This interpretation follows a path of least effort. The interpretation of "empty" as "containing nothing" is easily inferable from the context, since the bottles are empty and A and B can see that they are empty. A different interpretation of the word, e.g. "lacking meaning or value," would not be justified by the context and would require greater processing effort. The predicative interpretation of "is" is the most salient one in the context, provided that "every bottle" is taken to refer to individual bottles. The question then is: Is the

referential interpretation of “every bottle” the easiest to infer in the situation? Is it more available than the relational interpretation according to which “every bottle” expressess the set of bottles and the subset relation?

It would seem so. Here’s why. A utters (1) in reply to B’s question, (0). (0) contains the plural demonstrative phrase, “those bottles,” which refers to the four bottles. In using the demonstrative phrase, B has referred to the bottles herself; she knows which bottles she means.[167] She expects an answer having to do with those particular bottles. For (1) to answer (0), “every bottle” would have to be interpreted referentially, as being co-referential with “those bottles.” If “every bottle” is taken to expresses the set of bottles and the subset relation, the connection with B’s question and the context is not evident. The relational interpretation does not offer an answer to (0).

In Relevance-Theoretic terms, the relational interpretation of “every bottle” would require greater processing effort than the referential one, since the referential one immediately answers the question. So even

[167] This is easily seen if one assumes, with Neale (2004b), that there is an epistemic asymmetry in the perspectives of speaker and hearer. The speaker always knows what she means (e.g. what she is referring to and the message she intends to convey), but the hearer might fail to understand the speaker’s meaning.

if the relational interpretation were to have positive cognitive effects for B, it would not be as relevant to B. Therefore, the relational interpretation is not warranted from a Relevance-Theoretic point of view.

8. The Ambiguity Alternative

We shall close the chapter by briefly considering a possible alternative to the Relevance-Theoretic account outlined above. The alternative would be to postulate a semantic ambiguity for "every." The motivation for this proposal, let us assume, is the idea that if an expression makes different sorts of contributions to the truth conditions of a sentence, it is to be catalogued by different entries in our mental "lexicon." Each lexical entry would correspond to a different type of truth-conditional value. Thus, "every" would have two lexical entries, one corresponding to the relational reading of "every F," and the other to the referential interpretation of this phrase.

What is meant by "lexical entry"? According to Chomsky (1986, esp. §3.3.3.2), a speaker's knowledge of the words of her language is embodied in a mental lexicon. Each

word[168] the speaker knows is represented in this mental lexicon as a discrete lexical entry, which contains information about the word's linguistic meaning, syntactic characteristics, and how it sounds. As with print dictionaries, words of the same spelling and phonology, but differing in meaning or grammatical features, require separate entries specifying the different information. Thus, for example, "bank" has two lexical entries, one for "financial institution" and the other for "side of a river." In the same way, the verb "to hit" has two lexical entries, which correspond to two distinct syntactic features of this verb, transitivity (e.g. "John hit the baseball") and intransitivity (e.g. "John hit"). This information can be represented as follows:

Entry 1

Hit, Verb, [_ NP]

Entry 2

Hit, Verb, [_]

The entry gives the syntactic type of the word (a verb, in this case) and describes its syntactic behavior (in this

[168] Technically it is not a "word" in the ordinary sense that is stored in the lexicon (e.g. "crunchable"), but rather a "listeme." This term was introduced in (Di Sciullio & Williams, 1987) and is defined as a meaningful unit that cannot be generated by morphological rules and must therefore be memorized ("permanently stored") in the course of acquiring a language (e.g. "-able"). The distinction between words and listemes can be ignored for the purposes of this discussion.

case, whether the verb is transitive or intransitive). The two putative entries for "every," modeled on this example, would look like this:

Entry 1

Every, Determiner, [_ NP]

Truth-conditional value: subset relation

Entry 2

Every, Determiner, [_ NP]

Truth-conditional value: multiple objects

The entries say that "every" is a determiner, that it combines with noun phrases, and that it may express either the subset relation or multiple objects.

One difficulty with the ambiguity proposal, construed as a proposal about lexical entries, is that it does not seem to square with the notion of a lexical entry. From a Chomskyan point of view, the fact that an expression may have two different truth-conditional values would not justify the positing of two distinct lexical entries for that expression. Lexical entries are motivated by purely "internal" considerations: the reasons for positing two different entries for one word are phonological, syntactic or semantic, where this term is understood as pertaining to linguistic meaning and to the structural, word-to-word relationships licensed by an expression's meaning. But the

motivation for positing two lexical entries for “every” is “external:” it has to do with how communicators interpret words in contexts; and in particular, with how they judge the truth or falsity of what they hear.

Furthermore, it is not clear that the ambiguity proposal is incompatible with the Relevance-Theoretic explanation given above, and thus that it really counts as an alternative explanation to it. For even if we were to grant that an expression can have two different lexical entries because of the different truth-conditional contributions that it may be interpreted to make, this would not explain how a communicator selects the appropriate entry or is able to identify the correct truth-conditional contribution in a context. For this, an additional mechanism is needed. In our view, this mechanism is the relevance-guided interpretation process.

Works Cited

Abbot, B. (2005). Proper names and language. In G. Carlson and F. Pelletier (Eds.), *Reference and Quantification: The Partee Effect* (pp. 1–19). Stanford, CA: CSLI Publications.

Allers, M. (2006). Bruce, a predicate. Presented at the University of Texas Graduate Student Conference in April, 2006. Retrieved March 13, 2007, from http://www-personal.umich.edu/~mcallers/bruce.pdf.

Anscombre, J.-C. and Ducrot, O. (1983). *L'argumentation dans la langue*. Bruxelles: Pierre Mardaga.

Bach, K. (1987/1994a). *Thought and Reference*. Oxford: Oxford University Press.

Bach, K. (1994b). Semantic slack. In Tsohatzidis, S. (Ed.), *Foundations of Speech Act Theory* (pp. 267–91). London: Routledge.

Bach, K. (1994c). Conversational impliciture. *Mind and Language* 9: 124–62.

Bach, K. (1998). Standardization revisited. In A. Kasher (Ed.), *Pragmatics: Critical Concepts, Vol. IV* (pp. 712–22). London: Routledge.

Bach, K. (1999). The myth of conventional implicature. *Linguistics and Philosophy* 22: 327–66.

Bach, K. (2002). Giorgione was so-called because of his name. *Philosophical Perspectives* 16: 73–103.

Bach, K. (2006). Regressions in pragmatics (and semantics). In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 24–44). Basingstoke: Palgrave Macmillan.

Bach, K. (2008 forthcoming). Referentially used descriptions: a reply to Devitt. *European Journal of Analytic Philosophy*.

Bar-Lev, Z. and Palacas, A. (1980). Semantic command over pragmatic priority. *Lingua* 51: pp. 137–46.

Barwise, J. and Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4: 159-219.

Barwise, J. and Perry, J. (1981). Semantic innocence and uncompromising situations. In P. French et al. (Eds.), *Midwest Studies in Philosophy 6* (pp. 387-403). Minneapolis: University of Minnesota Press.

Bezuidenhout, A. (1997a). The communication of de re thoughts. *Noûs* 31: 197-225.

Bezuidenhout, A. (1997b). Semantics, pragmatics, semantic underdetermination, and the referential/attributive distinction. *Mind* 106: 375-409.

Bezuidenhout, A. (2002). Truth-conditional pragmatics. *Philosophical Perspectives* 16: 105-134.

Boolos, G. (1984). To be is to be a value of a variable (or to be some values of some variables). *Journal of Philosophy* 81: 430-49.

Borg, E. (2000). Complex demonstratives. *Philosophical Studies* 97: 229-49.

Braun, D. (1993). Empty names. *Noûs* 27: 449-69.

Braun, D. (1998). Understanding belief reports. *Philosophical Review* 105: 555-95.

Braun, D. (2001). Cognitive significance, attitude ascriptions, and ways of believing propositions. Presented at the Pacific APA meeting in San Francisco on March 31, 2001. Retrieved March 2, 2007, from http://www.ling.rochester.edu/~braun/Papers/cogsig.pdf

Burge, T. (1973). Reference and proper names. *Journal of Philosophy* 70: 425-39.

Cameron, J.R. (1999). Plural reference. *Ratio* 12: 128-147.

Cappelen, H. and Lepore, E. (2005). *Insensitive Semantics, A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Carston, R. (1988). Implicature, explicature, and truth-theoretic semantics. In R. M. Kempson, (Ed.), *Mental Representations: The Interface between Language and Reality* (pp. 155-81). Cambridge: Cambridge University Press.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction. *Synthèse*, forthcoming.

Carston, R. and Powell, G. (2006). Relevance theory—new directions and developments. In E. Lepore and B. Smith (Eds.), *The Oxford Handbook of Philosophy of Language* (pp. 341-60). Oxford: Oxford University Press.

Chastain, C. (1972). Reference and context. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science VII: Language, Mind, and Knowledge* (pp. 194-269). Minneapolis: University of Minnesota Press.

Chierchia, G. and McConnell-Ginet, S. (1990). *Meaning and Grammar*. Cambridge, MA: MIT Press.

Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger.

Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Clark, H. (1992). Making sense of nonce sense. In H. Clark, *Arenas of Language Use* (pp. 305-40). Chicago: University of Chicago Press.

Cohen, L.J. (1971). Some remarks on Grice's views about the logical particles of natural language. In Y. Bar-Hillel (Ed.), *Pragmatics of Natural Language* (pp. 50-68). Dordrecht: Reidel.

Collins, J. (2007). Syntax more or less. *Mind* 116 (464): 805-850.

Crimmins, M. and Perry, J. (1989). The prince and the phone booth. *Journal of Philosophy* 86: 685-711.

de Jong, F. and Verkuyl, H. (1985). Generalized quantifiers: the properness of their strength. In J. van Benthem and A. ter Meulen (Eds.), *Generalized Quantifiers in Natural Language* (pp. 21–43). Dordrecht: Foris.

Devitt, M. (1981). *Designation*. New York: Columbia University Press.

Devitt, M. (2004). The case for referential descriptions. In M. Reimer and A. Bezuidenhout (Eds.), *Descriptions and Beyond* (pp. 280–306). Oxford: Oxford University Press.

Devitt, M. and Sterelny, K. (1999). *Language and Reality*. Cambridge, MA: MIT Press.

Diesing, M. (1992). *Indefinites*. Cambridge, MA: MIT Press.

Di Sciullo, A.M. and Williams, E. (1987). *On the Definition of Word*. Cambridge, MA: MIT Press.

Dominicy, M. (1984). *La naissance de la grammaire moderne: langage, logique et philosophie à Port Royal*. Bruxelles: Pierre Mardaga.

Donnellan, K. (1966). Reference and definite descriptions. *Philosophical Review* 75: 281–304.

Ducrot, O. (1984). *Le dire et le dit*. Paris: Minuit.

Dummett, M. (1973). *Frege: Philosophy of Language*. Cambridge, MA: Harvard University Press.

Dummett, M. (1981). *The Interpretation of Frege's Philosophy*. Cambridge, MA: Harvard University Press.

Elbourne, P. (2005). *Situations and Individuals*. Cambridge, MA: MIT Press.

Englebretsen, G. (1997). *Something to Reckon With: The Logic of Terms*. Ottawa: University of Ottawa Press.

Evans, G. (1973). The causal theory of names. *Proceedings of the Aristotelian Society* 47: 187–208.

Fodor, J. (1975). *The Language of Thought*. New York: Cromwell.

Frege, G. (1879/1980). *The Foundations of Arithmetic*. Trans. by J.L. Austin. Evanston, IL: Northwestern University Press.

Frege, G. (1892a/1997a). On *Sinn* and *Bedeutung*. Trans. by M. Black. In M. Beaney (Ed.), *The Frege Reader* (pp. 151–71). Oxford: Blackwell.

Frege, G. (1892b/1997b). On concept and object. Trans. by M. Black. In M. Beaney (Ed.), *The Frege Reader* (pp. 181–93). Oxford: Blackwell.

Frege, G. (1914/1997). Letter to Jourdain, Jan. 1914. In M. Beaney (Ed.), *The Frege Reader* (pp. 319–21). Oxford: Blackwell.

Geach, P. (1962). *Reference and Generality*. Ithaca, NY: Cornell University Press.

Geurts, B. (1997). Good news about the description theory of names. *Journal of Semantics* 14: 319–48.

Glanzberg, M. (2006). Quantifiers. In E. Lepore and B. Smith (Eds.), *The Oxford Handbook of Philosophy of Language* (pp. 794–821). Oxford: Oxford University Press.

Green, M. (1998). Direct Reference and implicature. *Philosophical Studies* 91: 61–90.

Grice. H.P. (1957/1989). Meaning. *Philosophical Review* 66: 377–88. Reprinted in Grice, H.P. (1989), (pp. 213–23).

Grice, H.P. (1967/1989). Logic and conversation. William James lectures. Harvard University. Revised version in Grice, H.P. (1989), (pp. 1–143).

Grice, H.P. (1969/1989). Utterer’s meaning and intentions. *Philosophical Review* 78: 147–77. Reprinted in Grice, H.P. (1989), (pp. 86–116).

Grice, H.P. (1975/1989). Logic and conversation. In Cole, P. and Morgan, J. (Eds.), *Syntax and Semantics 3:*

Speech Acts (pp. 41–58). Reprinted in Grice, H.P. (1989), (pp. 22–40).

Grice, H.P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hampshire, S. (1959). *Thought and Action*. New York: Viking Press.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J. and May, R. (1981). Questions, quantifiers, and crossing. *Linguistics Review* 1: 41–79.

Hornsby, J. (1977). Singular terms in contexts of propositional attitudes. *Mind* 86: 31–48.

Horwich, P. (1998). *Meaning*. Oxford: Oxford University Press.

Justice, J. (2001). On sense and reflexivity. *Journal of Philosophy* 98: 351–64.

Kaplan, D. (1978). Dthat. *Syntax and Semantics* 9: 221–43.

Kaplan, D. (1989). Demonstratives. In J. Almog et al. (Eds.), *Themes from Kaplan* (pp. 481–564). Oxford: Oxford University Press.

Kaplan, D. (1990). Words. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 64: 93–119.

Katz, J. (1990). Has the description theory of names been refuted? In G. Boolos (Ed.), *Essays in Honor of Hilary Putnam* (pp. 31–61). Cambridge: Cambridge University Press.

Katz, J. (2001). The end of Millianism: multiple bearers, improper names, and compositional meaning. *Journal of Philosophy* 98: 137–66.

Keenan, E.L. and Stavi, J. (1986). A semantic characterization of natural language determiners. *Linguistics and Philosophy* 9: 253–326.

Keenan, E.L. and Westerståhl, D. (1997). Generalized quantifiers in linguistics and logic. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language* (pp. 837–93). Cambridge, MA: MIT Press.

King, J. (2001). *Complex Demonstratives: A Quantificational Account*. Cambridge, MA: MIT Press.

King, J. and Stanley, J. (2004). Semantics, pragmatics, and the role of semantic content. In Z. Szabó (Ed.), *Semantics vs. Pragmatics* (pp. 111–64). Oxford: Oxford University Press.

Kneale, W. (1962). Modality de dicto and de re. In P. Nagel et al. (Eds.), *Logic, Methodology and Philosophy of Science* (pp. 622–33). Stanford: Stanford University Press.

Kripke, S. (1971). Identity and necessity. In M. Munitz (Ed.), *Identity and Individuation* (pp. 135–64). New York: New York University Press.

Kripke, S. (1979). A puzzle about belief. In A. Margalit (Ed.), *Meaning and Use* (pp. 239–83). Dordrecht: Reidel.

Kripke, S. (1970/1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Larson, R. and Segal, G. (1995). *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

Levinson, S. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.

Lewis, D. (1970/1983). General semantics. In D. Lewis, *Philosophical Papers, Volume 1* (pp. 189–229). Oxford: Oxford University Press.

Lindström, P. (1966). First-order predicate logic with generalized quantifiers. *Theoria* 32: 186–95.

Link, G. (1983). The logical analysis of plurals and mass terms: a lattice-theoretical approach. In R. Bäurle, C. Schwarze, and A. von Stechow (Eds.), *Meaning, Use,*

and the Interpretation of Language (pp. 302–33). Berlin: Walter de Gruyter.

Loar, B. (1976). The semantics of singular terms. *Philosophical Review* 80: 43–62.

Lockwood, M. (1975). On predicating proper names. *Philosophical Review* 84: 471–98.

Marcus, R.B. (1961/1993). Modalities and intensional languages. *Synthèse* 13: 303–22. Reprinted in R.B. Marcus, *Modalities: Philosophical Essays* (pp. 5–23). New York: Oxford University Press.

Martinich, A.P. (1977). The attributive use of proper names. *Analysis* 37: 159–63.

McKinsey, M. (2005). Critical notice of Scott Soames, *Beyond Rigidity*. *Canadian Journal of Philosophy* 35: 149–68.

Mill, J.S. (1867/2001). Of names. In A.P. Martinich (Ed.), *The Philosophy of Language* (pp. 266–71). Oxford: Oxford University Press.

Mitchell, D. (1962). *An Introduction to Logic*. London: Hutchinson.

Montague, R. (1974). The proper treatment of quantification in ordinary English. In R. Thomason (Ed.), *Formal Philosophy* (pp. 247–70). New Haven: Yale University Press.

Moravcsik, J. (1998). *Meaning, Creativity, and the Partial Inscrutability of the Human Mind*. Stanford: CSLI Publications.

Mostowski, A. (1957). On a generalization of quantifiers. *Fundamenta Mathematicae* 44: 12–36.

Neale, S. (1990). *Descriptions*. Cambridge, MA: MIT Press.

Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy* 15: 509–59.

Neale, S. (1993). Term limits. *Philosophical Perspectives* 7: 89–114.

Neale, S. (2001). Implicature and Colouring. In G. Cosenza (Ed.), *Paul Grice's Heritage* (pp. 135–80). Rome: Brepols.

Neale, S. (2004a). This, that, and the other. In M. Reimer and A. Bezuidenhout (Eds.), *Descriptions and Beyond* (pp. 68–182). Oxford: Oxford University Press.

Neale, S. (2004b). Pragmatism and binding. In Z. Szabó (Ed.), *Semantics vs. Pragmatics* (pp. 165–285). Oxford: Oxford University Press.

Neale, S. (2007a). On location. In M. O'Rourke and C. Washington (Eds.), *Situating Semantics: Essays on the Philosophy of John Perry* (pp. 251–393). Cambridge, MA: MIT Press.

Neale, S. (2007b). Heavy hands, magic, and scene-reading traps. *European Journal of Analytic Philosophy* vol. 3(2): 77–132.

Neale, S. (2008). Term limits revisited. In J. Hawthorne and D. Zimmerman (Eds.), *Philosophical Perspectives: Philosophy of Language*. Forthcoming.

Partee, B. (1987). Noun phrase interpretation and type shifting principles. In J. Groenendijk et al. (Eds.), *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*. Dordrecht: Foris.

Peacocke, C. (1975). Proper names, reference, and rigid designation. In S. Blackburn (Ed.), *Meaning, Reference and Necessity* (pp. 109–32). Cambridge: Cambridge University Press.

Pelczar, M. and Rainsbury, J. (1998). The indexical character of names. *Synthèse* 114: 293–317.

Perry. J. (1986). Thought without representation. *Aristotelian Society Supplementary Volume* 60: 137–51.

Perry, J. (2001). *Reference and Reflexivity*. Stanford: CSLI Publications.

Plantinga, A. (1983). On existentialism. *Philosophical Studies* 44: 1–20.

Quine, W.V. (1940). *Mathematical Logic*. Cambridge, MA: Harvard University Press.

Quine, W.V. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Récanati, F. (1993). *Direct Reference: from Language to Thought*. Oxford: Basil Blackwell.

Récanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy* 25: 299-345.

Récanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Reimer, M. (1998). Donnellan's distinction/Kripke's test. *Analysis* 58: 89-100.

Reimer, M. (2003). Reference. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. Retrieved March 3, 2007, from http://plato.stanford.edu/entries/reference.

Rett, J. (2005). Context, compositionality, and calamity. *Mind and Language* 21: 541-552.

Rundle, B. (1965). Modality and quantification. In R.J. Butler (Ed.), *Analytical Philosophy, Second Series* (pp. 27-39). Oxford: Blackwell.

Russell, B. (1905). On denoting. *Mind* 14: 479-93.

Russell, B. (1910). Knowledge by acquaintance and knowledge by description. *Proceedings of the Aristotelian Society* 11: 108-128.

Russell, B. (1919). *Introduction to Mathematical Philosophy*. New York: Barnes and Noble Publishing.

Sainsbury, M. (2006). The essence of reference. In E. Lepore and B. Smith (Eds.), *The Oxford Handbook of Philosophy of Language* (pp. 393-421). Oxford: Oxford University Press.

Salmon, N. (1986). *Frege's Puzzle*. Cambridge, MA: Bradford Books/MIT Press.

Salmon, N. (1989). Illogical Belief. *Philosophical Perspectives* 3: 243–85.

Schiffer, S. (1981). Truth and the theory of content. In H. Parret and J. Bouveresse (Eds.), *Meaning and Understanding* (pp. 204–22). New York: de Gruyter.

Schiffer, S. (1987). The 'Fido'-Fido theory of belief. *Philosophical Perspectives* 1: 455–80.

Schiffer, S. (1995). Descriptions, indexicals, and belief reports: some dilemmas. *Mind* 104: 107–131.

Schiffer, S. (2006). A problem for a Direct Reference theory of belief reports. *Noûs* 40: 361–8.

Searle, J. (1958). Proper names. *Mind* 67: 166–73.

Searle, J. (1979). Literal meaning. In J. Searle, *Expression and Meaning* (pp. 117–36). Cambridge: Cambridge University Press.

Searle, J. (1980). The background of meaning. In J. Searle et al. (Eds.), *Speech Act Theory and Pragmatics* (pp. 221–32). Dordrecht: Reidel.

Sellars, W. (1954). Presupposing. *The Philosophical Review* 63: 197–215.

Soames, S. (2002). *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford: Oxford University Press.

Soames, S. (2004). Naming and asserting. In Z. Szabó, (Ed.), *Semantics vs. Pragmatics* (pp. 365–82). Oxford: Oxford University Press.

Soames, S. (2007). The gap between meaning and assertion: why what we literally say often differs from what our words literally mean. Forthcoming in M. Hackl and R. Thornton (Eds.), *Asserting, Meaning, and Implying*. Claremont, CA: Pomona College Press. Retrieved May 15, 2006, from http://www-rcf.usc.edu/~soames/forthcoming_papers.

Sperber, D. and Wilson, D. (1986/1995). *Relevance: Communication and Cognition*, Second Edition. Malden, MA: Blackwell.

Stanley, J. (1997). Names and rigid designation. In B. Hale and C. Wright (Eds.), *A Companion to the Philosophy of Language* (pp. 555–85). Oxford: Oxford University Press.

Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy* 23: 391–434.

Stanley, J. (2002a). Nominal restriction. In G. Preyer and G. Peter (Eds.), *Logical Form and Language,* (pp. 365–88). Oxford: Oxford University Press.

Stanley, J. (2002b). Making it articulated. *Mind and Language* 17: 149–168.

Stanley, J. and Szabó, Z. (2000). On quantifier domain restriction. *Mind and Language* 15: 219–261.

Strawson, P.F. (1950). On referring. *Mind* 59: 320–44.

Strawson, P.F. (1952). *Introduction to Logical Theory*. London: Methuen.

Szabó, Z.G. (2006). Sensitivity training. *Mind and Language* 21: 31–8.

Travis, C. (1981). *The True and the False*: *The Domain of the Pragmatic*. Amsterdam: John Benjamins.

Unger, C. (2001). *On the Cognitive Role of Genre: A Relevance-Theoretic Perspective*. Ph.D. dissertation, University of London.

von Fintel, K. (1994). *Restrictions on Quantifier Domains*. Ph.D. dissertation, Department of Linguistics and Philosophy, MIT.

Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem and A. ter Meulen (Eds.), *Generalized Quantifiers in Natural Language* (pp. 45–71). Dordrecht: Foris.

Westerståhl, D. (1989). Quantifiers in formal and natural languages. In D. Gabbay and F. Guenther (Eds.), *Handbook of Philosophical Logic, Vol. IV* (pp. 1–131). Dordrecht: Kluwer.

Wettstein, H. (1981). Demonstrative reference and definite descriptions. *Philosophical Studies* 40: 241–57.

Wilson, D. and Sperber, D. (1981). On Grice's theory of conversation. In Werth, P. (Ed.), *Conversation and Discourse* (pp. 155–78). London: Croom Helm.

Wilson, D. and Sperber, D. (2002). *Mind* 111: 583–632.

Wilson, D. and Sperber, D. (2004). Relevance theory. In G. Ward and L. Horn (Eds.), *Handbook of Pragmatics* (pp. 607–32). Oxford: Basil Blackwell.

Wittgenstein, L. (1953). *Philosophical Investigations*. Trans. by G.E.M. Anscombe. Oxford: Oxford University Press.

Ziff, P. (1960). *Semantic Analysis*. Ithaca, NY: Cornell University Press.