Lexical Meaning *in Context*

A WEB OF WORDS

NICHOLAS ASHER

CAMBRIDGE

Lexical Meaning in Context

A Web of Words

NICHOLAS ASHER
CNRS, Institut de Recherche en Informatique de Toulouse
and
University of Texas at Austin

CAMBRIDGE
UNIVERSITY PRESS

Copyrighted material

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9781107005396

© Nicholas Asher 2011

This publication is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without the written
permission of Cambridge University Press.

First published 2011

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this publication is available from the British Library

ISBN 978-1-107-00539-6 Hardback

Cambridge University Press has no responsibility for the persistence or
accuracy of URLs for external or third-party internet websites referred to in
this publication, and does not guarantee that any content on such websites is,
or will remain, accurate or appropriate.

Copyrighted material

Contents

Copyrighted material

Copyrighted material

Copyrighted material

Preface

Just over fifty years ago with the publication of "Two Dogmas of Empiricism", W. V. O. Quine launched a persuasive and devastating attack on the common sense notion of word meaning and synonymy, according to which two terms were synonymous just in case they had the same meaning. Quine's legacy continues to hold sway among much of the philosophical community today. The theory of word meaning is often thought either not to have a subject matter or to be trivial—*dog* means dog. What else is there to say? Well, it turns out, quite a lot. Linguists like Charles Fillmore, Igor Mel'cuk, Maurice Gross, Beth Levin, Ray Jackendoff, James Pustejovsky, and Len Talmy— to mention just a few, as well as researchers in AI who have built various on-line lexical resources like WORDNET and FRAMENET, have provided rich and suggestive descriptions of semantic relations between words that affect their behavior. And this has led to several proposals for a theory of word meaning.

Against this rich descriptive background, however, problems have emerged that make it not obvious how to proceed with the formalization of lexical meaning. In particular, something that is commonly acknowledged but rarely understood is that when word meanings are combined, the meaning of the result can differ from what standard compositional semantics has led us to expect: in applying, for instance, a property term ordinarily denoting a property P to an object term ordinarily denoting a, the content of the result sometimes involves a different but related property P' applied to an object b that is related to but distinct from the original denotation of a. While the choice of words obviously affects the content of a predication, the discourse context in which the predication occurs also affects it. The trick is to untangle from this flux a theory of the interactions of discourse, predication, and lexical content. That is what this book is about.[1]

[1] I owe many people thanks for help with this book: Alexandra Aramis, Alexis, Elizabeth, and Sheila Asher, Tijana Asic, Christian Bassac, David Beaver, Stephano Borgo, George

Copyrighted material

In this book, I argue that the proper way to understand the meaning of words is in terms of their denotations *and* the restrictions that other words impose on them. And it is the latter that govern how words interact semantically. I begin with the widely accepted observation according to which a predication will succeed only if the selectional restrictions the predicate imposes on its arguments are met. I provide an analysis of selectional restrictions by assigning words types. Meeting a selectional restriction is a matter of justifying a lexical presupposition, the presupposition that a term has a certain type. This analysis yields a theory of lexical meaning: to specify the type and the denotation of a word is to give its lexical meaning. The mechanisms of presupposition justification developed in dynamic semantics in recent years lead to an account of how predication adds content to the "ordinary" contents of the terms involved, which will provide my account of meaning shifting in context. The theory I will develop in this book has implications for compositional semantics, for example for the architecture of verbal and nominal modification. It also unifies analyzes in compositional semantics of presuppositions with my analysis of type presuppositions; for instance, the presuppositions of factive verbs or definite noun phrases are just special cases of type presuppositions.

The idea that there are non-trivial semantic interactions between words that affect the content of a predication is intuitive and perhaps obvious. But working out a precise theory, or even an imprecise one, of this phenomenon is difficult. I begin with some basic questions, distinctions, and observations.

What is a word? In some sense the answer is obvious: words are the things dictionaries try to define. On the other hand, the answer is not so simple. Words in many languages come with inflection for case, for number, for gender, among other things. Furthermore, there are morphological affixes that can transform one word into another like the nominalization affixes in English: an

Bronnikov, Robin Cooper, Denis Delfitto, Pascal Denis, Tim Fernando, Pierdaniele Giaretta, John Hawthorne, Mark Johnson, Hans Kamp, Chris Kennedy, Ofra Magidor, Alda Mari, Claudio Masolo, Bruno Mery, Friedericke Moltmann, Philippe Muller, David Nicolas, Barbara Partee, Sylvain Pogodalla, François Recanati, Christian Retoret, Antje Rossdeutscher, Sylvain Salvati, Magdalena Schwager, Stuart Schieber, Torgrim Solstad, Tony Veale, Laure Vieu, Kiki Wang, Laura Whitten, and participants of the seminars on lexical semantics at the University of Verona, the University of Texas at Austin, the University of Stuttgart, and the Summer Institute of the Linguistic Society of America at Stanford, where some of this material was presented. I want especially to thank Hans Kamp and members of the SFB 732 at the University of Stuttgart for their generous invitation to spend three months there to work on this project during the summer of 2008. Special thanks are also due to Julie Hunter and Renaud Marlet who reread much of the manuscript and offered many helpful comments and to James Pustejovsky, who got me to work on the subject of coercion and dot objects in the first place. Finally, I'd like to thank Andrew Winnard, Sarah Green, Gillian Dadd, Alison Mander, and Elizabeth Davey from Cambridge University Press for their help with the manuscript. The book is dedicated to Tasha, my darling little cat who didn't manage to live to see the end of this project.

Copyrighted material

affix like *-ion* turns a verb like *afflict* into the noun *affliction*. Morphological affixes and prefixes can often affect the meaning of a word; they can also determine how their host words combine with other words, as we shall see later on in this book. Even inflections like the plural are not always semantically innocent. Thus, the notion of a word quickly becomes a theoretical term; the meaningful parts of the lexicon may include things that we ordinarily would think of as bits of words, and basic word stems (the elements to which affixes and prefixes attach) may not end up looking like ordinary words at all.

Despite these complications, I will continue to speak (loosely) of words. What is it to give the meaning of a word? There are a number of answers in the literature on lexical semantics or theories of word meaning. Cognitive semanticists like Len Talmy and Tom Givon, among others, think that meanings are to be given via a set of cognitively primitive features—which might be pictorial rather than symbolic. According to these semanticists, a lexical theory should provide appropriate cognitive features and lexical entries defined in terms of them. Others in a more logical and formal framework like Dowty (1979) (but also Ray Jackendoff, Roger Shank, and other researchers in AI) take a specification of lexical meaning to be given in terms of a set of primitives whose meanings can be axiomatized or computationally implemented. Still others take a "direct" or denotational view; the function of a lexical semantics is to specify the denotation of the various terms, typically to be modelled within some model theoretic framework.

All of these approaches agree that a specification of lexical meaning consists in the specification of some element, whether representational or not, formal or not, that, when combined with elements associated with other words in a well formed sentence, yields a meaning for a sentence in a particular discourse context. Whatever theoretical reconstruction of meaning one chooses, however, it should be capable of modelling inferences in a precise manner so that the theory of lexical meaning proposed can be judged on its predictions. In addition, the theoretical reconstruction should provide predictions about when sentences that are capable of having a truth value are true and when they are not. This drastically reduces the options for specifying lexical meaning: such a specification must conform with one of the several ways of elaborating meaning within the domain of formal semantics; it must specify truth conditions, dynamic update conditions of the sort familiar from dynamic semantics (Kamp and Reyle (1993), Groenendijk and Stokhof (1991), Asher (1993), Veltman (1996)), or perhaps provability conditions of the sort advocated by Martin-Löf (1980) and Ranta (2004), among others.

For proponents of a direct interpretation of English, a denotational approach to lexical meaning suffices. Most semanticists, however, use a logical language

Copyrighted material

PART ONE

FOUNDATIONS

Copyrighted material

Copyrighted material

1

Lexical Meaning and Predication

To build a formal model of predication and to express lexical meaning, I will use the lambda calculus. The lambda calculus is the oldest, most expressive, and best understood framework for meaning representation; and its links to various syntactic formalisms have been thoroughly examined from the earliest days of Montague Grammar to recent work like that of de Groote (2001), Frank and van Genabith (2001). Its expressive power will more than suffice for our needs.[1]

The pure lambda calculus, or λ calculus, has a particularly simple syntax. Its language consists of variables together with an abstraction operator λ. The set of terms is closed under the following rules: (1) if v is a variable, then v is a term; (2) if t is a term and v a variable, then λvt is also a term; (3) if t and t' are terms, then *the application* of t to t', $t[t']$, is also a term. We can use this language to analyze the predication involved when we apply a predicate like an intransitive verb to its arguments. The meaning of an intransitive verb like *sleeps* is represented by a lambda term, $\lambda x\, \text{sleep}'(x)$; it is a function of one argument, another term like the constant j for *John* that will replace the λ bound variable x and yield a logical form for a larger unit of meaning under the operation of β reduction. β reduction, also known as β conversion, is a rule for inferring one term from another. β reduction is the formal counterpart in the λ calculus of the informal operation of predication. One can also think of reduction as the rule governing application, and so I shall call it the rule of *Application*.[2] I'll write such a rule in the usual natural deductive format.

[1] There are other formalisms that can be used—for instance, the formalism of attribute value matrices or typed feature structures with unification. This formalism, however, lacks the operation of abstraction, which is crucial for my proposals here.

[2] Besides Application, there are other rules standardly assumed for the λ calculus—for example, α conversion, which ensures the equivalence of bound variables, rules for equality, and the following rules which validate a rule of Substitution that I shall introduce subsequently:

Copyrighted material

- Application:

$$\frac{\lambda x \phi[\alpha]}{\phi(\frac{\alpha}{x})}$$

The λ calculus as our representational language tells us in principle what our lexical entries should look like. For example, if we decide that a word like *cat* is a one place predicate, then our lexical entry for this word should have the form $\lambda x\, \mathrm{cat}'(x)$, where *cat*′ is an expression in our language for logical forms that will, when interpreted, assign the right sort of denotation to the word and contribute to the right sort of truth conditions for sentences containing the word. Of course, there are lots of decisions to be made as to what *cat*′ should be exactly, but we will come back to this after we have taken a closer look at predication.

1.1 Types and presuppositions

Sometimes predications go wrong. This is something that lexical semantics has to explain.

(1.1) a. ?That person contains an interesting idea about Freud.
b. That person has an interesting idea about Freud.
c. That book contains an interesting idea about Freud.
d. That person is eating breakfast.
e. That book is red.
f. #That rumor is red.
g. # The number two is red.
h. # The number two is soft.
i. # The number two hit Bill.
j. The number two is prime.
k. John knows which number to call.
l. *John believes which number to call.

The predications in (1.1f,g,h) or (1.1i) are malformed—each contains what Gilbert Ryle would have called a category mistake. Numbers as abstract objects can't have colors or textures or hit people; it's nonsensical in a normal

- $t = t' \rightarrow t[t''] = t'[t'']$
- $t = t' \rightarrow f[t] = f[t']$
- $t = t' \rightarrow \lambda x\, t = \lambda x\, t'$

Church (1936) shows how to encode Boolean functions within the λ calculus, once we have decided on a way of coding up truth functions.

Copyrighted material

conversation to say something like the number two is red, soft, or that it hit Bill.[3] The mismatch between predicate and argument is even more blatant in (1.1l).

One has to exercise some care in understanding why a predication like (1.1a) sounds so much odder than (1.1b–d). In some sense people can contain information: spies have information that they give to their governments and that counter-spies want to elicit; teachers have information that they impart to their students. But one can't use the form of words in (1.1a) to straightforwardly convey these ideas. The predication is odd; it involves a misuse of the word *contain*. If it succeeds at all in making sense to the listener, it must be subject to reinterpretation.

It's important to distinguish between necessary falsity and the sort of semantic anomaly present in (1.1a) and (1.1f–i). In the history of mathematics, many people, including famous mathematicians, have believed necessarily false things. But competent speakers of a language do not believe propositions expressed by a sentence with a semantically anomalous predication. (1.1a) or (1.2c,d) are semantically anomalous in a way that (1.1b–d) or (1.2a,b) below are not.[4]

(1.2) a. Tigers are animals.
b. Tigers are robots.
c. #Tigers are financial institutions.
d. #Tigers are Zermelo-Frankel sets.

Many philosophers take (1.2a) to be necessarily true and (1.2b) to be necessarily false.[5] Nevertheless, according to most people's intuitions, a competent speaker could entertain or even believe that tigers are robots; he or she could go about trying to figure this out (e.g., by dissecting a tiger). It is much harder to accept the possibility, or even to make sense of, a competent speaker's believing or even entertaining that tigers are literally financial institutions, let alone ZF style sets. Thinking about whether a competent speaker could entertain or believe the proposition expressed by a sentence gives us another means to distinguish between those sentences containing semantically anomalous expressions and those that do not.

[3] As the attentive reader may have already guessed, besides "normal" conversations, there are also "abnormal" discourse contexts—contexts that would enables us to understand these odd sentences in some metaphorical or indirect way, or that even enable us to reset the types of words. More on this later.

[4] Thanks to Dan Korman for the first two examples.

[5] The reason for this has to do with a widely accepted semantics of natural kinds due to Hilary Putnam and Saul Kripke, according to which *tigers* picks out a non-artifactual species in every possible world.

Copyrighted material

when embedded under negation, interrogative mood or modal operators. This is indeed the case:

(1.4) a. # The number two could have been red.
b. # Is the number two soft?
c. # The number two didn't hit Bill.[7]

The sentences in (1.4) all convey presuppositions that are absurd and that cannot be met—namely, that the number two is a physical object. Other tests for presuppositions concern the non-redundancy of presupposed content and the inability to make certain discourse continuations on presupposed content.[8] These tests apply to type requirements of predicates as well. It is not redundant to say *the abstract object two is prime* instead of *two is prime*, and it seems impossible to make discourse continuations on the type requirements, since the latter are not even propositional contents. Thus, it seems that the type requirements of predicates provide a kind of presupposed content. I shall call these *type presuppositions*.

Two features of presuppositions will be very important for the study of predication in this book. The first is the variability among terms that generate presuppositions to license accommodation. It is standardly assumed that the adverb *too* generates a presupposition that must be satisfied in the given discourse context by some linguistically expressed or otherwise saliently marked content. Thus, in an "out of the blue" context, it makes no sense to say

(1.6) Kate lives in New York too.

even though as a matter of world knowledge it is clear that the presupposition of *too* in this sentence is satisfied—namely, that there are other people besides Kate who live in New York. Even if the proposition that there are other people besides Kate who live in New York is manifestly true to the audience of (1.6), (1.6) is still awkward, unless the presupposed content has been made salient somehow in the context. The presupposed, typing requirements of the predicates in (1.1) and (1.4) resemble the behavior of the presupposition of *too*; they

[7] A presuppositional view should allow that this sentence has a perfectly fine reading where the negation holds over the type requirements as well. But typically such readings are induced by marked intonation. If this sentence is read with standard assertion prosody, then it is as anomalous as the rest.

[8] The continuation test says that one cannot elaborate or explain or continue a narrative sequence on presupposed content. Thus, one cannot understand the example below as conveying that John regretted that he yelled at his girlfriend and that then after fighting with her he went to have a drink.

(1.5) John regrets that he fought with his girlfriend. Then he went to have a drink.

Copyrighted material

have to be satisfied in their "predicative" context in order for the sentences containing them to receive a truth value. Accommodation of these type presuppositions is impossible. The sentences that fail to express a coherent proposition capable of having a truth value do so, because the relevant type presuppositions cannot be satisfied, given that the arguments and predicates therein mean what they standardly mean and have the types that they standardly do.

On the other hand, some presupposition introducing phrases like possessive DPs readily submit to accommodation. For instance, *Sylvain's son* presupposes that Sylvain has a son, but this information is readily accommodated into the discourse context when the context does not satisfy the presupposition.

(1.7) Sylvain's son is almost three years old.

Other definite descriptions can be satisfied via complex inferences. The example below, which features a phenomenon known as "bridging," features such an inference; the definite *the engine* is "satisfied" by the presence of a car in the context—the engine is taken to be the engine of the car:

(1.8) I went to start my car. The engine made a funny noise.

In the following chapters we will see cases of type presuppositions that can either be satisfied in complex ways like the bridging cases or can be accommodated via a "rearrangement" or modification of the predicative context, if the latter fails to satisfy the type presuppositions in a straightforward way. Figuring out when presupposed typing requirements can be accommodated and when they cannot will be a central task of this book.

Another important property of presuppositions is their sensitivity to discourse context. For instance, if we embed (1.7) in the consequent of a conditional, the presupposition that projects out from the consequent can be bound in the antecedent and fails to project out further as a presupposition of the whole sentence (1.9):

(1.9) If Sylvain has a son, then Sylvain's son is almost three years old.

A similar phenomenon holds for type presuppositions. Consider (1.4a) embedded as a consequent of the following (admittedly rather strange) counterfactual.

(1.10) If numbers were physical objects, then the number two could have been red.

The presupposition projected out from (1.4a) is here satisfied by the antecedent of the counterfactual and rendered harmless. Thus, category mistakes for the most part must be understood relative to a background, contextually supplied set of types, a background that may itself shift in discourse.

Copyrighted material

1.2 Different sorts of predication

Having introduced types as part of the apparatus of predication, let me come back to predication itself. I have spoken so far of predication as a single operation of applying a predicate to its arguments. But in fact predication takes many forms in natural languages, some particular to particular languages, others more general. Even among ordinary predications, linguists distinguish between:

- predication of a verb phrase to a subject or a transitive verb to an object
- adjectival modification with different types of adjectives—e.g., evaluative adjectives like *good rock*, *bad violinist*, material adjectives like *bronze statue*, *paper airplane*, and manner adjectives like *fast car*, *slow cigar*
- adverbial modification and modification of a verb phrase with different prepositional phrases or PPs—e.g., the distinction between *load the wagon with hay* and *load the hay on the wagon*.

Beyond these are more exotic forms of predication:

- metaphorical usage (extended predication)

  (1.11) John is a rock.

- restricted predication

  (1.12) John as a banker makes $50K a year but as a plumber he makes only $20K a year.

- copredication

  (1.13) The lunch was delicious but took forever.

  (1.14) The book has a purple cover and is the most intelligible introduction to category theory.

  (1.15) #The bank is rising and specializes in IPOs.

- loose predication

  (1.16) That's a square (pointing to an unpracticed drawing in the sand).

- resultative constructions

  (1.17) a. Kim hammered the metal flat.
  b. * Kim hammered the metal gleaming.

  (1.18) depictives

  a. Pat swims naked.
  b. *Pat cooks hot.

- the genitive construction

Copyrighted material

(1.19) a. Kim's mother
b. Kim's fish

- noun noun compounds

(1.20) a. lunch counter
b. party favor

Each one of these forms of predication presents its own challenges for lexical and compositional semantics; the lexical theory must assign to the words in these constructions the right sort of meaning and postulate the right sort of composition rules for predication so as to get the right result. In addition, a lexical theory must specify what morphological processes and elements affect meaning and how; it must give those processes and elements a meaning. A lexical theory using the typed lambda calculus can provide the right sort of picture to tackle these issues.

Let's consider these forms of predication in a bit more detail. Loose predication is a difficult and well-known problem in philosophy.[9] But other forms of predication mentioned above, which linguists think also provide challenges for lexical theory, have not received so much philosophical scrutiny or formal analysis. Copredication, for instance, which is a grammatical construction in which two predicates jointly apply to the same argument, has proved a major challenge. Languages, as we shall see in the next chapter, distinguish between events and objects; the predicates that apply the one type do not apply in general to the other type literally. It turns out that some objects, however, are considered both events *and* physical objects in some sense. Consider, for instance, lunches. Lunches can be events but they are also meals and as such physical objects. As a result, *lunch* supports felicitous copredications in which one predicate selects for the event sense of *lunch* while the other selects for the physical object or meal sense.

(1.21) Lunch was delicious but took forever.

It turns out that many words behave like *lunch* in (1.21) and denote objects with multiple senses or aspects. I will call predications like those in (1.21) *aspect selections*, and I will analyze these predications as predications that apply to selected aspects of the object denoted by the surface argument.

In trying to account for instances of copredication that involve aspect selection like (1.21), standard, typed theories of predication and lexical semantics confront some difficult if not unanswerable questions. How can a term have two incompatible types, as is apparently the case here? How can one term

[9] Loose predication is related to vagueness, and vague predication might be considered another form to be studied. But I shall not do that here.

Copyrighted material

denote an object or set of objects to which apply two properties demanding different, even incompatible types of their bearers? It would have to be the case then that such an object must have, or belong to, two incompatible types. But how is that possible? Proponents of standard type theory have only one clear recourse, and that is to claim that terms associated with two incompatible types are ambiguous. But that deepens the mystery about copredications involving aspect selections: if *lunch* in (1.21) is ambiguous between a physical object reading and an event reading, then we must disambiguate the term in one way to make sense of the first predication but disambiguate it in a second way to make sense of the second predication; and the problem is that, on the surface at least, we have only one term to disambiguate—we have to choose a disambiguation, but such a choice will inevitably cause one of the predications in (1.21) to fail. At this point we might try a strategy of desperation and postulate a hidden "copy" of the problematic term, rewriting (1.21) in effect as

(1.21′) Lunch was delicious but lunch took forever.

This copying strategy now allows the proponent of standard type theory to proceed to disambiguate the two occurrences of *lunch* in different ways allowing the two predications to succeed. But the promise of the copying strategy is shortlived. Copying expressions will get us incorrect truth conditions in many cases. Consider (1.22), where *last* applies to events, while *tasted* applies only to objects (you can't taste events except metaphorically):

(1.22) A lunch was gingerly tasted by Mary and then lasted three hours.

The copying strategy forces us to interpret (1.22) as

(1.23) A lunch was gingerly tasted by Mary and then a lunch lasted three hours.

It's easy to see that (1.22) and (1.23) have different truth conditions; (1.22) is true only in those situations where Mary gingerly tasted the same lunch that lasted three hours, while (1.23) can be true in situations where Mary gingerly tastes one lunch but another lunch lasts three hours. Thus, it is not obvious how to deal with examples of copredication even from the standpoint of compositionality, if one's lexical theory produces a rich system of types. Montague himself noted that there were copredications that were puzzling even within his much more impoverished system of types. In (1.24) *temperature* seems to have two aspects, one of which is a number on a scale, while the other is a function from times to numerical values.

(1.24) The temperature is 90 and rising.

Copyrighted material

activity involving the subject of the aspectual verb serves as its internal argument. Thus, we understand *John started the car* as John's started the running of the car's engine. To start a cigarette is typically to start to smoke a cigarette. (1.27d) shows that coercions can happen with prepositional phrases—*the dress with the flowers* has at least one interpretation where a representation of flowers is stitched, printed, or drawn on the fabric of the dress, while *the garden with the flowers* does not have that interpretation, at least not nearly so saliently.[11] As we shall see, there are subtle differences with respect to the presuppositions in the typing requirements of various aspectual verbs and other coercing predicates.

Are such coercions really part of lexical *semantics*? That is, is it a defeasible but *a priori* inference that if John started the car, John started the engine of the car or that if Julie enjoyed the book, then (defeasibly) she enjoyed reading it? Do such inferences follow solely from one's linguistic mastery of the language? Fodor and Lepore think that none of these inferences belong to lexical semantics but are rather part of encyclopaedic or world knowledge. However, most people can distinguish between the largely automatic interpretations that these predications seem to entail and those that require more conscious effort. One might take that to be a mark of the information as being present even during predication rather than inferred afterwards using background, nonlinguistic beliefs.

It is notoriously difficult to distinguish between what is properly a part of lexical meaning and what is world knowledge. Quine's attack on lexical meaning can be seen as starting from the point that one cannot make this distinction in a principled way. Part of the difficulty is that, to some extent, the division between word meaning and world knowledge is a theory-internal distinction. For instance, if you're an externalist for whom the meanings of two singular or natural kind terms t and t' are determined by their reference, it may be a fact of meaning that $t = t'$ or not. Thus, *water is* H_20 would be a fact of meaning, and hence analytic, on such a view! Despite these difficulties, there is are tests one can use to see whether it is certain information conventionally associated with particular word meanings rather than just general world knowledge that gives rise to these inferences. For one thing, it seems pretty clear that the infierences given in (1.27c) are tied to particular predicates, particular verbs. Let's suppose that *cigarette*, like *lunch*, always has associated with it a possible event reading. It should then be possible to access that appropriate event reading with other predicates that take events.

[11] This example is due to Marliese Kluck.

Copyrighted material

(1.28) a. Nicholas's smoking of that cigarette will begin in 2 minutes.
b. Nicholas's cigarette break will begin in 2 minutes.
c. ??Nicholas's cigarette will begin in 2 minutes.

It's quite clear that (1.28c) is semantically strange. The event associated with *cigarette* in *enjoy the cigarette*, *begin the cigarette*, *finish the cigarette*, just isn't available with other event predicates. This strongly suggests that there is some particular conventional meaning that issues from the predication of the properties these verbs denote to the objects denoted by their arguments that isn't available in other predicational contexts. That is, the eventuality of smoking isn't just accessible with any predication involving *cigarette*; it is the result of combining *cigarette* as an object or internal argument of an aspectual verb or a verb like *enjoy*.

It is not only that such inferences are tied to particular verbs; they are tied to them independently of what the verb's object is. Consider the use of a nonsense word like *zibzab*. [12] To say

(1.29) John enjoyed the zibzab

is to say that John enjoyed doing something to the zibzab. At this point, it's really hard to understand how the inference to an event reading is part of world knowledge. It becomes clear that it is an *a priori* truth that to enjoy something is to be involved in some interaction with it—some eventuality. When the direct object argument of a verb like *enjoy* does not denote an eventuality as part of its standard meaning, coercion introduces somehow an appropriate eventuality. This militates strongly for placing coercion within the realm of linguistic knowledge, not contingent factual information about what the world is actually like.

A further question concerns how this eventuality involved in predications like those in (1.27c) is specified. While we'll see that this is not true in all cases, some nouns like *cigarette* help to specify the eventuality induced by coercion when they are in the direct object of a coercing verb. To see this, consider replacing the word *cigarette* with the relevant part of its entry in *Webster's New World College Dictionary*, which should at least roughly have the same content as the word *cigarette*.

(1.30) a. Nicholas enjoyed a cigarette.
b. Nicholas enjoyed a small roll of finely cut tobacco wrapped in thin paper.

[12] Thanks to Chris Kennedy for these sorts of examples.

Copyrighted material

Speakers immediately get the defeasible interpretation of (1.30a) where Nicholas smokes a cigarette but not in the second case. These alternations appear pretty systematic with a small class of words like *book*, *novel*, *sonata*, and so on, indicating that indeed such defeasible interpretations are a part of lexical semantics.[13]

Coercions, as Aristotle said of all familiar things, are easy to see but hard to understand. Some linguists have argued that in fact coercions are what they appear to be. They indicate that the meaning of terms fluctuates from context to context, and some have taken the moral of these observations to be some sort of radical contextualism about meaning (for instance, Recanati (2004, 2002, 2005)). Pustejovsky (1995) also seems to endorse such a view in emphasizing the "generativity" of the lexicon. But these conclusions do not follow from the evidence. They are also vastly counterintuitive: when I say that I enjoyed the cigarette, does the word *cigarette* now all of a sudden change its meaning to mean *smoking a cigarette*? It does not seem so. Fodor and Lepore (1998) and more recently Cappelen and Lepore (2004) correctly, in my view, crticize such an approach to coercion in their criticisms of Pustejovsky (1995). Furthermore, despite many claims that the lexicon is in fact generative or context sensitive in some radical way, I do not know of any formally worked out proposal of this view.[14] As we shall see in chapter 3, the generative lexicon of Pustejovsky (1995) has static lexical entries that do not change during coercion. When it comes to technical developments, I shall show that in fact basic word meanings cannot change if we are to be able to derive any predictions at all about lexical meaning.

If that approach to coercion is wrong, however, what is the right approach? Coercion is a ubiquitous, attested phenomenon in natural language. One has to be able to give an analysis of it in any remotely viable theory of predication and lexical meaning. In order to explain the data one has to do one of two things: either one has to develop a theory of lexical meaning where the lexical entries themselves change in context, *or* one has to complicate one's notion of predica-

[13] Laura Whitten and Magda Schwager independently observed to me that using such dictionary definitions is problematic, because Gricean maxims would predict that there are some special reason that the more complex formulation is used. This special reason would block the standard associations with the content. But if the inference here concerns nonlinguistic knowledge, we wouldn't expect the flouting of the Gricean maxim to block such an inference. A Gricean explanation of why the inference fails for (1.30b) occurs precisely because that inference is based on lexical content, *not* world knowledge.

[14] Although one possibility for formalisation would be in a connectionist approach, where word meanings are thought of as vectors of strengths of associations with other words that get recalculated every time the word occurs. This is very far from either Pustejovsky's generative lexicon or Recanati's relevance theory approach. It's also philosophically and conceptually extremely unsatisfying, as such an approach doesn't begin to tell us anything about lexical meaning or lexical inference or about how meanings compose together.

Copyrighted material

tion and logical form. Since I do not see how the first option can be developed in any detail that does justice to the composition of sentential or discourse meaning from lexical content, I will investigate the alternative that coercion phenomena call for a reanalysis of predication. When I say that I enjoyed a cigarette, the word *cigarette* does not change its meaning but what I enjoyed is doing something with the cigarette. That is, coercions involve a more complex act of predication than one might have thought. The reason that this is so, I will argue, is that in coercions type presuppositions have to be accommodated and, as with accommodation at the level of contents, type accommodation typically introduces new material into logical form.[15] Viewed from this perspective, coercion is not really a problem about meaning change in the lexicon; it's a problem about compositionality—about how lexically given meanings combine together in the right sort of way. I argue for a similar conclusion for dual or multiple aspect nouns like *lunch*, *book*, *temperature*, and so on; the process of justifying type presuppositions with the types assigned to these nouns will complicate the process of logical form construction and may add new content to logical form. To account for many features of predication, the logic of meaning composition has to be rethought and revamped considerably from the standard approach to predication that underlies Montague Grammar. This is the task to which I devote myself for most of this book.

1.3 The context sensitivity of types

We want a theory of lexical information that offers a framework within which empirical research will yield a correct account of lexical content. I have argued that a lexical theory has to do two things to reach this goal: give an account of the meanings of lexical items and an account of the operation of predication needed to derive truth conditions for clauses. But in order to capture the observations and intuitions of most of those who have worked in the field of lexical semantics, we need to do this in a particular way: we need to construct a theory of lexical meaning and predication that can exploit features of the discourse context. In particular this means a context sensitive theory of typing.

Observations that confirm this last claim have been around for years. Nevertheless, there have been few attempts in the literature to account for these observations. Lascarides and Copestake (1995) noticed that the event readings of the object of a verb like *enjoy* can depend on discourse factors. Normally *enjoy* coerces its object or theme argument that has the type BOOK into an expression

[15] GL says little about predication. Because it fails to carry through on either way of analyzing the data, we shall see that GL fails to account for coercion.

Copyrighted material

that involves an event of reading the book as in (1.31a). But this reading depends on the assumption that Julie is a person, someone capable of reading the book, because this reading is not available with (1.31b). It becomes available if (1.31b) occurs in the context, say, of a fairy tale in which goats have been established as capable of reading. This is a matter of *linguistic knowledge*, of how the discourse context affects the analysis of (1.31b), *not* a matter of world knowledge—it is completely unintuitive to assume as a general matter of world knowledge that goats can read in fairy tales (it has to depend on the particular fairy tale!).

(1.31) a. Julie enjoyed the book.
b. The goat enjoyed the book.

Most stories about coercion (Pustejovsky (1995), Nunberg (1979), Egg (2003), Asher and Pustejovsky (2006)) assume that the object argument of *enjoy* is some sort of eventuality, which is the result of a typing adjustment due to a clash between the type of argument the verb demands and the type of argument that is in fact its direct object; the verb *enjoy* requires an event as object argument and so coerces the direct object into giving an argument that is of some event type. Regardless of the details of how this coercion process actually works, (1.31a,b) shows two things. First, the inference from *enjoy the book* to *enjoy reading the book* must be defeasible. Second, the fact that we can get the reading that the goat enjoyed the book in (1.31b), given a discourse context in which goats talk, shows that the typing and typing adjustment rules must be sensitive to information in discourse.

Danlos (2007) has shown that aspectual verbs are also sensitive to discourse context. Aspectual verbs take some sort of eventuality as an object or theme argument. The GL framework claims that these eventualities are given by the lexical entries of nouns. But Danlos's examples show that this is not the case:

(1.32) a. ??Yesterday, Sabrina began with the kitchen. She then proceeded to the living room and bedroom and finished up with the bathroom.
b. Yesterday Sabrina cleaned her house. She began with the kitchen. She then proceeded to the living room and bedroom and finished up with the bathroom.
c. Last week Sabrina painted her house. She started with the kitchen. She then proceeded to the living room and bedroom and finished up with the bathroom.

The examples in (1.32) show that the eventuality is *not*, at least in all cases, given by the lexical entry of a noun in the theme argument of the verb or some

Copyrighted material

theory of predication and lexical meaning and an account of how these interact with discourse structure in this book. I will use the theory to investigate phenomena like copredication, restricted predication, and coercion.

The guiding idea, already implicit to some extent in Pustejovsky (1995), but made much more explicit in Asher and Pustejovsky (2006), is that almost all words will have single and simple lexical entries. Words like the nouns *cat*, *lunch*, *book* or the verbs *kill*, *read*, and *master* denote simple properties or relations, and so accordingly the logical forms that specify their denotations are very simple. For instance, *cat* has the lexical entry $\lambda x\ \mathrm{cat}(x)$, while *master* is a simple transitive verb that denotes a relation between its subject (or external) and object (or internal) arguments.

If words have simple entries and make simple contributions to truth conditional content, they come with a rich amount of information about the types assigned to the lambda terms and the variables within them. These types will guide predication and be responsible for fine-grained differences in lexical meaning. When words are combined together to form clauses, sentences, and discourses, the types associated with various terms will interact in complex ways. I will introduce operations of type adjustment in response to type mismatches between predicate and argument that correspond to the accommodation of a type presupposition, or more generally speaking the justification of such a presupposition. I will tell a similar story about morphological processes with semantic import; in effect the application of semantically rich morphology to a word stem is also a matter of type driven predication, which may bring with it various type adjustments.

The effects of these type adjustments at logical form is that the logical form for the clause will contain elements that are not present in the lambda terms for the constituent words themselves. Predication involves not only applying a function to an argument but also operations of adjustment corresponding to type presupposition justification. Coercion and the sort of problematic copredications that I introduced earlier will invoke particular sorts of presupposition justification. As we will see, however, predication is not simply a matter of putting well formed lexical meanings together and adjusting them when they do not fit; type information is also to some extent dependent on the discourse context. The theory I will present will provide a framework for investigating the context sensitivity of type assignments.

While I will use the framework of the typed lambda calculus, I will extend the typed lambda calculus beyond the usual set of simple types and functional types to include two complex types. These complex types furnish the basis for my analysis of problematic cases of copredication and coercion. One complex type, the $\bullet$ type, will be used to analyze terms of which we can predicate

Copyrighted material

properties of two different aspects of the same thing (what I call *aspect selection*). Most of these terms are nouns, and I will call them *dual aspect nouns*. I will argue that this kind of predication requires a special metaphysical conception of the objects whose aspects are the bearers of the properties predicated. The type adjustment with dual aspect nouns is, in some sense, just a shifting of emphasis or a reconceptualization of the very same object. Coercions like those induced by verbs like *enjoy*, shift the predication entirely away from the original object to some other object of a different type—for instance, an eventuality associated with the original object. Such coercions, unlike aspect selection, do not affect the way the objects denoted by the term that is subject to the coercion are counted or individuated. I will model coercions with another sort of complex type, something that I shall call a *polymorphic* or *dependent type*.

I will provide rules for dealing with the types I introduce in the analysis of predication when they occur either as type presuppositions of predicates or as types on arguments of predicates that require some other type. Rules for type presupposition justification will allow us to select the appropriate type of the argument for the requirements of the predication at hand. In addition I will show how the type system is sensitive to the discourse context. By integrating a theory of discourse structure and discourse contexts from earlier work (Asher (1993), Asher and Lascarides (2003)) within the theory of predication, I will show how discourse can transform and constrain type assignments and type transformations. The type system is dynamic and evolves as discourse proceeds in a way similar to the way that linguists and philosophers have argued that the semantics of discourse evolves dynamically.

My approach crucially distinguishes the logical forms constructed during predication from the types that guide and constrain predication. When we accommodate, say, a type presupposition of a predicate that demands an event as an argument but is given something of type PHYSICAL OBJECT, justifying the presupposition will require not only an adjustment in types but, typically, an adjustment at the level of logical form as well. We need types to construct logical form, but we also need the logical forms as distinct semantic citizens, for it is they, not the types, that are the vehicles of model theoretic content. I will argue for a "two stage" or two level semantics for lexical meaning: a level with the usual intensions for the expressions of logical form, and a level with a proof theoretic semantics for the types.

The nature of types and the argument for my two level theory will occupy much of the next chapter. The development of the system of complex types and the two-stage theory of semantics and its applications will occupy the rest of the book. Besides the analyzes of copredication and coercion, I will show how the system yields an analysis of restricted predication and the genitive

Copyrighted material

construction. Though topics like metaphor and poetic license outrun the scope of this book, I will tentatively offer an application of my system to these topics at the end of the book.

We'll see that types in the theory of predication are closely linked to metaphysical principles of individuation and counting. Thus, the types used to guide predication will be of a quite general nature. The system of types, however, involves more types than those just needed for checking predication. Fine-grained differences in types can affect the content of a predication and can account for at least some analytical entailments—entailments that are *a priori* and follow from the meanings of the "non logical" words of natural language. More speculatively, types will provide a linguistic foundation for a theory of concepts and of internally available contents and inferences. The approach is thus anti-Quinean (and also contra Fodor and Lepore (1998), Cappelen and Lepore (2004)). Nevertheless, I have taken to heart the warnings of Fodor and Lepore about lexical semantics. For instance, my type-driven theory of predication is agnostic about lexical decomposition beyond what is demanded by morphology and syntax.

Copyrighted material

2

Types and Lexical Meaning

The typed lambda calculus and its operation of type restricted application are familiar to anyone who has worked in formal semantics. But there are largely unexamined questions about the nature of types, their relations to formulas of logical form, and the effect of rules of type shifting on logical form. We need to look at these questions in detail.

2.1 Questions about types

Let us first turn to questions about types. For one thing, are our types all atoms or are there types that have a structure and that are constructed from "type constructors" together with other types? Another question is, what is the interpretation of our types?

Montague Grammar has an answer to our questions. Montague Grammar starts with two basic types, the type of entities E and the type of truth values T and then closes the collection of types under the recursive rule that if a and b are types, then so is $a \Rightarrow b$.[1] The type $a \Rightarrow b$ is one that, given an argument of type a, produces an object of type b. Montague Grammar converts these extensional types into intensional types as follows: if a is an extensional type, then $s \Rightarrow a$ is its intensional correlate, where s is the type of worlds or more generally indices of evaluation. Montague Grammar has an extensional set theoretic model of types: the primitive types are identified by their *inhabitants*, the set of objects of that type relative to some domain of interpretation, while the set

[1] The notation $\rightarrow$ for the functional type constructor is standard. But as I will distinguish this functional type constructor from its related cousin, the implication constructor for logical forms, $\rightarrow$ and from the usual way of defining a function ($f: a \rightarrow b$), I will use the slightly nonstandard $\Rightarrow$.

Copyrighted material

of functional types over types a and b is modelled as the function space or set of all functions from a into b, $\{f: a \rightarrow b\}$.

When Montague developed his theory, his use of the typed lambda calculus served a logical purpose. Turing (1937) had shown that the untyped lambda calculus had a model in the set of computable functions, but the application of such a theory in formal semantics was problematic. When terms of the untyped lambda calculus include the standard truth functional operators essential to semantics, it is easy to form terms like $\lambda x \neg x[x]$, which is the property of not applying to oneself—the Russell Property. Applying the Russell property to itself (note that we're using untyped Application here) produces the following result, which is uninterpretable when negation is understood as in classical logic:

(2.1) $\lambda x \neg x(x)[\lambda x \neg x(x)] = \neg(\lambda x \neg(x)[\lambda x \neg x(x)])$

The typed lambda calculus avoids this problem, since the Russell Property does not have a consistent type in the typed lambda calculus. It is, in other words, not a well-formed term. The typing of expressions allows one to combine the lambda calculus with the operators of classical logic; it also ensures that the theory is consistent if set theory is. The typed lambda calculus avoids paradox in a simple and pretty much cost free way. Since the work of Dana Scott and Gordon Plotkin in the early seventies (Scott (1972)), we have abstract models of the type-free lambda calculus. But they have certain drawbacks for the purpose of studying natural language semantics. The models used by Scott require that the values of lambda terms be continuous functions in the sense that one can compute their value in the limit given some long enough run of values. But it is precisely the operators of classical logic like $\neg$, $\forall$, or $\exists$ that fail to be continuous in the requisite sense.

While the untyped lambda calculus has various uses in mathematics and computer science,[2] there are compelling linguistic reasons to adopt a typed lambda calculus in constructing logical forms. The models of a Montagovian typed lambda calculus together with the standard quantifiers and connectives are unproblematic set theoretic constructions. Furthermore, the typed lambda calculus provides a tight connection between syntactic categories and semantic

[2] For example, in representing fixed points or recursion. In recursion, a function or term takes itself as an argument. Turner (1989) has argued that in natural language terms, in particular property terms, can take themselves as arguments as well:

(2.2a) Being nice is nice.
(2.2b) It's bad to be bad.

I will argue in chapter 10, however, that examples like (**??**) don't argue against a typed notion of predication.

Copyrighted material

2.2.2 Locations vs. objects

Another example of a grammatically grounded distinction between types in natural language concerns the distinction between locations and physical objects. Places are fixed elements in the terrestial reference frame, while objects typically have a complex internal structure and can move with respect to the terrestial reference frame. Some evidence for this distinction comes from Basque, where the grammar encodes differences between location and objects via two genitive cases *-ko* and *-ren*; locations in general easily take the genitive *-ko* but not *-ren*, while objects in general do the reverse. Aurnague (1998) distinguishes the following types: PLACES (e.g., valley, field, river, mountain, hill), OBJECTS (e.g., apple, glass, chair, car), and MIXED OBJECTS (e.g., house, church, town hall). Of particular interest are the "mixed objects" and the behavior of their expressions in Basque. The terms for mixed objects readily accept both forms of the Basque genitive. So if we accept the encoding hypothesis for Basque, mixed objects like houses would appear to belong to two types, or two ontological categories, at the same time—LOCATION and PHYSICAL-OBJ—neither of which is a subtype of the other (it is neither the case that the properties associated with physical objects are inherited as properties of places nor that the properties associated with places are inherited as properties of physical objects).

(2.6) Maite dut etxeko atea haren paretak harriz eginak direlariak.
(Michel Aurnague p.c.)
I like the door of (locational genitive) the house the walls of (physical object genitive) which are made of stone.

Prepositions in English serve to distinguish between places or locations and physical objects, though the distinctions are less clear cut than in Basque (Asher 2006).

2.2.3 Mass vs. count

The mass/count distinction is another type distinction marked in many languages. Chinese marks this distinction in its system of classifiers, whereas other languages like English mark the distinction with determiners and to some extent within nouns (although many nouns can receive both a count and noun interpretation). Certain determiners in English are designated as mass determiners—e.g., *much*, as in *much water*, *much meat*. They do not go with count nouns in general—e.g., *much person*, *much people* are malformed. Other determiners like *many, every, the* apply both to count nouns and mass nouns but

Copyrighted material

require special interpretations when put with mass nouns.[4] Thus *every water*, *many waters* must range over contextually given portions of water or perhaps kinds of water.

Languages also distinguish "determiner phrases" (DPs) denoting quantized portions of matter like *many waters* or *the water* from DPs with ordinary count nouns.

(2.7) You can take two piles and put them together to make a bigger pile.
(2.8) Two little waters make one large water. (in a restaurant)

Piles pick out portions of matter as does the expression *two waters*. But you can't put two dogs together to make a bigger dog. This last observation holds for all count nouns.

2.2.4 Kinds vs. individuals

Another universal distinction in language is the distinction between kinds and individuals. Kinds are often expressed in English with a bare plural noun phrase (e.g., *cats*, *numbers*, *people*) but can also be expressed with other constructions:

(2.9) The Mexican fruit bat is common in this area.

Linguists take the definite noun phrase in sentences like (2.9) to refer to a kind rather than to range over individual members of the kind. They argue, quite sensibly, that the predicate *is common in this area* (but also others like *is extinct*, *is widespread*) cannot hold of an individual but only of kinds or species. Thus, most languages encode a three way distinction between the types of masses, countable individuals, and kinds. Chinese, once again, encodes the distinction between kinds and individuals in the classifier system.

(2.10) Moby Dick shi yi tiao / *zhong jing.
Moby Dick be one Cl_{tail}/*Cl_{kind} whale.
Moby Dick is a whale (individual).

2.2.5 Containers vs. containables

A more subtle type distinction encoded in the system of prepositions within English involves containers and containables. In general anything that desig-

[4] Borer (2005a) argues that the mass/count distinction is to be located in a classifier-like syntactic projection and that nouns are by default mass. She cites as evidence the fact that nouns without classifiers in Chinese are interpreted as mass.

Copyrighted material

nates a specific volume or enclosure can be a container. Many physical objects can serve as containers, although some cannot.

(2.11) a. The water is inside the pitcher.
b. The keys are inside the car.
c. John put the keys inside his pocket/inside the drawer.
d. # John threw the keys inside the air. (Versus: John threw the keys in the air.)
e. # John put the wine inside the water. (Versus: John put the wine in the water.)

2.2.6 Plurality

Plurality also introduces type distinctions. Research on plurals, as I briefly outlined in chapter one, distinguishes at least two types of plural predication: distributive and collective. Distributive predication occurs when a property or relation is predicated of each element of a set as in

(2.12) The boys each worked hard.

On the other hand, (2.13a,b) exemplify collective predications where a property is predicated of the whole set of students but not of each student individually.

(2.13) a. The students surrounded the building.
b. The students mowed the whole meadow.

Sometimes collective predication occurs with singular nouns (and so this semantic phenomenon must be distinguished from the syntactic phenomenon of number).

(2.14) The committee is meeting in the lounge. (collective predication)

Some predicates, finally, put no requirements on how their plural arguments are to be understood. [5]

(2.15) Three girls danced with four boys.

(2.15) makes no claims about whether the boys distributively or collectively danced with the girls, only that there were three girls and four boys and that dancing went on between them.

[5] These are called *cumulative* predications.

Copyrighted material

The distinctions between types of predication mark distinctions in type presuppositions. For instance, an account of the lexical meaning of the word *disperse* or *surround* must mark it as requiring an external argument that is of group or collective type, which means that the argument must denote a group of individuals or range over groups of individuals. Verbs like *disperse* or *surround* do not go well with inherently distributive quantifiers like *most students*, whereas they apply perfectly to plural noun phrases that can be interpreted as denoting groups:

(2.16) a. ?Most students surrounded the building.[6]
b. The students in the square surrounded the building.

Other predicates like *work* must be interpreted distributively and impose type presuppositions on their arguments to that effect.

2.2.7 Types and specific lexical items

Some verbs have quite specific type requirements encoded in their selectional restrictions.

(2.17) a. John weeded (mulched, hoed ...) the garden (lawn, area, tomatoes, peas, plants ...).
b. John hoed (mulched) the weeds.
c. # John mulched, hoed, weeded the water.
d. #John shoveled the closet.

You can't weed, hoe, or mulch certain types of locations—bodies of water, for instance. Locations that can be weeded have to have dirt or soil in them. And you can only shovel a location open to the elements. This is not a matter of world knowledge but a matter of grammar, broadly construed: one can perfectly well imagine someone cleaning a body of water like a lake of floating algae or water plants, but we don't call that *weeding*.

Given the fine-grained type distinctions made by the language, it should not be surprising that when predicates combine with arguments of different types, the meaning of the predicate shifts in the resulting predication. Consider the following.

(2.18) a. John swept (shoveled, wiped ...) the closet (room, walkway, kitchen, fireplace, floor, counter ...).
b. John swept (shoveled, wiped ...) the dirt (debris, manure, sand, slush, litter, shavings, cinders, dust ...).

[6] Note that the partitive DP *most of the students* does admit of a collective interpretation, unlike the straight quantified DP *most students*.

Copyrighted material

Predications involving *sweep* have a different content depending on whether the direct object is a location, a place, or a surface (we can assimilate all of these here to the type LOCATION), or whether the direct object is a portion of matter. The resulting meanings for *sweep*, for instance, are so different that they don't license copredication or ellipsis:

(2.19) a. John swept the kitchen and Mary the entryway.
b. John swept the dust and Mary the leaves.
c. #John swept the kitchen and Mary the leaves.
d. # John swept the kitchen and the dust.

Another example of an apparent meaning shift in a predicate because of a shift in type of its argument comes from communication verbs. The meaning of verbs like whisper, whistle, whine, etc. varies with respect to whether it has an object argument; furthermore, the type of this argument can also affect the verbal meaning.

(2.20) a. John shouted (whispered, whistled, whined ...). (activity)
b. John shouted (whispered, whistled, whined ...) a warning. (accomplishment)
c. John shouted (whispered, whistled, whined ...) at the animal. (accomplishment or activity)
d. The bullets whistled past John. (accomplishment)

Such so called *verbal alternations* have occupied linguists for many years.[7] They offer us another means for seeing how languages encode a sophisticated system of types. But they also clearly pose a challenge for lexical semantics that resorts to types: how can we account for the shifts in the meaning of a verb given its different arguments?

Other well-known verbal alternations involving prepositions show subtler shifts in meaning.

(2.21) a. John treated Mary to dinner.
b. John treated Sam for cancer.
c. John treated the cancer.
(2.22) a. John loaded the hay on the wagon.
b. John loaded the wagon with the hay.
c. John sprayed the paint on the wall.
d. John sprayed the wall with the paint.

[7] For a comprehensive bibliography and discussion, see Levin (1993).

Copyrighted material

part of the domain of interpretation. Further, fictional objects not only don't exist in the actual world; they don't exist in possible worlds either—i.e., other ways in which the world could be. A fictional creature like a hobbit is not at all like the sister that I might have had; the latter exists in a possible world, the former does not. That's what it is to be fictional. So on a view that identifies types with their inhabitants, types corresponding to fictional objects and the absurd type would be the same type, since they have the same extension or the same set of inhabitants, namely the empty set. But types of fictional objects are intuitively distinct from $\perp$.

Whether a term describes a fictional character or not certainly appears to make a difference as to how predications are understood. Within fiction, there is no question of checking or wondering whether the predication actually results in a literal truth. It is even quite controversial among philosophers who have written on fiction whether terms that appear to refer to fictional entities refer in fact to anything at all. On the other hand, fictional talk differs from metaphorical or loose talk; fictional talk is literal—the trees in *The Lord of the Rings* literally speak (see 2.25) whereas in metaphorical talk (2.26) the predications aren't to be taken at face value.

(2.25) Look! the trees are speaking. (*Lord of the Rings*)

(2.26) These trees are really speaking to me. I'm going to paint the living room green.

To make sense of this difference in predicational behavior, we should distinguish a type of fictional objects. And it should be distinct from the absurd type no matter what the circumstances of the actual world are. This leads me to adopt the thesis that types are neither to be identified with their actual inhabitants (extensions) nor even their possible inhabitants (standard semantic intensions). They are "hyper-intensional."[8]

If types are intensional entities, they are not intensions as semantics standardly conceives of them—i.e., as functions from indices (possible worlds or sequences consisting of a world, time, context, and other appropriate elements) to extensions.[9] *Modulo* a certain understanding of fictional objects, I have established that types aren't to be identified with extensions, intensions, or sets thereof. So where do types fit into an ontology of abstract entitites? We need to think about the relations between the following sorts of abstract entities:

- types

[8] Reinhard Muskens in Muskens (2007) also argues for a hyper-intensional construal of types.

[9] The extension (at an index) of an individual constant is an individual; the extension of a 1-place predicate is a set; and an extension of a closed formula is a truth value.

Copyrighted material

- concepts
- properties

Given the relatively well-understood analysis of properties in formal semantics and pragmatics as semantic intensions (or as functions from indices to extensions), types cannot be properties. In addition, properties are typically understood to be mind-independent entities, entities whose existence is not dependent on the existence of minds. Types, however, given their role in guiding predication, are part of the conceptual apparatus necessary for linguistic understanding. They are *mind-dependent representations* of mind-independent properties and individuals. This leads us to the hypothesis that types are concepts, which I take to be mind-dependent entities as well.

This hypothesis has some promising support. Concepts come at different levels and granularities. They form a hierarchy, just as types do. There are concepts of what it is to be a property, what it is to be an individual, what it is to be a physical object, and so on. There are also much more specific concepts: the concept of red, the concept of Ségolène Royale or of Hillary Clinton. Like types, concepts are the internal, mind dependent reflection of mind independent properties and individuals they are concepts of. Concepts (and types) have their own "internal" semantics which has to "track," in the appropriate way, the properties and individuals they are concepts of. It is in virtue of such tracking that a concept is a concept of some object or property. For example, the concept of red tracks the property of being red. A concept RED of the color red is triggered by something at a particular location l in the conceiver's visual field, typically when the conceiver is perceptually aware of something at l that is in fact red in color. Though this tracking is generally reliable, it can occasionally fail. For instance, if there were something wrong with a conceiver's visual apparatus or the circumstances of the perceptual event were very non-standard, the concept RED could be introduced as holding of some object when the object of which the conceiver is perceptually aware is in fact green in color. Another way the tracking could go wrong is that the object is in fact red but the concept fails to be triggered.

Although making this notion of "tracking" precise is a book-length project on its own, I can say a few words here. I understand this notion in terms of how the rules for the application of the concept in the conceptual system function. In fact it is these rules that define the concepts and give them their content. These rules look something like natural deduction rules. A concept has certain "introduction rules" and certain inference rules that it licenses. The introduction rules stipulate that an object must satisfy certain conditions for its falling under a certain concept; these conditions may be determined by the sensory system

Copyrighted material

of the organism or by other associated concepts. Talk of satisfying the application conditions of a concept can be replaced by the notion of something's being provably of that type, giving an intuitionist flavor to the interpretation of types.[10] There are more complex combination rules as well that determine how one concept interacts with others. It is these same rules that determine whether an object is of a basic type like CAT or not. Type presuppositions and the rules for type presupposition justification are instances of rules of type and concept combination. This system of rules supplies what computer scientists call an *internal* proof theoretic semantics of a term and these rules define or give the content of concepts and *a fortiori* of types. This internal semantics contrasts with the mind-external, denotational semantics of the terms, which involve real world objects, properties, and so on. The two are connected by the tracking mechanism. Linking concepts and types together helps us understand both: concepts get a rigorous framework from type theory, while types are now linked to the agent's sensory interactions with his environment and as well as the interactions between other concepts/types in the linguistic system.

It is plausible that humans in a given speech community share concepts and a type system. They must do so in order to communicate, to exchange information.[11] The internal semantics conceived as a system of proof or computation rules allows us to make sense of a shared conceptual system. If your concept of red and my concept of red have the same internal semantics, then we can be said to have the same concept of red, and similarly for other concepts. We can prove of two such proof or computational systems whether they are the same or not, using several different criteria. The crudest one is input/output equivalence; roughly two systems are input/output equivalent, just in case they give the same results in the same cases. With respect to the conceptual/type system, this would mean that the same linguistic actions and judgments are observed in the same contexts. Demonstrably, members of the same speech community have systems that are largely input/output equivalent. There is also the criterion of trace equivalence, where for each computation, the same sequence of actions is observed in the two systems. Finally, there is the criterion of bisimulation, according to which for each point in the computation there are the same possible continuations.

Because the rules that make up the internal content of concepts are in general defeasible, concepts cannot determine reference to mind independent entities or properties independently of the context of their application. My notion of a concept thus differs from Frege's notion of a sense.[12] On a standard Fregean

[10] More on this below.

[11] A proof of this fact is to be had in Lewis (1969).

[12] Peacocke (1992) uses concepts as constituents of thoughts to account for informativeness and

Copyrighted material

view, senses compose together to yield thoughts, which determine on their own the truth conditions of sentences and discourses. I need a different view. A proposition is the result of the compositional interpretation of logical forms for the words that make up the sentence or discourse. But the compositional interpretation of a logical form results for me, as for most semanticists, in an intension—a function from indices to truth values. Since concepts don't determine extensions, let alone intensions, concepts cannot be the constituents of propositions.[13] Philosophers might take intensions to be simply formal stand-ins for what propositions "really are." But even then, if sentences are typically about mind-independent objects and the properties and relations these objects stand in, then the "real propositions" such sentences express will not contain concepts either— at least not of the sort I have in mind, mind-dependent entities with an internal semantics.

More concretely, consider basic referential expressions: indexicals, proper names, demonstratives. The content of these terms, intuitively, has to do with the individuals they denote, not some proof object or set of rules for defeasibly determining whether a given object is in their denotation. The content of the type associated with *you* consists of rules for determining who the audience is in a particular context. But that's not the contribution of *you* to the content of a clause in which it occurs. Its semantic content in this sense is the audience itself.

Forceful externalist arguments given by Kripke (1980), Putnam (1975), and others show that our concepts associated with names of individuals and natural kinds do not suffice to determine the extensions or intensions of these expressions. If one looks to the behavior of such terms in modal contexts, there is compelling evidence that their meanings are not in general determined by "what is in the head" of a competent speaker of the language. In this respect too, types resemble concepts; they are tied via the expressions they type to properties and real world entities, but they are not identical to properties or real world entities, nor to sets thereof. They are part of our conceptual apparatus used to guide predication.

To show how the externalist arguments affect types, let's consider a typical Twin Earth scenario, familiar from the externalist literature cited above. Oscar on Earth and his twin "Twin Oscar" on Twin Earth speak syntatically identical languages and are type identical down to their molecular constitution. In

Frege style puzzles about the substitution of coreferential terms; he takes concepts to be something like Fregean senses.

[13] In fact, on the standard semantic conception of propositions, propositions don't have "constituents" except in a set theoretic sense—and these would be sets of n-tuples of worlds, other indices, and truth values.

Copyrighted material

keeping with general physicalist principles then, they have the same internal make up, the same thoughts, the same conceptual system. In particular their linguistic judgements about semantic well-formedness will be the same. Thus, when Oscar and Twin Oscar each interpret the strings *water is wet* and *water is a tree*, they assign the same syntactic form and the same semantic types to each expression; for the first string they will each construct a coherent logical form using the tools of the theory of predication, while for the second they will not. But whereas they marshal the same type system and conceptual resources when dealing with their languages, the languages of Oscar and Twin Oscar are different—they have a different semantics. On Earth *water* picks out the kind H_20, or real water, whereas on Twin Earth the string *water* has a different semantics—it denotes a chemical compound distinct from H_20, which, so the story goes, is *XYZ* or "twin water." *Water* in English picks out a different substance from *water* in Twin English. Thus *water* makes dramatically different contributions to truth conditions in sentences of English and Twin English like:

(2.27) Water is H_20.

(2.27) is true (and necessarily true) in English but false (and indeed necessarily false) in Twin English. Such Twin Earth scenarios are well established in the philosophical literature, and intuitions about them are relatively robust. They constitute powerful evidence that as internal reflections of properties and individuals, concepts, and types are not identical with mind-independent properties or individuals nor do they determine them, although they are associated with them through the tracking mechanism.

If concepts, and *a fortiori* types, are not constituents of propositions, they can nevertheless compose together. Making the linguistically relevant types a subset of the set of concepts allows us to use the logical framework of types to explore concepts. Types associated with properties are functions from one type into another; when given an appropriate type as argument, types associated with properties return a new type. We can even compose types or concepts together to give us types associated with propositions or semantic intensions.[14] I shall call the type corresponding to a proposition a thought. If there are as many types as there are distinct word stems in the language, then the hypothesis that types are concepts and compose together to yield thoughts gains in plausibility.

Let us look a bit more closely at the composition of thoughts. Once we

[14] Composition allows us to talk of concepts as constituents of thoughts, though this talk should be interpreted with care. It is true that concepts compose together to form thoughts, but that does not necessarily mean that the finished product will actually contain those concepts. Nevertheless, because I shall identify thought contents at least in part with the derivation of the thought from its constituents through composition, that's pretty close to constituency.

Copyrighted material