

Nominalization, Specification, and Investigation

By

Richard Wyley Lawrence

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

in

Philosophy

in the

Graduate Division
of the
University of California, Berkeley

Committee in Charge

Professor John MacFarlane, Co-chair
Professor Paolo Mancosu, Co-chair
Professor Hannah Ginsborg
Professor Line Mikkelsen

Summer 2017

Nominalization, Specification, and Investigation

This dissertation, in its entirety, is the original work of the author.

Copyright

Richard Wyley Lawrence

2017

Abstract

Nominalization, Specification, and Investigation

by

Richard Wyley Lawrence

Doctor of Philosophy in Philosophy
University of California, Berkeley
Professor John MacFarlane, co-chair
Professor Paolo Mancosu, co-chair

What does it mean for something to be an *object*, in the broad sense in which numbers, persons, physical substances, and reasons all play the role of objects in our language and thought? I argue for an epistemological answer to this question in this dissertation. These things are objects simply in the sense that they are answers to questions: they are the sort of thing we search for and specify during investigation or inquiry. They share this epistemological role, but do not necessarily belong to any common ontological category.

I argue for this conclusion by developing the concept of an *investigation*, and describing the meaning of nouns like ‘number’ in terms of investigations. An investigation is an activity structured by a particular question. For example, consider an elementary algebra problem: what is the number x such that $x^2 - 6x + 9 = 0$? Beginning from this question, one carries out an investigation by searching for and giving its answer: $x = 3$. On the view I develop, nouns like ‘number’ signify the *kind* of question an investigation addresses, since they express the range of its possible answers. ‘Number’ corresponds to a ‘how many?’ question; ‘person’ corresponds to ‘who?’; ‘substance’ to one sense of ‘what?’; ‘reason’ to one sense of ‘why?’; and so on.

I make use of this idea, which has its roots in Aristotle’s *Categories*, to solve a puzzle about what these nouns mean. As Frege pointed out in

the *Foundations of Arithmetic*, it seems to be impossible for

(1) The number of Jupiter's moons is four.

to be true while

(2) Jupiter has four moons.

is false, or vice versa. These sentences are just two different ways of expressing the same thought. But on a standard analysis, it is puzzling how that can be so. Every contentful expression in (1) has an analogue in (2), except for the noun 'number'. If the thought is the same whether or not it is expressed using 'number', what does that noun contribute? Is the concept it expresses wholly empty? That can't be right: 'number' is a meaningful expression, and its presence in (1) seems to make that sentence *about* numbers, in addition to Jupiter and its moons. So why doesn't it make a difference to the truth conditions of the sentence?

The equivalence between these two sentences is famous, but it is hardly a unique example. To say that Galileo discovered Jupiter's moons is just to say that the *person* who discovered them was Galileo. Likewise, to say that Jupiter spins rapidly because it is gaseous is just to say that the *reason* it spins rapidly is that it is gaseous. So the same puzzle that arises for 'number' also arises for 'person', 'reason', and other nouns of philosophical interest. If they are significant, what contribution do they make?

Because the problem is general, I pursue a general solution. The sentences which introduce the nouns in these examples are known as *specificational* sentences, because the second part specifies what the first part describes. In (1), for example, 'four' specifies the number of Jupiter's moons. I argue that we should analyze specificational sentences as pairing questions with their answers. At a semantic level, a sentence like (1) is analogous to a short dialogue: "How many moons does Jupiter have? Four." This analysis is empirically well supported, and it unifies the theoretical insights behind other approaches. Most importantly, it solves the puzzle. According to this analysis, (1) asserts no more or less than the answer it gives, which could also be given by (2); that is why they are equivalent. But it differs from (2) by explicitly marking this assertion as an answer to the 'how many?' question expressed by 'the number of Jupiter's moons'. That is why the two sentences address different subject matters and have different uses.

In order to formulate this analysis in a contemporary logical framework, I apply the concept of an investigation in the setting of game-theoretical semantics for first-order logic. I argue that quantifier moves in semantic games consist of investigations. A straightforward first-order representation of the truth conditions of specificational sentences then suffices to explicate the question-answer analysis. In the semantic games which characterize the truth conditions of a specificational sentence, players carry out investigations structured by the question expressed in the first part of the sentence. When they can conclude those investigations by giving the answer expressed in the second part, the sentence is true.

The game semantics characterizes objects by their role in investigations: objects are whatever players can search for and specify as values for quantified variables in the investigations that constitute quantifier moves in the game. This semantics thus captures the sense in which objects are answers to questions. I use this account to offer a new interpretation of Frege's claim that numbers are objects. His claim is not about the syntax of number words in natural language, but about the epistemological role of numbers: numbers are the sort of thing we can search for and specify in scientific investigations, as sentences like (1) reveal.

The substantive form belongs originally only to things, the adjective form to qualities, the verb form to events. But, of course, language could not in its judgments always begin with the thing, and annex qualities and action to this as the subject; it had to make the qualities in themselves and action in itself also object of its reflection. Hence it severed their connection with things, gave *them* a substantive form. . . . Almost invariably we find a tendency to make the newly acquired syntactic dignity of words convertible with a new metaphysical dignity acquired by their matter. . . . Language creates for us a mythology, from which, of course, in the use of language we can never wholly set ourselves free without becoming pedantically precise, but against the influence of which on the moulding of our thoughts we ought to be carefully on our guard.

— Hermann Lotze (1887, pp. 629–630)

If I want to speak of a concept, language, with an almost irresistible force, compels me to use an inappropriate expression which obscures—I might almost say falsifies—the thought. . . . We cannot avoid words like ‘the concept’, but where we use them we must always bear their inappropriateness in mind.

— Gottlob Frege (1891/1997a, p. 174)

Acknowledgments

This dissertation has been a long time in the making, and I am indebted to many people for their conversation, feedback, encouragement, and support.

I must thank first and foremost my advisors, who have had an enormous influence on the dissertation. Their feedback and advice shaped its overall narrative, the contents of particular chapters, and the details of its arguments. The rough edges that remain are, of course, my own. (I must also thank them for their patience where my own inability, or authorial stubbornness, have left their comments unaddressed.) Hannah Ginsborg's interest and encouragement during the early stages of the project helped make it as ambitious as it is, and my conversations with her kept me focused throughout on the philosophical issues at the heart of the matter. Line Mikkelsen's expertise and guidance through the linguistic literature was immensely helpful as I worked out my study of specificational sentences in Part I. Paolo Mancosu's knowledge of the history of logic and mathematics had a strong influence on the chapters in Part II, and my discussions with him have helped me conceive how the project will develop into further research. And John MacFarlane, whose incisive comments have sharpened my thinking since my first day of graduate school, was unfailingly attentive to both my biggest ideas and the smallest details of my work at every stage. Over the years, each of my advisors has supported me in too many other ways to list here, and I will never forget the warmth, generosity, faith, and pride they have shown toward me as I brought this project to fruition.

Many other Philosophy faculty at Berkeley, especially Timothy Clarke, Klaus Corcilius, Michael Martin, Hans Sluga, and Seth Yalcin, deserve my thanks for their words of encouragement and the conversations that they had with me about my work. I also received valuable comments from faculty members at other institutions, including Ivano Caponigro,

Robert May, and Gila Sher. I am grateful to Janet Groome and David Ly-
naugh for relieving me of a thousand administrative tasks, for patiently
enduring my constant presence in their office, and for providing me with
sweets and conversation whenever I needed a break. And I will always
be grateful for the time and financial support I have received from the
Department and the University as a whole, which enabled me to pursue
this project and see it through.

Before I came to graduate school, Thomas Ricketts told me that most
of my education would come from my fellow graduate students. He
was right, and I couldn't have made a better choice in coming to Berke-
ley. My fellow graduate students have been my friends, my support net-
work, and the first sounding board for all of my ideas. I am especially
grateful to the following students for their comments on my writing and
the conversations they had with me about my work, though I am sure
that this list is incomplete: Austin Andrews, Eugene Chislenko, Lind-
say Crawford, Caitlin Dolan, Peter Epstein, Melissa Fusco, Nick Good-
ing, Jim Hutchinson, Ethan Jerzak, Jeff Kaplan, Alex Kerr, Arc Kocurek,
Katie Mantoan, Luke Misenheimer, Ethan Nowak, Emily Perry, Kirsten
Pickering, Michael Rieppel, Rachel Rudolph, Pia Schneider, Janum Sethi,
Umrao Sethi, Justin Vlasits, and James Walsh.

Of these, Michael Rieppel deserves special mention. Mike and I dis-
covered early on in my graduate career that our philosophical interests
and outlook were closely aligned. In the years since, we have enjoyed
many, many conversations about philosophy, particularly over beers at
the Heart and Dagger. He has read and commented on more of this dis-
sertation than anyone apart from my advisors, and it is possible that his
influence on my thinking exceeds even theirs; I only hope that my con-
versation has been half as useful to him as his was to me. Mike is also
responsible for introducing me to the circle of friends who have been my
closest and most trusted companions during graduate school, including
Anna Bailey, Bill Campbell, and Lindsay Crawford. Their friendship has
sustained me all throughout the years I've spent writing this dissertation,
and I would never have completed it without their encouragement and
commiseration—not to mention their cooking, their wine and whiskey,
their songs, and their adventures.

I also owe much to my family for their patience and love as I made
my way through graduate school. I rarely faced the dreaded question
from them of when I would finish—or what would happen if I didn't.

I am especially grateful to my sister Kendra and her husband Robby, who allowed me to stay in their beautiful home in the Oakland hills for a month during the summer of 2016. Most of the final chapter of this dissertation was written over the course of that month. My stay there provided the seclusion and atmosphere that I needed to work out how I should finally conclude the project. That was an otherwise difficult period for me, and staying in their home was the most welcome retreat I can imagine; without them, I am not sure whether I would have reached the conclusion I did, or any conclusion at all.

Finally, I am grateful to Jillian Budd, without whom I would never have started this great adventure, and never would have believed I could finish it. No other person has done more to support me in the years since I began. Now that I have finished, it is time to begin again.

Contents

Acknowledgments	ii
Contents	v
I Nominalization and Specification	1
1 Numbers, questions, and (some) nouns	2
1.1 Statements of numbers as focus constructions	4
1.2 A syntactic proposal	8
1.3 Nouns corresponding to questions	18
1.4 The plan	32
2 Specification	34
2.1 Ambiguities in pseudo-cleft sentences	35
2.2 Other types of specificational sentences	49
3 Semantic analyses of specification	64
3.1 The inversion analysis	66
3.2 The equative analysis	74
3.3 The question-answer analysis	88
II Investigation	114
4 Investigation	115
4.1 The structure of investigations	118
4.2 Investigatory semantics	131
4.3 Back to natural language	147

CONTENTS

5 Numbers, again	155
5.1 Frege's understanding of objects and concepts	160
5.2 The investigatory conception of concepts and objects	185
5.3 Platonism and the independence of objects	209
Bibliography	219

Part I

Nominalization and Specification

Chapter 1

Numbers, questions, and (some) nouns

Frege thought that sentences like the following should be analyzed as identity statements:

- (1) The number of moons of Jupiter is four.

He took both ‘the number of moons of Jupiter’ and ‘four’ as proper names, and ‘is’ as the sign of identity. His motivation for doing so seems to have been primarily mathematical. Frege was convinced that the numbers must in some sense be self-subsistent objects if arithmetic is to be a genuine science.¹ Thus, our words for numbers must be like other words that stand for objects: names or singular terms. There are some reasons to be suspicious of the identity analysis, though, that I want to examine here.

The first, as Frege had observed already, is that number words are also used in attributive constructions. In modern syntactic terms, they are *determiners*. For example:

- (2) Jupiter has four moons.

In this sentence, the words one can substitute for ‘four’ while preserving well-formedness are not names or singular terms, but words like ‘the’, ‘some’, ‘most’, ‘its’ or ‘those’.² This should make us suspicious that ‘four’

¹See Frege (1884/1980, § 57).

²Of course, one may also substitute adjectives, like ‘bright’ or ‘beautiful’, in this position. I treat number words in attributive position as determiners, rather than adjectives, because this is one way of accommodating Frege’s insight that “the content of a statement of number is an assertion about a concept” in modern terminology (Frege,

as it occurs in (1) is a singular term, since substituting one singular term for another usually preserves well-formedness. Furthermore, (1) and (2) seem to have the same truth conditions; it is difficult to see how one could be true while the other was false. This makes it unlikely that ‘four’ is simply ambiguous between a singular term use and a determiner use. Positing that ‘four’ is ambiguous would leave the close semantic relationship between these two sentences unexplained.

Another cause for suspicion is the fact that the definite description in (1) doesn’t have the relationship to the corresponding indefinite description that we would expect. Normally one can exchange the definite article for the indefinite:

- (3) a. The composer of *Tannhäuser* is Wagner.
 b. A composer of *Tannhäuser* is Wagner.

Both of these sentences are grammatical, and (3-b) says much the same as (3-a). The difference is that (3-b) does not presuppose that *Tannhäuser* has a unique composer, whereas (3-a) does. By contrast,

- (4) ?A number of moons of Jupiter is four.

is not even clearly grammatical, and if it is, it’s not clear what it says or how its meaning is related to (1).³

Finally, there is some evidence that the definite description in (1) might be more like a question than a singular term. Definite descriptions are usually ambiguous between an *acquaintance* reading and an *indirect question* reading when they occur as the direct object of ‘knows’. For example, (5-a) is ambiguous between the two readings (5-b) and (5-c):

- (5) a. Ralph knows the teacher.
 b. Ralph knows who the teacher is. (indirect question)
 c. Ralph is acquainted with the teacher. (acquaintance)

But when a definite description headed by ‘the number’ occurs as the object of ‘knows’, the indirect question reading is strongly preferred.

- (6) a. Frege knows the number of sections in the *Grundlagen*.

1884/1980, § 46).

³This observation, and the example in (3), are due to Hofweber (2007). I discuss Hofweber’s view further below.

- b. Frege knows how many sections are in the *Grundlagen*.
- c. ?Frege is acquainted with the number of sections in the *Grundlagen*.

It is much more natural to read (6-a) as synonymous with (6-b) than as saying that Frege is acquainted with the number 109. To my ear, the acquaintance reading is not even clearly possible. Even if it is, we need an explanation of why the ambiguity is so lopsided in the case of ‘the number’. Why is the indirect question reading so strongly preferred? If definite descriptions headed by ‘the number’ are not semantically singular terms, it would explain our preference for the indirect question reading of (6-a). This is another cause for suspicion about the identity analysis of (1).

I think the suspicions are right: (1) is not an identity sentence, at least not in the sense that it should be thought of as an identity sign flanked by two semantically singular terms. The apparent definite description ‘the number of moons of Jupiter’ should be interpreted more like an indirect question than a singular term which refers to an object. Likewise, ‘four’ should be interpreted more like an answer to that question. This means it also needn’t be construed as a singular term, since pronounced answers to questions can come from many parts of speech.

This claim about (1), if it is correct, has interesting consequences for the philosophy of mathematics. I am more interested here, however, in generalizing the lessons of this example to other sentences, about things other than numbers. I will argue that ‘number’ is part of a special class of nouns, nouns which have an important relationship to questions. It is this fact that accounts for the suspicious behavior of ‘the number of moons of Jupiter’ in the examples above. I will characterize this class of nouns, and say more about what they mean, in the following chapters. Here, I focus mainly on bringing the puzzle posed by these nouns more clearly into view.

1.1 STATEMENTS OF NUMBERS AS FOCUS CONSTRUCTIONS

Thomas Hofweber (Hofweber, 2007, 2005) has also recently taken issue with Frege’s analysis of (1), and developed an alternative analysis. In response to some of the suspicions I pointed to above, he proposes that (1) is not an identity statement, but a *focus construction*. It is merely a syntac-

tic variant of sentence (2). The syntactic variation communicates stress or focus on the ‘how many’ aspect of what is communicated, but (1) and (2) are semantically equivalent, and communicate the same information. To support this claim, Hofweber makes some very illuminating observations. I agree that these observations motivate a different analysis of (1). But as we will see in the next section, I don’t think Hofweber’s proposal that (1) is a focus construction is quite right. The problem is that it does not account for the role of ‘number’ in sentences like (1).

Hofweber’s argument deals with the use of number words like ‘four’ in ordinary, non-arithmetical language.⁴ As we have seen, there are two such uses. Number words sometimes occur syntactically as determiners, as in (2), and sometimes as singular terms, as in (1). Hofweber argues that the latter uses are not *semantically* singular terms, as Frege thought. We should not semantically construe such uses as identity statements.

Hofweber cites a family of substitution problems, like those I explained above, as evidence that sentences like (1) are not identity statements containing semantically singular terms. Specifically, he appeals to the oddness of substituting the indefinite article for the definite article in number-descriptions, like in (4). He also points out that one cannot substitute a number-description for a number word in its determiner use. If we try to substitute ‘the number of moons of Jupiter’ for ‘four’ in (2), the resulting sentence is not grammatical:

(7) *Jupiter has the number of moons of Jupiter moons.

This is odd, if we assume that ‘the number of moons of Jupiter’ and ‘four’ are always co-referring singular terms, for we should be able to substitute one for the other in extensional contexts like (2). It’s not odd, of course, if ‘four’ is ambiguous between a determiner use and a singular term use, since there’s no reason to expect that ‘the number of moons of Jupiter’ could be substituted for the determiner uses of ‘four’. But as I pointed out above, positing that ‘four’ is ambiguous leaves the felt equivalence of (2) and (1) unexplained.

Hofweber concludes from this evidence that (1) should not be read as an identity statement. Instead he claims that, in their ordinary usage,

⁴The use of number words in arithmetical discourse is somewhat different. Though Hofweber also has an account of arithmetical discourse, articulated in Hofweber (2005), it is outside the scope of my interests in this chapter.

statements of number like (1) are best understood as focus constructions. The central idea of his proposal is that the relationship between (2) and (1) is much like the relationship between (8-a) and (8-b), or between (9-a) and (9-b):

- (8) a. Johan likes soccer.
- b. It is soccer that Johan likes.
- (9) a. Mary entered quietly.
- b. Quietly is how Mary entered.

The distinction here is between *neutral* sentences, like (8-a) and (9-a), and *focused* sentences, like (8-b) and (9-b). Intuitively, a focused sentence places stress or emphasis on some part of what's said, where the neutral sentence does not. The idea is that (1) is a focused variant of the neutral sentence (2).

Hofweber argues that these focus constructions are truth-conditionally equivalent to their neutral variants.⁵ For how could it be true that Johan likes soccer, but false that it is soccer that Johan likes? Or true that Mary entered quietly, but false that quietly is how she entered? Or true that Jupiter has four moons, but false that the number of moons of Jupiter is four? In each case, the focused sentence communicates just the same information as its neutral variant; it just does so in a different way. So the neutral and focused versions of these sentences are semantically equivalent. How, then, should we account for the differences between them?

Hofweber's answer is that the neutral and focused versions of a sentence have different communicative uses, which accounts for the differences in syntax. The most important communicative difference is that neutral sentences can be used to answer more questions than their focused variants. For example, (8-a) can be an answer either to 'Who likes soccer?' or 'What does Johan like?', while (8-b) is only an appropriate answer to the latter. The second sentence exhibits a focus on some *aspect* of the information being communicated. When that aspect of the information is not in question, or when some other aspect of the information was asked for, this focus is communicatively inappropriate. Hofweber

⁵Hofweber points out that this is not true in general. Focus can affect truth conditions. But it doesn't seem to in the examples above, including the case of numbers. Indeed, it is the felt equivalence of (2) and (1) that Hofweber thinks needs to be explained.

claims that (1) communicates the same information as (2), but focuses on the ‘how many’ aspect of that information. Both (1) and (2) can answer a question like ‘How many moons does Jupiter have?’. But due to its focus, (1) is not an appropriate answer to other questions, like ‘Which planet has four moons?’.

Part of what makes this account plausible is the parallel between *structural* focus constructions, like (8-b) and (9-b), and *tonal* focus constructions, where focus is achieved through intonation. For example, a tonally focused variant of (8-a) would be:

(10) Johan likes SOCCER.

where the capital letters indicate a stress that is pronounced. It is hard to deny that such tonal focus constructions have the same truth conditions as their neutral counterparts.⁶ On the other hand, the communicative import of tonal focus constructions is evident from everyday conversation. Why, then, do we also have structural focus constructions, like (8-b) and (9-b)? Because not all language is pronounced. We have a communicative need for focus constructions even in settings where we can’t use words with different intonations. In such settings, we impart focus through syntax rather than intonation.

This helps explain why the significant differences in syntax between structural focus constructions and their neutral counterparts does not prevent them from having the same truth conditions. In structural focus constructions, the syntax has a communicative role over and above its usual role as the starting point of semantic interpretation, so it is not necessarily a good guide to the semantic structure of the sentence at the sub-sentential level. In particular, if (1) is a focus construction, ‘four’ there might still have the semantic role of a determiner, extracted away from the noun that it modifies (‘moons’) for the sake of focus. In the same way, ‘quietly’ in (9-b) has been extracted away from the verb it modifies (‘entered’), but it remains semantically an adverb modifying that verb. In this case, there is little temptation to treat ‘quietly’ as a singular term or anything other than an adverb. Hofweber thinks there should be equally little temptation to treat ‘four’ as a singular term in (1).

When we treat ‘four’ as a determiner rather than a singular term, and explain (1) as a focus construction rather than an identity statement, the

⁶Again, this is not quite true in general, but the exceptions need not concern us here.

substitution problems dissolve, because they are based on the assumption that both ‘four’ and ‘the number of moons of Jupiter’ are semantically singular terms. This is a good reason to think Frege’s analysis is wrong, and that we can better account for the semantics of (1) by treating it as a merely syntactic variant of (2). The syntactic difference is communicatively important, but for just this reason, the syntax of (1) is a misleading guide to its semantics.

1.2 A SYNTACTIC PROPOSAL

I agree with Hofweber’s assessment that (1) should not be construed as an identity statement. And I think Hofweber is right to draw attention to the truth-conditional equivalence of (1) and (2), as well as their different communicative roles. I think Hofweber’s observation that (2) can answer a wider range of questions than (1) is particularly important. In fact, I think it is the key to understanding the semantic role of ‘the number of moons of Jupiter’ in (1).

In this section, I want to raise a problem for Hofweber’s view. The problem is this: what is the role of ‘number’ in (1)? Attempting to solve this problem will show us that we cannot view (1) as *merely* a syntactic or structural variant of (2). This means that Hofweber’s claim that (1) is a focus construction is not quite right. But our attempt to make it work will lay the foundations for a different account, one which gives us a clearer picture of the role of ‘number’ and a family of other nouns.

1.2.1 Whence ‘the number’?

Hofweber’s analysis of the singular term uses of number words as focus constructions explains both the felt equivalence of (2) and (1) and the reason we have such constructions: they serve an important communicative function. Still, I think the analysis does not address one important question. If non-arithmetical uses of number words like (1) are merely focus constructions, why do they contain apparently contentful words that do not appear in their unfocused counterparts? Specifically, why do they contain the expression ‘the number’? The problem here is that ‘the number’ seems to contribute some important content to such statements. This is not true of Hofweber’s other examples of focus constructions, like (8-b) or (9-b), where the words which appear in the focus construction but not

the neutral sentence do not seem to introduce any new concepts. We need a way of accounting for the fact that (1) and (2) are equivalent, *despite* the fact that (1) employs the concept of number, while (2) apparently does not. Otherwise, there is room to doubt that (1) is merely a syntactic variant of (2).

One way to see the problem here is to compare Hofweber's proposal to a suggestion of Frege's. Hofweber intends his proposal to be a development of Frege's idea that we may 'recarve' the content of a sentence in a way that uses different concepts. Frege says, for example, that (11-b) is a 'recarved' version of (11-a)⁷:

- (11) a. Line *a* is parallel to line *b*.
 b. The direction of line *a* is the direction of line *b*.

(11-b) employs the concept of *direction*, where (11-a) does not, but the two sentences have the same truth conditions. Frege's idea is that the truth conditions of the sentence are prior to its analysis into individual terms and incomplete expressions, and we have some choice in the concepts we use to express these truth conditions. Thus we can, as it were, shift some of the content expressed by 'parallel' in (11-a) over to the concept expressed by 'direction' in (11-b), in order to make use of the concept of identity in the latter, without changing the truth conditions. This 'recarving' metaphor is suggestive, but Frege does not develop it further in the *Grundlagen*.

If we could further articulate what it is to 'recarve' the content of a sentence, we would have an explanation of why the concept of number occurs in (1) but not (2), despite their equivalence. Hofweber's idea is that 'recarving' is structural focus. Thus, seeing (1) as a focused version of (2) spells out Frege's recarving metaphor, so it provides such an explanation. The problem is that Hofweber argues for his proposal by assimilating (1) to examples of focus constructions which do *not* introduce new concepts into the sentence, and it is difficult to see how we could extend his proposal to cases like (11). (What aspect of the information communicated by (11-a) does (11-b) focus on?) Thus the focus construction proposal provides little insight into the sense in which recarving can introduce a new concept while holding truth conditions fixed. For that reason, it cannot fully account for the relationship of (1) to (2).

⁷See Frege (1884/1980, § 64).

As an answer to this problem, I would like to make a proposal on Hofweber's behalf. When we use a structural focus construction like (8-b), we typically have (at least) two communicative goals: to give the information that Johan likes soccer, and to stress the part or aspect of this information that answers the question of what Johan likes. To achieve these goals, we must typically use an expression that satisfies two constraints: it must be a grammatical sentence, and it must have a structural focus. The extra words are required in order to simultaneously satisfy these constraints, but they carry no independent semantic content.

Could the same be true of (1)? This sentence has several words that do not appear in (2): 'the', 'number', 'of', and 'is'. Can these be understood as merely required to meet the demands of grammaticality while achieving focus? In particular, can 'number' be understood this way? Let us see how one might spell out such a proposal. I will call this the *syntactic proposal* about 'number'.

1.2.2 *The syntactic proposal about 'number'*

We have seen that one common way of achieving structural focus is by extraction: movement of the word or phrase which contributes the information the speaker wants to focus on into a privileged position. In (8-b) and (9-b), this is a position at or near the grammatical subject of the sentence. According to the syntactic proposal about 'number', a parallel phenomenon is involved in (1). The typical communicative role of this sentence, we are granting, is as a focus construction, which communicates the same information as (2) but additionally has a structural focus on the 'how many' aspect of that information. This focus effect is achieved by extracting the part of the sentence that carries the 'how many' information—namely, the determiner 'four'—to a privileged position. In this case, the privileged position is at the end of the sentence, but we could also put it in subject position without too much awkwardness:

(12) Four is the number of moons of Jupiter.

But extracting 'four' is not as simple a syntactic operation as extracting a direct object (as in (8-b)) or an adverb (as in (9-b)). We cannot, as it were, simply move it to the front of the sentence, and replace it with a

cross-referencing word like ‘that’ or ‘how’, while preserving grammaticality:

- (13) a. *Four is Jupiter has ... moons.
b. *Four is ... Jupiter has moons.

There is evidently no way to fill in the blank in (13-a) or (13-b) that would yield a grammatical sentence. We can come close if we follow the model of (9-b), and notice that a question and its focused answer employ the same question word:

- (14) *How* did Mary enter?
(9-b) Quietly is *how* Mary entered.

A direct analogy with (9-b) is (15-a), but it is ungrammatical. Grammaticality is restored if we also move ‘moons’ to complement the question word, as in (15-b):

- (15) *How many* moons does Jupiter have?
a. *Four is *how many* Jupiter has moons.
b. Four is *how many moons* Jupiter has.

This is the first observation that the syntactic proposal about ‘number’ draws on: extracting number determiners to a privileged position requires both moving their complements (‘moons’) and introducing a new construction, such as a question word. Let’s call this new construction the *focus-forming* expression. The focus-forming expression in (15-b) is the question word ‘how many’.

Of course, (15-b) is not yet a sentence containing ‘number’. We have not yet arrived at a syntactic explanation for the occurrence of ‘number’ in (2). To get there, we need a second observation about (14), namely, that we can form structurally-focused answers using other types of focus-forming expressions besides the question word ‘how’. We can also form a structurally-focused answer using the determiner ‘the’ as part of the focus-forming expression, rather than the question word ‘how’:

- (14) How did Mary enter?
(16) Quietly is *the way* Mary entered.
(or: *The way* Mary entered is quietly.)

Why would we do this? Why not just stick to ‘how’ as a focus-forming expression when giving an answer to (14)? One answer is that, as a strategy for achieving structural focus, this approach is more flexible than ‘how’. A variety of determiners can be used in focus-forming expressions, not just ‘the’:

(17) Quietly is {a/one/the first...} way Mary entered.

The sentences in this family have different, but related, communicative uses. You might wish to tell me that Mary entered quietly, and additionally to imply that relative noiselessness was the only salient feature of her entering. You can do this with (16). But you might also want to avoid that implication; in that case, a sentence in the (17) family is more appropriate.

Using a determiner in a focus-forming expression also has important consequences for the subsequent discourse. Depending on how our conversation about Mary develops, you might also wish to highlight the similarity between how Mary entered and how Sean entered, for which you’d use (18-a). Or you might wish to contrast how Mary entered with how Sean entered, in which case you’d use a sentence like (18-b).

(18) a. Sean entered in {that/the same} way, too.
b. Sean entered in {a different/another} way.

And so on. The point is that being able to employ different determiners in your answers to questions like (14) can be quite useful. Different determiners can impart or avoid certain immediate implications, as well as enable certain types of anaphora later in the discourse.

The flexibility of using determiners in focus-forming expressions, as opposed to question words, carries a small price. Using a determiner in a focus-forming expression requires introducing a new word into the sentence. In the cases above, that word is ‘way’. This word is required for grammaticality, as a complement to the determiner. But that seems to be its only purpose. We don’t use (16) instead of (9-b) because we wish to talk about something other than how Mary entered. ‘Way’ is just a dummy noun here. We use it for the effects its determiner has on communication, not because we want to talk about a new family of objects, the ways. Thus, it is plausible that it occurs in (16) merely as a syntactic side-effect of the structural focus construction, not as a word

which makes a semantic contribution.

The idea of the syntactic proposal about ‘number’ is that ‘number’ plays a similar role in (1) that ‘way’ plays in (16) and (17). It provides a dummy noun phrase, an appropriate syntactic complement for a determiner in a focus-forming expression. We use the determiner with ‘number’ as a focus-forming expression, as opposed to the question phrase ‘how many’, because of its relationship to a family of determiners which enable more flexible communication. Specifically, sentences using ‘number’ as part of a focus-forming expression imply weaker sentences, are implied by stronger sentences, and enable contrasting sentences and discourse anaphora. It is more difficult to achieve such effects using ‘how many’ as a focus-forming expression.

To see this more clearly, consider the conversation that might develop when someone asks:

(19) How many moons orbit each planet?

Here are two ways you might answer this question in the case of Jupiter:

- (20) a. Four is how many moons orbit Jupiter.
 b. Four is the number of moons orbiting Jupiter.

Notice, however, that the latter answer is more flexible, in much the same way as (16). It is implied by stronger sentences like (21-a), implies weaker sentences like (21-b), has contrasts like (21-c) and enables discourse anaphora as in (21-d) and (21-e):

- (21) a. Four is the number of moons orbiting each planet.
 b. Four is {a/one} number of moons orbiting a planet.
 c. Two is {another/a different} number of moons orbiting a planet.
 d. Venus has that same number of moons, too.
 e. Saturn has at least that number of moons.

It is difficult to make the natural relationships between such sentences felt if you are limited to using ‘how many’ instead of using ‘number’ with a determiner in the subsequent discourse:

- (22) a. Four is how many moons orbit each planet.
 b. Four is {*one/*a} how many moons orbit a planet.

- c. Two is {*another/*a different} how many moons orbiting a planet.
- d. Venus has {*that same} how many moons, too.
- e. Saturn has {*at least that} how many moons.

In particular, there doesn't seem to be a weaker sentence (22-b) implied by (20-a) in the way that (21-b) is implied by (20-b). Nor does there seem to be a contrasting sentence (22-c).⁸

Thus, we can tentatively conclude that (20-b) is more flexible than (20-a). In a conversation where one can choose either sentence as a way of focusing on the 'how many' aspect of the information communicated, using (20-b) seems to allow the conversation to develop in ways that (20-a) does not.

To summarize the syntactic proposal about 'number', then: 'number' is a dummy noun phrase, much like 'way'. It is introduced into focus constructions like (1) because extracting a number determiner to a privileged position requires moving the noun phrase it modifies and introducing a focus-forming expression. There are at least two options for the focus-forming expression: a question word like 'how many', or a determiner-noun combination. The latter option is communicatively more flexible because of the relationships between different determiners, but using a determiner requires a noun phrase complement. In the case of (1), 'number' is that complement. It is introduced to meet the demands of grammaticality when a determiner is chosen over a question word as the means of achieving structural focus.

⁸Actually, things are not quite this simple. In some cases, there are workarounds that allow one to achieve similar effects on communication when 'how many' is the focus-forming expression. For example, the intended meaning of (22-d) can be grammatically expressed as "Venus has (exactly) *that* many moons, too" or "Venus has *just as many* moons", while the intent of (22-e) might be expressed as "Venus has *at least that many* moons". These workarounds exploit the parallels between 'the same number' and '(exactly) that many', and 'at least that number' and 'at least that many' to achieve the same effect as the focused sentences using 'number'. This works because 'that many moons' can serve as an anaphor whose antecedent is 'how many moons' in (20-a). But this strategy won't help in the cases of the weaker and contrasting sentences (22-b) and (22-c), because there are no parallel expressions with 'many' that do the work of 'a number', 'one number', 'another number' and 'a different number'. This shows that there is still an expressive advantage gained by using 'the number' as a focus-forming expression.

1.2.3 *A problem with the syntactic proposal*

I think the syntactic proposal about ‘number’, sketchy though it may be, provides an illuminating view of the role of ‘number’. If it is right, it provides a solution to the problem I raised above, that we have no account of how (1) can be equivalent to (2) given that (1) introduces the concept of number. The solution is deflationary: ‘number’ only appears to introduce a new concept; in fact, it is semantically idle, serving only to help achieve focus by syntactic means. Thus, there is no disanalogy between (1) and more straightforward cases of focus constructions like (8-b) and (9-b): none of these sentences introduces a concept absent from its neutral counterpart. So there is no reason to doubt that (1) is a structural focus construction after all.

Nevertheless, I think there is a problem with the syntactic proposal. The problem is this: even if ‘number’ occurs in (1) and related sentences merely for the sake of grammaticality, this does not show it is semantically idle. There is an important semantic difference between focused sentences which use ‘number’ as part of the focus-forming expression and sentences which don’t.

The syntactic proposal says that ‘number’ is introduced into (1) to preserve grammaticality in the face of two factors: the extraction of the number word ‘four’ to a privileged syntactic position, and the use of a determiner rather than a question word in the focus-forming expression. It claims that ‘number’ is just a dummy noun like ‘way’, which makes no semantic contribution to focus constructions like (1). Against this claim, we can make two observations. First, ‘number’ is not intersubstitutable with ‘way’, as we might expect if both are just dummy nouns. Second, ‘number’ is sensitive to the semantic properties of the extracted determiner. These observations suggest that ‘number’ has important semantic features that the syntactic proposal overlooks.

According to the syntactic proposal, there are at least two ‘dummy’ nouns in English, ‘number’ and ‘way’. If these are just dummy nouns, which serve only the grammatical purpose of providing a complement to words like ‘the’ or ‘that’, it seems we should be able to use one wherever we can use the other. This is clearly not so, however:

- (23) a. #Four is the way of moons Jupiter has.
 b. #Quietly is the number Mary entered.

These sentences, it seems, are syntactically well-formed but semantically bad. It is not clear why this should be so, if both ‘number’ and ‘way’ are merely uninterpreted bits of syntax. Intuitively, the problem with these sentences is that there is some kind of semantic disagreement between the focused expression and the noun in the focus-forming expression: four is not a *way* you can have moons, and quietly is not a *number* you can enter. This suggests that ‘number’ and ‘way’ have semantic features which can agree or disagree with the semantic features of extracted expressions.

This conclusion is further supported by a second observation. ‘Number’ cannot be used as part of a focus-forming expression when the extracted determiner is not a number word. We can see this as follows. In (2), ‘four’ is a determiner modifying ‘moons’. There is a family of related sentences that have other determiners in the place of ‘four’:

- (24)
- a. Jupiter has some moons.
 - b. Jupiter has a few moons.
 - c. Jupiter has many moons.
 - d. Jupiter has every moon.
 - e. Jupiter has a moon.
 - f. Jupiter has the moon.
 - g. Jupiter has that moon.

These sentences are all perfectly grammatical, and can be used to communicate information in a variety of contexts. But notice what happens if we try to extract the determiner, in a way analogous to (1), for the sake of focus:

- (25)
- a. Some is {?how many/*the number of} moons Jupiter has.
 - b. A few is {?how many/*the number of} moons Jupiter has.
 - c. Many is {?how many/?the number of} moons Jupiter has.
 - d. A is {*how many/*the number of} moons Jupiter has.
 - e. Every is {*how many/*the number of} moon Jupiter has.
 - f. The is {*how many/*the number of} moon Jupiter has.
 - g. That is {*how many/*the number of} moon Jupiter has.

In most of these cases, both ‘how many’ and ‘the number of’ fail as focus-forming expressions. ‘How many’ seems to work best, though still somewhat awkwardly, as a focus-forming construction when ‘some’, ‘many’

or ‘a few’ is the extracted determiner. (This is unsurprising, given the intuitive semantic relationship between these determiners and number words. Like number words, they express quantities, though not definite quantities.) But ‘the number of’ seems to be even more selective than ‘how many’. I do not see any way of rendering any of these examples clearly grammatical when ‘the number of’ is employed as a focus-forming expression. ‘The number of’ seems to be appropriate *only* when the extracted determiner is a number word.

If that is right, ‘number’ is sensitive to the semantic features of the extracted determiner, and not merely the syntax. ‘Number’ distinguishes between number determiners and other determiners, and the only plausible way of drawing this distinction is along semantic lines: number determiners *express definite quantity*, and other determiners do not. If the use of ‘number’ in focus-forming expressions is sensitive to this semantic distinction, it means that ‘number’ has some semantic features. And similar considerations show that ‘way’ has semantic features, too.⁹ These semantic features explain why ‘number’ is not intersubstitutable with ‘way’, and why it can only be used in a focus-forming expression when the focused determiner expresses a definite quantity.

Thus, we cannot plausibly maintain that ‘number’ is simply an uninterpreted piece of syntax. Unfortunately, this leaves us with the question of exactly what its semantic contribution is. Without an answer to this question, we still lack an explanation of why (2) and (1) are truth-conditionally equivalent. Because (1) uses ‘number’ while (2) does not, and because ‘number’ is not merely a dummy noun, we do not seem to have grounds to claim that (1) is *merely* a structurally-focused variant of (2). So we need a different approach to understanding the relationship between (2) and (1).

⁹For example, consider what happens if ‘the way’ is used to focus on ‘occasionally’ in “Mary entered occasionally”. Though ‘quietly’ and ‘occasionally’ are syntactically both adverbs, “Occasionally is the way Mary entered” makes no sense, presumably because ‘occasionally’ does not express anything about the *manner* of Mary’s action. Like ‘number’, ‘way’ is sensitive to such semantic distinctions among adverbs when used as a focus-forming expression.

1.3 NOUNS CORRESPONDING TO QUESTIONS

In light of these problems, I do not see any plausible way of maintaining Hofweber's proposal that (1) is a focus construction. In this section, I'd like to say where I think we went wrong, and lay out a different proposal for understanding (1) and the role that 'number' plays within it.

In spelling out the syntactic proposal about 'number', we made two crucial observations, which I think still hold good. First, we observed that there are two strategies for achieving structural focus: by using either a question word or a determiner-noun pair as a focus-forming expression. Second, we compared these two strategies, and observed that while using a determiner-noun pair did not seem to alter the *subject matter* of a focus construction, it did alter its *communicative import*. A focus construction like (16), which uses a determiner-noun pair, says the same thing as a question-based focus construction like (9-b), but it has a different effect on the subsequent discourse.

We went wrong in concluding from these observations that nouns used in focus-forming expressions, like 'number' and 'way', were just dummy nouns, syntactically required in focus constructions but making no interesting semantic contribution. As we have just seen, this conclusion was too hasty; it is unlikely that 'number' occurs in sentences like (1) for purely syntactic reasons. In fact, this conclusion was only attractive because we were attempting to see how (1) could be a structurally-focused variant of (2). Intuitively, both 'number' and 'way' *do* mean something, and they mean quite different things. Thus, we should put aside the focus construction proposal, and ask more directly: what do these nouns mean?

1.3.1 Other question words and nouns

The observations which supported the syntactic proposal about 'number' are really instances of a more general pattern. To understand the role of 'number' and other nouns, it will be helpful to have this pattern more fully in view.

Nearly every question word has the same relationship to some other noun as 'how many' has to 'number', or 'how' has to 'way'. The relationship is this: a given use of a question in an indicative sentence can be replaced by a definite description employing a noun corresponding to

that question, and the resulting sentence is felt to be truth-conditionally equivalent to the original. These question-noun relationships are summarized in Table 1.1.¹⁰ Examples of sentences which have two equivalent versions, one using a question word and one using a corresponding noun, are provided in (26)–(33).

Table 1.1: Relationships between question words and nouns

Question word	Corresponding Nouns
how	way
how many	number
how much	quantity, amount
when	time, moment, day, year...
where	place, location position
why	reason, cause, explanation
who	person
what	thing

- (26) a. Quietly is how Mary entered. (=9-b)
 b. Quietly is the way Mary entered. (=16)
- (27) a. Four is how many moons Jupiter has. (=15-b)
 b. Four is the number of moons Jupiter has. (=12)
- (28) a. One cup is how much milk is needed.
 b. One cup is the amount of milk needed.
- (29) a. Ten p.m. is when I go to bed.
 b. Ten p.m. is the time I go to bed.
- (30) a. Charlotte went back to where she first saw the spy.
 b. Charlotte went back to the place she first saw the spy.
- (31) a. We'll never know why Oscar fled the country.
 b. We'll never know the reason Oscar fled the country.
- (32) a. Who Oedipus married was, tragically, his mother.
 b. The person Oedipus married was, tragically, his mother.

¹⁰There are two question words which appear to be exceptions to this pattern: 'which' and 'whether'. Thus, they do not appear in Table 1.1 or the examples in (26)–(33). These two question words are rather special; I will have more to say about them in Chapter 4.

- (33) a. What Harry most desired was his dinner.
 b. The thing Harry most desired was his dinner.

Some terminology will be useful for characterizing the relationship we see here. *Question words* are the words (and phrases) which appear in the left hand column of Table 1.1; their *corresponding nouns* appear in the right hand column. For example, ‘why’ is a question word, and ‘reason’ one of its corresponding nouns. The relationship between question words and their corresponding nouns shows itself primarily in complete phrases. A *question* is a complete phrase headed by a question word, such as ‘how Mary entered’ in (26-a), or ‘why Oscar fled the country’ in (31-a).¹¹ A *question-replacing description* is likewise a complete phrase, headed by ‘the’ plus a noun phrase headed by a corresponding noun, such as ‘the way Mary entered’ in (26-b) or ‘the reason Oscar fled the country’ in (31-b). Each of the examples in (26)–(33) contains two versions of an indicative sentence. The *question version* contains a question. In the *description version*, a question-replacing description replaces the question. Specifically, the question is replaced by a description formed using a noun corresponding to the question word which heads the question.

Several features of this question word-noun relationship seem noteworthy here. First of all, the examples are far from exhaustive. The questions and descriptions in (26)–(33) occur in different environments, and in sentences which differ in their overall surface syntax. So far as I can see, it would not be difficult to multiply these examples; a description can be substituted for a question in many syntactic environments.

Second of all, it seems to me that in each example, the two versions of the sentence are truth-conditionally equivalent. I do not know how it could be true that one cup is how much milk is needed, but false that one cup is the amount of milk needed, or true that we’ll never know why Oscar fled but false that we’ll never know the reason he fled. Nor could I suppose that ten p.m. is when I go to bed without supposing that ten p.m. is the time I go to bed. And so on. Substituting the question with an appropriate description seems to yield a completely equivalent sentence, at least as far as truth conditions are concerned.

¹¹Note that questions differ from relative clauses, which are also headed by question words in English.

These observations generalize the observations we made earlier. They suggest that the relationship between questions and question-replacing descriptions is a productive phenomenon.¹² Question words are systematically correlated with certain nouns, and replacing a question with an appropriate description does not affect the truth conditions of the containing sentence. This relationship seems to hold for many different question words and nouns, across many containing sentences. It is not an isolated or limited phenomenon, and it is unlikely to be a coincidence.

In particular, the relation between questions and descriptions is not limited to sentences which can be construed as focus constructions. Neither sentence in (30) or (31), for example, appears to be a focus construction. These sentences do not limit the questions they can answer as narrowly as focus constructions, nor do they seem to have neutral variants. (30-b), for example, can answer both “Who went back to where she first saw the spy?” and “Where did Charlotte go?”. Thus, these sentences do not seem to stress any particular aspect of the information they convey, either through syntax or intonation. Even if some of the sentences in (26)–(33) should be regarded as focus constructions, the notion of a focus construction is not general enough to account for all instances of the relationship between question words and nouns. We cannot hold that the nouns in Table 1.1 are merely ‘dummy’ nouns which serve only to help achieve structural focus, because there are cases where they do not serve this purpose. This is another strike against the syntactic proposal about ‘number’.

To sum up: we have discovered a class of nouns which have an important relationship to question words. Questions can be replaced by descriptions formed using these nouns in a wide variety of sentences. This is the same relationship we earlier characterized in terms of two different strategies for achieving structural focus, but it now appears to extend more widely than sentences which can be construed as focus constructions. The nouns in this family include ‘number’ and ‘way’, but these are only two examples among many. There is at least one such noun corresponding to every question word in Table 1.1; some question words have several corresponding nouns.

¹²Schlenker (2003) argues that definite descriptions can be productively interpreted as questions, using evidence from French.

1.3.2 *Do nouns corresponding to questions denote objects?*

With the relationship between question words and their corresponding nouns more fully in view, we can now return to the question I raised at the beginning of this section: what do ‘number’ and ‘way’ mean? More generally, if we take it that ‘number’ and ‘way’ are part of the family of nouns which correspond to question words, what do the nouns in this family mean?

Our usual model for understanding the semantics of nouns says that nouns denote classes of objects. ‘Horse’ denotes the (class of) horses, ‘potato’ the (class of) potatoes, and ‘silicate’ the (class of) silicates. Does it make sense to say the same thing about the nouns in Table 1.1? Should we say that ‘way’ denotes the (class of) ways, ‘amount’ (the class) of amounts, ‘time’ the (class of) times, and so on?

We have already seen some intuitive reasons why doing so is unattractive. In motivating the syntactic proposal about ‘number’, for example, we noted that we do not seem to use ‘way’ because we wish to talk about a class of objects, the ways. This led us to conclude that ‘way’ was a dummy noun, at least as it appears in sentences like (16). Though this conclusion was incorrect, our willingness to draw it does point to a contrast between nouns like ‘way’ and nouns like ‘horse’.

Part of the point of saying that nouns denote classes of objects is to indicate their contribution to the subject matter of clauses where they appear. It seems obvious that there are such things as horses, and that we use the word ‘horse’ to talk about them. By saying that ‘horse’ denotes the class of horses, we indicate what sentences containing ‘horse’ are about, and which objects we must look to when determining if they are true. But nouns corresponding to questions seem to differ from ‘horse’ in just this respect. It is not so clear that sentences containing ‘way’, for example, are about ways, or that we must ‘look to the ways’ to determine if they are true. At any rate, it is not clear that we talk about them or look to them *in the same sense* as we talk about or look to horses. So we might doubt that these two kinds of noun should be given the same treatment by a semantic model for our language.

Another consideration against the usual model comes from grammar. Individual ways, amounts, times and so on are often not specified by expressions that should be treated as singular terms. Instead, they are often specified by expressions from other parts of speech, such as the adverb

‘quietly’ in (16). Individual horses and potatoes, by contrast, are typically specified using expressions which are clear examples of singular terms, such as names or deictic demonstratives. I shall return to this point in Section 1.3.3.

These intuitive considerations are not strong enough to show that the usual model does not apply to the nouns in Table 1.1, though. For one thing, they do not hold equally well for all the nouns that appear there. It is just as obvious that there are persons as that there are horses, for example, and that we use ‘person’ to talk about them and proper names to specify them. For another, we might simply deny that part of the point of saying a noun denotes a class of objects is to capture our intuitions about their relation to subject matter. Perhaps the notion of ‘object’ which is most useful for semantic theorizing has nothing to do with subject matter, and as far as *that* sense of ‘object’ goes, ordinary nouns and nouns corresponding to questions denote objects in exactly the same sense.

Still, applying the usual model to nouns corresponding to questions will not help us to understand their special features. The puzzle presented by these nouns is that they can appear in the description version of a sentence, and not in the question version, without preventing the two versions from being truth-conditionally equivalent. These nouns appear to make a semantic contribution, but not a contribution that can’t be had in their absence. The usual model won’t capture this puzzling feature of these nouns, even for nouns like ‘person’ to which it uncontroversially applies, precisely because it applies equally well to nouns that *don’t* have this feature.

To see this, consider again the sentences in (32):

- (32) a. Who Oedipus married was, tragically, his mother.
 b. The person Oedipus married was, tragically, his mother.

Suppose we follow the usual model, and grant that ‘person’ denotes the (class of) persons, and that both ‘Who Oedipus married’ in (33-a) and ‘The person Oedipus married’ in (33-b) are semantically singular terms which refer to a certain person. Even granting this, a version of the problem I raised for Hofweber’s account remains: why should it be obvious that (33-a) and (33-b) are equivalent—and hence that these singular terms co-refer—given that the latter employs ‘person’ and the former does not?

Observing that ‘person’ denotes the class of persons does not resolve

this problem. In fact, it makes it more acute, because saying that ‘person’ denotes the class of persons assimilates our semantic explanation of (33-b) to our explanations of sentences where another noun occurs in the place of ‘person’:

- (34) a. The *queen* Oedipus married was, tragically, his mother.
 b. The *woman* Oedipus married was, tragically, his mother.
 c. The *relative* Oedipus married was, tragically, his mother.
 d. The *Greek* Oedipus married was, tragically, his mother.
 e. ...

A typical semantic explanation of sentence (34-a) would compositionally derive a statement of its truth conditions from, among other things, a clause that says ‘queen’ denotes the class of queens. This explanation could be carried over to any of the other sentences in the family in (34) simply by replacing this clause with a clause for another noun. Likewise, it could be carried over to (33-b) simply by replacing the clause for ‘queen’ with a clause that says ‘person’ denotes the class of persons.

The trouble is that none of the sentences in (34) is equivalent to (33-a), and any reasonable explanation of their semantics will predict this. Each of the sentences in (34) could be false while (33-a) is true. Intuitively, the reason in each case is that the italicized noun restricts which objects the description ‘the NOUN Oedipus married’ might refer to, while ‘who Oedipus married’ is not similarly restricted: it might have been true of someone that she was *who* Oedipus married, yet false that she was a queen, a woman, a relative, or a Greek. This restricting effect of nouns will typically be represented in a semantic theory in the way that clauses for nouns compose in larger expressions.¹³ But ‘person’ differs from ‘queen’, ‘woman’, ‘relative’ and so on precisely in that it does not impose a similar restriction in this context, or any other context where a question-replacing description can occupy the same position as a ‘who’-question. Thus, saying that ‘person’ denotes the class of persons just as ‘queen’ denotes the class of queens makes the need for an explanation of

¹³For example, the denotation of the expression *queen Oedipus married* might be derived by intersecting the denotation of *queen*—that is, the class of queens—with the denotation of *Oedipus married t*. By taking the intersection with the class of queens, the theory ensures that candidates for the reference of *The queen Oedipus married* are restricted to queens.

this difference all the more pressing.

One might object here that the difference between ‘person’ and other nouns like ‘queen’ is not really so stark, because ‘person’-descriptions are not always equivalent to ‘who’-questions. Consider, for example, your exclamation to the guilty-looking dog in a kitchen where trash has been strewn about: “I see who got in the garbage!” It would be a little strange to say instead: “I see the person that got in the garbage!” Thus, it may seem, even ‘person’ introduces more restrictions than ‘who’, and so ‘who ...’ and ‘the person ...’ are not equivalent. I reply that such cases clearly involve derivative uses of ‘who’. ‘Who’-questions can *by default* be replaced by ‘person’-descriptions, though particular contextual or pragmatic factors can sometimes undermine this relationship, like the familiar relationship we have with non-human pets. At any rate, it seems clear that ‘who’ bears a *much* closer semantic relationship to ‘person’ than to any other noun. This is the phenomenon that needs to be explained, and assimilating ‘person’ to other nouns won’t help explain it.

The conclusion is not that we cannot apply our usual model for the semantics of nouns to the nouns in Table 1.1. I am only claiming that if we apply it, we will not have made any progress toward understanding the special features of these nouns. When a question is replaced by a description headed by an appropriate corresponding noun, the resulting sentence is truth-conditionally equivalent; when the same question is replaced by a description headed by an ordinary noun, it is not. Since all the usual model says about nouns is that they denote classes of objects, it provides no resources for explaining this difference.

So we still have not solved our puzzle. I leave it open, for now, whether and in what sense the usual model can be applied to nouns corresponding to questions. My point is simply that the usual model cannot explain the relationship of question words and their corresponding nouns on its own, so we have to look elsewhere if we want to understand this relationship.

1.3.3 *Meanings via questions*

We don’t yet have an account of what the nouns in Table 1.1 mean, or an explanation of why the question and description versions of a sentence are truth-conditionally equivalent. Nor do we know how far this class of nouns extends, or why our language has them at all. I will investigate

these issues in subsequent chapters. Here, I want to conclude by making a few suggestions about where that investigation should begin. We have seen that these nouns are related to question words. Perhaps we can get a grip on what they mean by looking further at how questions work.

Questions have answers, and we ask them in order to receive answers. Of course, sometimes things turn out badly: a question can be confused, or have no answer, or have no good answer so far as anyone knows. But these exceptions prove the rule. Questions are asked in *expectation* of getting an answer. In the normal case, an answer can be supplied, though perhaps only after some investigation or thought.

The remarkable thing about questions is that one need not already know the answer to a question to ask it. An answer supplies *more information* about a certain topic than the question it answers. One way to put this is to say that, whereas a question raises an issue, an answer resolves it. Another is to say that a question admits of a *range* of answers. To ask a question, you do not need to know what its specific answer is, but you do need to know what would count as an acceptable answer, or what in general the range of answers is like. That is, you must be able to distinguish between expressions which provide an answer and expressions which don't.

This is where things get interesting. How do you recognize what counts as an answer to a question you've asked? As far as surface-level syntax is concerned, answers to questions can be expressions from many parts of speech. To see this, we can look at how questions are answered in discourse. Here are some examples:

- (35) a. How many moons does Jupiter have?
b. Four.
- (36) a. How did Mary enter?
b. Quietly.
- (37) a. What does Harry want most?
b. His dinner.
- (38) a. Where did you leave your bicycle?
b. At the park.
- (39) a. When are we meeting?
b. At noon.

Here, (35-a) is answered by a determiner in (35-b), (36-a) is answered by an adverb in (36-b), (37-a) is answered by a singular noun phrase in (37-b), and (38-a) and (39-a) are both answered by prepositional phrases in (38-b) and (39-b), respectively. The discourse participant who supplies an answer to a question does not generally need to do so using any particular part of speech.

In some sense, of course, the answer to a question must be a complete sentence, because to answer a question is to assert something—to say something true or false. The most proper answer to (39-a), for example, is something like, “We are meeting at noon.” Only sentences are asserted; when subsentential expressions are used to make an assertion, as in the examples in (35)–(39), we may think of the rest of the sentence as being elided. But this should not distract us from the fact that certain parts of those sentences are more relevant than others in answering the question. They are generally the parts which supply more information than the question itself does, which is why they must be pronounced rather than elided.

What the examples in (35)–(39) show is that an expression cannot be recognized as an answer to a question by its syntax alone. There is no general syntactic category of ‘answers to questions’. Nor is there much hope of recognizing the answers to a *particular* question by their syntax. Consider, for example, that both (38-b) and (39-b) are prepositional phrases headed by ‘at’, yet (39-b) cannot answer (38-a), and (38-b) cannot answer (39-a): ‘at the park’ does not specify *when* we can meet, and ‘at noon’ does not specify *where* we can meet. Although perhaps we could reject some expressions as appropriate answers to a question just by their syntax, syntax will not in general be enough. Recognizing when an answer to a question has been given likely requires being able to make semantic distinctions between candidate answers.

The questions in (35)–(39) are all uttered in the interrogative mood, but there are also indicative sentences that pair questions with their answers. These are known to linguists as *specificational* sentences. Specificational sentences are a species of copular sentence.¹⁴ Like other copular sentences, specificational sentences have two significant parts: a pre-

¹⁴The terminology here is due to Higgins (1979). For a discussion of the taxonomy of copular sentences and different issues surrounding their semantics, see Mikkelsen (2011).

copular element, and a post-copular element. The distinguishing semantic feature of specificational sentences is that one of these elements *specifies* what is asked for by the other.

Pseudo-clefts provide one important example of a kind of specificational sentence. Their form is by now familiar:

- (40) a. How many moons Jupiter has is four.
 b. How Mary entered was quietly.
 c. What Harry wants most is his dinner.
 d. Where you left your bicycle is at the park.
 e. When we are meeting is at noon.

In a pseudo-cleft sentence, the pre-copular element is like a question, and the post-copular element supplies its answer, as the parallel between the sentences in (40) and the question-answer pairs in (35)–(39) shows. As we saw with answers to questions in discourse, the post-copular element in a pseudo-cleft sentence is not restricted to any given part of speech. Determiners, adverbs, singular terms, and prepositional phrases can all occupy the post-copular position. And as with questions in discourse, the part of speech of the post-copular element does not determine whether it is an appropriate answer to the question in the pre-copular element:

- (41) a. #Where you left your bicycle is at noon.
 b. #When we are meeting is at the park.

These sentences are syntactically well-formed but semantically bad, precisely because the post-copular phrase cannot specify the information the pre-copular phrase asks for.

The sentences in (40), of course, are much like those we were looking at above, such as the sentences in (26)–(33). The questions in their pre-copular phrases can be replaced by question-replacing descriptions. The result is a new set of specificational sentences:

- (42) a. The number of moons Jupiter has is four.
 b. The way Mary entered was quietly.
 c. The thing Harry wants most is his dinner.
 d. The place you left your bicycle is at the park.
 e. The time we are meeting is at noon.

Notice that in each case, when replacing the question, we must choose an appropriate corresponding noun to head the question-replacing description. In (42-e), for example, only ‘time’ works as a noun corresponding to ‘when’, not ‘day’, ‘year’ or ‘interval’.¹⁵

Here, then, is my suggestion. Question-replacing descriptions are more like questions than like singular terms or names. Let’s hypothesize that a question and a description which can replace it both express the same *issue*, which can be resolved by an answer in discourse. A speaker can use a specificational sentence to say that a specific answer resolves that issue. As we have seen, answers to questions may come from many parts of speech. For this reason, construing the pre- and post-copular elements in a specificational sentence as expressing issues and their answers, respectively, does not require us to construe either as a singular term.

If we adopt the suggestion that question-replacing descriptions express issues, a natural account of nouns corresponding to question words also becomes available. To use a question, a speaker does not need to know its answer. Similarly, to use a question-replacing description, she does not need to know the answer to the issue it expresses. But in both cases, she does need to have some grasp of the range of answers, and this grasp requires her to make semantic distinctions between answers and non-answers. The role of nouns corresponding to questions is to make these semantic distinctions explicit. They indicate what range of expressions will count as answers to a question or as resolving an issue.

For example, consider again the sentences in (23):

- (23) a. #Four is the way of moons Jupiter has.
b. #Quietly is the number Mary entered.

We observed above that these sentences, just like those in (41), are syntac-

¹⁵Higgins had already argued that sentences like “The number of planets is nine”, and so presumably also (42-a), are specificational sentences rather than identity statements in Higgins (1979, pp. 215–218). Since this chapter was originally written, Higgins’ point has become much more widely appreciated in the philosophical literature, in part due to interest in Hofweber’s proposal about (42-a). See, for example, Brogaard (2007), Moltmann (2013), Jackson (2013), Felka (2014), Hofweber (2014), Knowles (2015), Felka (2016). Below, I propose an alternative analysis which is similar to those of Moltmann (2013) and Felka (2016) in holding (42-a) to be a specificational sentence, and in adopting a question-answer analysis of such sentences.

tically well-formed but semantically bad. Intuitively, this is because four is not a *way* you can have moons, and quietly is not a *number* you can enter. My suggestion is that this intuition can be cashed out as follows: ‘quietly’ cannot provide an answer in the range of expected answers that ‘number’ allows, and ‘four’ cannot provide an answer in the range of expected answers that ‘way’ allows.

We can see why it is important to be able to express distinctions among ranges of possible answers by considering that some questions admit different ranges of answers, depending on the context. Some expressions, even those that can answer questions expressed by the same question word, may or may not count as expressing an answer to a particular question in a discourse, or resolving a particular issue. We use multiple nouns corresponding to a single question word to help us draw more precise boundaries between answers and non-answers than the question word itself does. To see this, consider three ways of answering a ‘when’-question:

- (43) When did you start your dissertation in earnest?
- a. The *year* I started my dissertation in earnest was 2013.
 - b. The *semester* I started my dissertation in earnest was last spring.
 - c. The *month* I started my dissertation in earnest was April of last year.

In each of these responses, the speaker could have used a ‘when’-question instead of a question-replacing description. But by using a description, she expresses her understanding of the issue that the questioner has raised with (43). With (43-a), for example, she expresses that she takes the issue to have answers in the range allowed by ‘year’, such as ‘2013’. Expressing her understanding this way serves at least one important communicative purpose: it helps her coordinate with the questioner on the issue. If his understanding does not agree with hers, he can rephrase: “Ah, no, I meant to ask, which *semester*...”. This is one way that using a noun in a question-replacing description can have an expressive advantage over using a question word in a question.

Finally, I think these suggestions also point the way toward explanations of the data which motivated an alternative analysis of (1):

1. 'How many'-questions only admit of one correct answer, so exchanging the definite article for the indefinite in a 'how many'-replacing description, as in (4), results in a mismatch of presuppositions. 'Number' indicates that a 'how many' issue is being addressed, so only one answer is possible, but 'a' indicates that more than one answer is possible. The indefinite article can be substituted for the definite in question-replacing descriptions when the corresponding question does not presuppose that only one answer is possible.
2. When embedded under 'knows', the question reading of a question-replacing description is much more natural than an acquaintance reading, because a question-replacing description expresses an issue, rather than denoting an individual with which one might be acquainted. (When a question-replacing description can also be used to refer to an individual, such as 'the person...', the preference for the question reading disappears.)
3. We should not expect to be able to substitute a question-replacing description for an answer, even in extensional contexts, any more than we would expect to be able to substitute a question for its answer. Questions and their answers are generally of different syntactic and semantic types. The specificational sentences which connect them should not be understood as identities that license substituting an issue-denoting expression for an answer, so it's no surprise that 'the number of moons of Jupiter' cannot be substituted for 'four', as seen in (7).
4. Specificational sentences like (1) can be used to answer fewer questions than neutral sentences like (2), not because they are structurally focused, but because they contain statements of the particular issue they address. They say that a particular answer is the answer to a particular issue. It is incongruous to use a sentence which specifies an answer to one issue to answer a different issue.

Thus, it seems a view of the sort sketched above provides many of the advantages of the focus construction view, without the negative consequences that follow from treating nouns corresponding to question words as having purely syntactic occurrence.

The syntactic proposal about ‘number’, then, did not lead us entirely astray. It gave us the idea that nouns like ‘number’ have an important relationship to questions, and that there may be expressive or communicative reasons to use descriptions containing these nouns rather than questions. Once we drop the problematic assumption that nouns corresponding to questions occur for purely syntactic reasons, the way is clear for developing an account of their meanings based on their relationship to questions.

There is obviously much more to be said to develop these thoughts, and to show that such an account can solve the problems with which I began. But I hope I have made it plausible that such an account has good prospects.

1.4 THE PLAN

In this chapter, I argued that we should not take Frege’s analysis of statements of number at face value. Given some of the problems we looked at, it is not obvious that sentences like

- (1) The number of moons of Jupiter is four.

should be analyzed as statements of identity, that is, as an identity sign flanked by two semantically singular terms.

I made two suggestions about how we might give statements of number an alternative analysis. First, I suggested that we should think of statements of number as part of the class of *specificational* sentences. The copula in these sentences separates expressions which behave like a question and its answer. I observed that ‘the number of moons of Jupiter’ belongs to the class of *question-replacing descriptions*. These descriptions can be productively interpreted like questions across many contexts, and in some contexts (such as under ‘know’ or ‘tell’) this interpretation is strongly preferred. But in general, neither a question nor its answer need be construed as a semantically singular term. Thus, we might doubt that ‘the number of moons of Jupiter’ and ‘four’ need to be construed as semantically singular terms in (1), contrary to Frege’s view.

Of course, if we accept that descriptions like ‘the number of moons of Jupiter’ should be treated semantically like questions, another issue arises: why does our language have these question-replacing descrip-

tions at all? Why don't we simply use the questions themselves? My second suggestion was that there may be a communicative or *expressive* advantage to using question-replacing descriptions in place of questions. Factoring the sense of a question word into a determiner and a corresponding noun makes it possible to use other determiners and nouns in their places, both of which may be communicatively advantageous. Being able to use other determiners allows a speaker more fine-grained control over the effects of her question on the discourse, such as avoiding presuppositions or enabling certain kinds of anaphora. Being able to use other nouns allows a speaker to be more explicit about what kind of answer she is looking for, and to coordinate with others about how to understand an issue.

So far, these suggestions do not add up to a concrete proposal for semantic analysis of sentences like (1). My goal in the next few chapters is to develop these two suggestions into a more precise account of the semantics of specificational sentences and the various kinds of expressions that appear within them, and the role that these expressions play in our broader linguistic practice. I will be particularly concerned with nominalizations, such as 'the number of moons of Jupiter' or 'the reason Oscar fled', which can appear as the pre-copular phrase in a specificational sentence, and which are often of significant philosophical interest.

I will then leverage the account in those chapters to answer some of the outstanding questions from the present one: do nominalizations like 'the number of moons of Jupiter' denote objects, or not? And is (1) an identity statement? It will turn out that there *is* a sense in which I think 'the number of moons of Jupiter' stands for an object, and (1) is an identity statement. But I will try to develop a sense for these claims which is compatible with the observations and suggestions I have made here, and different from the platonism that has often been attributed to Frege.

Chapter 2

Specification

Specificational sentences are a class of copular sentence. This class includes sentences like:

- (1) a. What Caesar wants is his ball.
- b. The thing Caesar wants is his ball.

One type of specificational sentence consists of pseudo-clefts like (1-a), which have an explicit question in the subject position. Another type consists of sentences like (1-b), which have descriptions in the subject position. Not all copular sentences are specificational, though. For example,

- (2) a. What Caesar wants is red and squishy.
- b. The thing Caesar wants is red and squishy.

are predicational, rather than specificational.

My goal in this chapter is to get clear on what specification is. What makes the sentences in (1) cases of specification, as opposed to predication, or something else? I will argue in Section 2.1 that specification is at least partly a semantic or interpretive phenomenon: it is a way of interpreting a copular sentence, a particular kind of meaning that copular sentences can have.

My argument that specification should be considered a semantic phenomenon is simple. There are certain well-formed sentences that have multiple interpretations or readings, only one of which is specificational. One such sentence is this example, from Higgins (1979, p. 7):

- (3) What John is is enviable.

On a specificational reading, this sentence specifies one of John's qualities: it says that John is enviable. On a predicational reading, it says that John's position or role is enviable. If John is vice president, for example, it says that being vice president is enviable, but it does not imply that John himself is enviable. This shows that one cannot in general determine whether a sentence counts as specificational just by looking at its syntax or grammar. Instead, whether a sentence is specificational is a matter of how it is interpreted in context. Thus, characterizing the specificational reading is a task for semantic theory.

This argument, which I discuss in more detail in Section 2.1, appeals only to the existence of pseudo-cleft sentences which are ambiguous between a specificational reading and a predicational reading to show that specification is a semantic phenomenon. But this conclusion is only really interesting because there are many other kinds of sentences which have a specificational reading, and because there are other readings for copular sentences besides the specificational and predicational ones. The latter part of the chapter lays the foundations for a systematic semantic account of specification, one which makes sense of its relationship to the other readings, and which accounts for the wide variety of sentences that have a specificational reading. I survey the other readings in Section 2.1.3, and the other types of specificational sentences in Section 2.2. These surveys provide the data which show that the question-answer analysis is a more illuminating model of the specificational reading than two existing alternatives, as I argue in Chapter 3. It also points the way to an account of the epistemological and communicative role of specificational sentences, which I will develop in Chapter 4.

2.1 AMBIGUITIES IN PSEUDO-CLEFT SENTENCES

I introduced the idea of a specificational sentence in the last chapter by looking at pseudo-cleft sentences, such as those in (40). But while some pseudo-cleft sentences are paradigmatic examples of specificational sentences, not every pseudo-cleft sentence has a specificational reading. Since at least Higgins (1979), linguists have recognized a distinction between *specificational* and *predicational* readings for pseudo-cleft sentences, and observed that some sentences are ambiguous between the two. It is

worth looking at these different interpretations in order to see what distinguishes the specificational interpretation.

The following pair of pseudo-clefts illustrates the distinction between the specificational and predicational readings:

- (4) What I bought for mother was some flowers. (specificational)
 (5) What I bought for mother was expensive. (predicational)

These sentences are unambiguous; each has only the reading indicated. Sentence (4) is a specificational sentence: it specifies that what I bought for mother was flowers. Sentence (5), on the other hand, does not specify what I bought for mother, but rather describes it as being expensive. It predicates expense of what I bought for mother, whatever that was; but it does not say which thing or things I bought.

My main goal in this section is to show that some pseudo-clefts are ambiguous between these two readings. To that end, I provide some tests that help distinguish the two readings, which will also be useful later on. I then argue that the existence of pseudo-clefts which are ambiguous between these readings shows that whether or not a sentence is specificational is not a matter of its grammar, but of how it is interpreted. This shows that specification is a semantic phenomenon: saying what it is for a sentence to have a specificational interpretation, and what that interpretation consists in, falls within the domain of semantic explanation. At the end of this section, I will introduce two other readings for pseudo-clefts, the equative reading and the issue-predicational reading. These are not crucial for my main line of argument here, but they help strengthen it. They will also help reveal the virtues (and vices) of some semantic analyses of specification in Chapter 3.

2.1.1 *Two intuitive tests*

The difference between the specificational and predicational readings of a copular sentence has to do with how the pre-copular phrase relates to the post-copular phrase. For the sake of brevity, I will call the pre-copular phrase the *subject* and the post-copular phrase the *complement*. In a predicational sentence, the complement describes or is true of the subject; in a specificational sentence, the complement specifies the subject. But what is the difference between ‘describing’ and ‘specifying’? There are two

ways of bringing out this difference that I will rely on here, the list analogy and the directness test. These give us some intuitive tests for whether a sentence has a specificational reading.

In formulating these tests, I am drawing on discussion in Higgins (1979, Ch. 1 and 5), Mikkelsen (2005), and den Dikken, Meinunger, and Wilder (2000). The reader should keep in mind that these tests are meant to be intuitive. They can help English speakers who already intuitively understand the difference between specification and predication to disentangle these two readings in particular cases, and bring the distinction more sharply into focus. But they will not always yield unambiguous results, and they are not intended to unexceptionably track the theoretical distinction. They are instead a ladder to be thrown away once a theoretical account of specification is available.

With that caveat, here are the tests:

The list analogy. Sentences with a specificational reading often seem to work like lists, where the subject gives the heading of the list, and the complement gives the item or items on the list. When expressed in sentences, lists are often introduced by a colon. To get the specificational reading of a sentence like (4), it can help to think of it as a one-item list, and imagine a colon appearing after the copula:

(6) What I bought for mother was: some flowers.

There is a clear parallel here between the specificational sentence (6) and sentences which give multi-item lists:

(7) What I bought for mother was: some flowers, a card, and an extraordinary broach.

Indeed, for my purposes here, lists like (7) are clear cases of specificational sentences, with or without the colon.

In contrast to specificational sentences, predicational sentences like (5) do not work like lists, as shown by the fact that inserting a colon is ungrammatical:

(8) *What I bought for mother was: expensive.

The problem here seems to be that ‘expensive’ does not say what or which things go under the heading of ‘what I bought for mother’. What belongs on this list is left open by (5). Instead, the role of the complement is to say something about *whatever* belongs on that list, namely, that it cost rather a lot.

The directness test. Another way to think of the contrast between the specificational and predicational readings is to consider the contrast between direct and indirect answers to questions. With respect to the right question, a specificational sentence gives a direct answer, while a predicational sentence gives only an indirect one.

When a speaker asks a question, a *direct* answer is an assertion made in response that is fully cooperative and provides some information about the issue expressed by the question. For example, in response to

(9) What did you buy for mother?

paradigmatically direct answers would be:

- (10) a. Some flowers.
 b. I bought some flowers for mother.
 c. One thing I bought was some flowers.

Indirect answers, by contrast, do not provide information directly relevant to answering a question. Clear cases of indirect answers to (9) are:

- (11) a. I didn’t buy anything for mother yet.
 b. I’m not telling you.
 c. I’m not saying, but it was expensive.

None of these sentences directly answers (9). The first rejects a presupposition of the question. The second and third are pragmatically uncooperative: though they do not reject the question, they do not provide information that helps answer it. In the last case, the speaker provides some information that may or may not help the audience *infer* what she bought for mother (namely, that what she bought was expensive), depending on what other knowledge they have. But what she says leaves that question just as open as if she had not given this information, so her answer still counts as indirect.

Which question should we look at to determine if a sentence directly answers it, and therefore has a specificational reading? In the case of a pseudo-cleft, this is easy: the relevant question appears explicitly as its subject. The relevant question for (4) and (5) can be expressed by (9), for example (after adjusting the indexical 'I'). With respect to this question, the specificational sentence (4) is a direct answer, very much like (10-b) or (10-c), while the predicational sentence (5) is an indirect answer, akin to (11-c):

- (9) What did you buy for mother?
- (4) What I bought for mother was some flowers. (specificational, direct)
- (5) What I bought for mother was expensive. (predicational, indirect)

This holds true generally: a specificational pseudo-cleft provides a direct answer to the question that appears as its subject; a predicational pseudo-cleft provides at best an indirect answer.

We will shortly see examples of specificational sentences where the subject is not an explicit question. In these cases, finding a question which reveals the sentence to have a specificational rather than predicational reading is a more subtle matter, because there are some questions which admit predicational sentences as direct answers. I shall describe how to find such a question in more detail below.

The list analogy and the directness test give us two ways to show that a sentence has a specificational reading, and to distinguish that reading from a predicational one. They make use of different grammatical devices to bring out the specificational reading, and in this sense they are complementary. At the same time, it seems clear that they reveal the same phenomenon. Higgins' terminology is helpful here: questions and list headings both express *constraints*. Directly answering a question, or placing an item on a list, are ways of *saying what satisfies* such a constraint. A specificational pseudo-cleft says what or which things satisfy a constraint. A predicational pseudo-cleft describes or says something about whatever satisfies the constraint expressed by its subject, but it does not specify what satisfies that constraint.

2.1.2 *An ambiguity between specification and predication*

With these tests in hand, we are now in a good position to see that some pseudo-clefts are ambiguous between the specificational and predicational readings. Consider again Higgins' example (3):

(3) What John is is enviable.

Let's look at the directness test first. The relevant question is:

(12) What is John?

This question is ambiguous between several readings. On one reading, it is asking for John's qualities. In that case, it could be directly answered by sentences like the following.

- (13) John is...
- a. tall and rich.
 - b. a good friend.
 - c. enviable.

On this understanding of the question, (3) directly answers it, and so counts as a specificational sentence.

In another context, the question (12) might be interpreted as asking for John's position or role; in that case it can be answered by sentences like:

- (14) John is...
- a. the vice president.
 - b. the night watchman.
 - c. the first chair violinist.

Understood this way, the question is *not* directly answered by (3). An assertion of (3) in response to this question is an indirect answer, and has the predicational reading. It says that, whatever position John has, that position is enviable; but it does not say which position this is.

The list analogy confirms these results. Thought of as the heading of a list, the subject phrase 'what John is' is ambiguous in just the same way as the question (12). When the subject is thought of in the first way, as the heading of a list of John's qualities, then 'enviable' potentially belongs on

the list, and inserting a colon in (3) is appropriate:

(15) What John is is: enviable.

But when the subject is thought of in the second way, as the heading of a list of roles or positions which John might have, (3) can only be interpreted as describing this position as enviable, in which case inserting the colon is inappropriate.

It's easiest to see the ambiguity in pseudo-clefts which have 'what' questions as the subject and an adjective for a complement, like (3). But these are not the only examples. Higgins also gave some examples of ambiguous 'what' pseudo-clefts which have a noun phrase as the complement. Sentence (16) has this feature¹:

(16) What we saw in the park was a Labrador and a Poodle.

The list analogy is helpful with this example. In (16), the specificational reading lists two dogs we saw in the park: it says we saw at least one Labrador, and at least one Poodle, in the park. The predicational reading says that we saw a dog in the park which was both a Labrador and a Poodle. It says we saw a cross-bred Labradoodle in the park.

There are also ambiguous pseudo-clefts where the subject employs a question word other than 'what':

(17) Where John's going is the pits.

(18) How John got hurt was an accident.

On the specificational reading, (17) says John is going to the pits; it specifies the place John is traveling to. On the (slangy) predicational reading, the sentence says that the place John is going is depressing or boring, though it leaves open which place that is.² Likewise, the specificational reading of (18) specifies the cause of John's injury, namely, an ac-

¹This example is a modified version of one from Higgins (1979, p. 11).

²If the metaphorical character of the predicational reading in this example seems worrisome, here is another example: "Where John's going is the best restaurant in town." On the specificational reading, when the interpreter understands which restaurant is meant, the sentence specifies John's destination. On the predicational reading, it says that some restaurant—whichever one John is going to—is the best in town. The two readings are easier to pry apart in (17), so I will rely on that example below, but 'where' pseudo-clefts can be ambiguous even when no metaphor is present.

cident. The predicational reading says that John's injury was not caused intentionally. As with 'what' pseudo-clefts, the different readings can be brought out by the intuitive tests, though I leave this to the reader.

The ambiguity exhibited by sentences like (3) or (18) should be distinguished from a simpler kind of ambiguity, of the sort in (19):

(19) Where I went first was the bank.

In this sentence, the lexical ambiguity of 'bank' does not give rise to different kinds of readings for the whole sentence. The ambiguity is *localized* to the complement. That is, neither reading for the complement forces a different interpretation of the subject, or a different interpretation of the relationship of the complement to the subject. The natural readings of this sentence are both specificational: they differ only in whether they specify my destination to be a financial institution or a spot down by the river.

The examples above are not like this. (16), (17) and (18) have ambiguous complements, but the ambiguity 'infects' the whole clause: interpreting the complement one way or the other also requires interpreting its relationship to the subject differently.³ This observation is in line with what I suggested above: specification and predication appear to be different ways of relating the complement to the subject when interpreting a whole copular clause. In this sense, the ambiguity is *structural* rather than merely lexical. Specification, like predication, is a way of interpreting a whole sentence, not just a way of interpreting one of its parts.

This brings me to the main claim of this section: specification is a semantic phenomenon, something that requires an explanation within the domain of semantic theory. I think a simple argument is enough to establish this claim in the sense I intend it. The task of semantic theory is to systematically describe and explain the different ways sentences can be interpreted. Since there are sentences, such as (3), (16), (17), and (18), that admit of both specificational and predicational interpretations, it is clear

³In (3), the ambiguity seems more localized to the subject, but I think this is not the best way to describe the case. To see this, consider the following schema: x envies y because y is F . When a sentence of this form is true, both y and F might be said to be 'enviable', but in different senses. On the specificational reading of (3), the subject plays the role of y in this schema, while on the predicational reading, it plays the role of F . The two interpretations of the subject are thus of different types; if the interpretation of 'enviable' is held constant, its relation to the subject must vary.

that specification is one way of interpreting a sentence, but not the only way. Thus, it falls to semantic theory to describe the specificational reading, and explain how it differs from other ways of interpreting a copular sentence.

This may seem too quick. How do we know that there is not some other, non-semantic explanation of the ambiguity between specification and predication? Perhaps, for example, the ambiguity can be explained in purely syntactic terms. To make this plausible, consider a parallel sort of ambiguity in English: quantifier scope ambiguity. A sentence like

(20) Some dog is loved by every boy.

is ambiguous between two readings, on which the universal and existential quantifiers take different relative scopes:

(21) $(\forall x : Bx)(\exists y : Dy)(xLy)$
(Every boy loves some dog or other.)

(22) $(\exists y : Dy)(\forall x : Bx)(xLy)$
(There is some particular dog which every boy loves.)

Like the ambiguity between specification and predication, such scope ambiguity is structural, not merely lexical. The ambiguity is not localized to any particular expression; rather, it concerns how two different expressions in the sentence are interpreted relative to one another in the whole clause.

It is typically assumed that quantifier scope ambiguity can be explained in purely syntactic terms. Such an explanation makes two important claims. First, it posits two different parses, or ‘logical forms’, for the sentence in (20). These logical forms are distinct syntactic structures, licensed by different syntactic rules in the grammar of English, but pronounced and written in the same way. The structures will differ at least with respect to the relative scopes of the quantifiers. Second, it assumes that the semantic component of the grammar interprets logical forms, rather than the pronounced form of the sentence. Quantifier scope ambiguity thus arises in a sentence like (20) because there are two different starting points for interpreting it, not because there are two ways of interpreting the same starting point. The semantic ambiguity traces back to a syntactic ambiguity, a difference in logical form. So quantifier scope ambiguity is a syntactic, rather than semantic, phenomenon.

This standard picture motivates an objection to my argument that specification is a semantic phenomenon. Couldn't there be a similar explanation of the ambiguity between specification and predication? Such an explanation would claim there are distinct logical forms for the specificational and predicational readings of sentences that exhibit the ambiguity, licensed by different syntactic principles in the grammar of English. Second, it would assume that these logical forms are the starting points of semantic interpretation. Thus, the ambiguity between specification and predication in sentences like (17) or (18) would again trace back to a syntactic ambiguity between different logical forms. If such an explanation were right, it would show that specification could be distinguished from predication in purely syntactic terms. In that case, we might not need any particular semantic account of what it is to interpret a sentence specificationally; we might only need to apply general semantic principles to distinctly specificational logical forms.

I do not mean to rule out an explanation of specification which involves distinct logical forms for sentences ambiguous between the specificational and predicational readings. But I do think any such explanation is unlikely to count as 'purely syntactic' in a way that threatens my claim that specification is a semantic phenomenon. The problem with the objection is that it assumes all ambiguities in logical form have a purely syntactic explanation, and can be derived just from rules governing the syntactic categories of English expressions, without appealing essentially to facts about how these expressions are interpreted. Such an explanation of the specification-predication ambiguity seems unlikely, given facts we have already reviewed.

To see this, consider that the syntactic category of the complement does not predict whether a pseudo-cleft exhibits the structural ambiguity between specification and predication.

- (23) a. Where John is going is the pits. (specification/predication)
 b. Where John is going is the bank. (lexical only)
 c. Where John is going is the park. (no ambiguity)
- (24) a. What John is is enviable. (specification/predication)
 b. What John is is slow. (lexical only)
 c. What John is is angry. (no ambiguity)

In the examples in (23) and (24), the first sentence exhibits the structural ambiguity between specification and predication. The second exhibits a simple lexical ambiguity, but no structural ambiguity. The third exhibits no ambiguity at all. But in each case, the three sentences differ only at the final word.

These examples illustrate an important difference between quantifier scope ambiguity and the ambiguity between specification and predication. Quantifier scope ambiguity is isolated to sentences containing words in a particular syntactic category. This makes it possible to formulate syntactic rules (such as ‘quantifier raising’ rules) that apply to all and only the expressions in that category, which allows licensing multiple logical forms for sentences containing quantifiers, without over-generating logical forms for sentences that don’t.

But as (23) and (24) show, the ambiguity between specification and predication arises in pseudo-clefts with complements in different syntactic categories (namely, NPs in (23) and APs in (24)), and only for some expressions within those categories. A purely syntactic explanation of the ambiguity between specification and predication would need to license distinct logical forms for the first sentence in each example without over-generating distinct logical forms for the other two. Doing so requires that the syntactic rules distinguish the first sentence from the other two just based on the syntactic categories governing the final word. But what purely syntactic criterion could simultaneously distinguish ‘pits’ from ‘bank’ and ‘park’, and ‘enviable’ from ‘slow’ and ‘angry’? It seems unlikely that we could justify rules that make such a distinction solely by appealing to syntactic facts.

Instead, the most plausible justification for distinguishing the first sentence from the other two in (23) and (24) would appeal to semantic facts, such as that ‘the pits’ can be *interpreted* both as specifying a place or as a predicate true of places, while ‘the park’ can only be interpreted as specifying a place. But if semantic facts like these are part of the reason for assigning such sentences two different logical forms, then positing distinct logical forms does not yield a purely syntactic explanation of the ambiguity between specification and predication. Rather, the fact that some sentences are ambiguous between the specificational and predicational interpretations is prior: the semantic ambiguity explains our identification of distinct logical forms, not the other way around. Insofar as we posit different logical forms for these sentences in order to represent

the semantic ambiguity, we give up on the idea that ambiguities in logical form can be explained purely by appealing to syntactic principles. And since we most likely must appeal to semantic facts to justify licensing distinct logical forms for sentences like (23-a), there is room for an independent semantic description of what it is to interpret such sentences specificationally, rather than predicationally. That is the project I am pursuing here.

All this is not to say that specificational sentences do not have interesting syntactic, grammatical, or other non-semantic properties. In fact, they do. Many of these properties have been studied by linguists, and there is ongoing debate about the proper syntactic analysis of specificational pseudo-clefts.⁴ In claiming that specification is a semantic phenomenon, I am not claiming that these investigations are ungrounded or uninteresting. I am only claiming that syntactic features do not *exhaust* what it is to be a specificational sentence. The ambiguity between specification and predication which I have examined in this section shows that specification is a way of interpreting the relation between the subject and the complement in a copular sentence, not merely a unique syntactic construct. Thus, whatever the best account of the syntax of specificational sentences turns out to be, the way we interpret them should also be described and characterized by semantic theory.

2.1.3 *Two other readings for pseudo-cleft sentences*

Recognizing sentences like (3) as ambiguous between the specificational and predicational readings suffices to make the point that specification is a semantic phenomenon. The point may be strengthened, however, by noting that the specificational and predicational readings do not exhaust the possible interpretations of pseudo-clefts, or copular sentences more generally. This underscores the point that the specificational reading is just one reading among many, and hence that it belongs in the domain of semantic explanation. It thus helps motivate a systematic semantic

⁴Apart from Higgins (1979), the syntax of specificational pseudo-clefts has been discussed by Ross (1972), Moro (1997), Heycock and Kroch (1999), den Dikken, Meinunger, and Wilder (2000), Schlenker (2003), Mikkelsen (2005), Romero (2005), and many others. Most of this discussion revolves around the so-called *connectivity problem*, the fact that constructions such as negative polarity items and reflexive pronouns are unexpectedly licensed in the complement position of a specificational sentence.

analysis. The existence of other readings for copular sentences raises the question of how all these interpretations are related, and whether they can be described by a unified semantic account.

I would like to call attention to two additional readings here, because they will be important later on. One is an *equative* or ‘generalized identity’ reading. A paradigmatic example might be:

(25) What I did yesterday is what you’ll do tomorrow.

In this sentence, the relationship of the complement to the subject does not appear to be the same as in the predicational or specificational readings. Rather, the sentence appears to equate two constraints. It says that whatever satisfies ‘what I did yesterday’ will also satisfy ‘what you’ll do tomorrow’, and vice versa. But the sentence leaves open exactly what sort of action satisfies these constraints, as our intuitive tests show: in most contexts, it would be strange to write ‘what you’ll do tomorrow’ on a list of things I did yesterday; and (25) does not seem like a direct answer to “What did you do yesterday?” Thus this sentence does not quite seem specificational—though it could be in a context where what you’ll do tomorrow is commonly known. On the other hand, the sentence does not exhibit the semantic asymmetry of a predicational sentence: it does not seem right to say that the complement describes or is true of what the subject denotes, any more than the other way around. Thus the sentence does not seem predicational.

Linguists debate how the equative reading relates to the specificational and predicational readings. Indeed, one possible semantic analysis of specificational sentences is that they are a species of equatives. I think there is something right about this analysis, and I will discuss it further in Chapter 3, after I have presented some additional kinds of specificational sentences in Section 2.2. For now, I will assume that the equative reading is *not* distinct from the specificational reading, mostly to ease exposition, although I will qualify this claim later.

Another reading that will play a role in my argument is one that I call *issue-predicational*. I have not seen this reading discussed elsewhere in the literature. A typical example is:

(26) What I bought for mother is not a question I’m willing to discuss.

This sentence is predicational, in that the complement is predicated of

what the subject denotes. But it differs from predicational sentences like (5) in that the subject denotes a question or issue, rather than something I bought for mother. Suppose that what I bought for mother was some flowers. In that case, (26) does not imply that

(27) #Some flowers is not a question I'm willing to discuss.

whereas the predicational sentence (5) *does* imply that

(28) Some flowers are expensive.

This indicates that the subject 'what I bought for mother' is interpreted differently on the predicational and issue-predicational readings.

A few sentences are ambiguous between the specificational, predicational and issue-predicational readings simultaneously. An example is:

(29) What John is is superfluous.

The ambiguity can be brought out via three different paraphrases:

- (30) a. John is superfluous. (specificational)
 b. John's position is superfluous. (predicational)
 c. The question of what John is is superfluous. (issue-pred.)

The ambiguity between the specificational and predicational readings of (29) is parallel to the ambiguity of (3), which I discussed above. But unlike (3), (29) also has an issue-predicational reading, as shown by the paraphrase in (30-c). One might use (29) to say that the *question* of what John is is superfluous, and so we should talk about something else.

To sum up: we have found that the specificational reading is but one possible reading among several for pseudo-cleft sentences, which shows that whether a sentence is specificational is a matter of how it is interpreted. On the specificational interpretation, a sentence says what or which thing satisfies a constraint. The distinction between the specificational and predicational readings is particularly important, and I have discussed two intuitive tests, the list analogy and the directness test, to help make this distinction clear. In the next section, I will extend this distinction to some other copular sentences by applying these tests. The equative and issue-predicational readings will not be important for that

discussion, but they will re-appear later on, because I aim to provide a semantic analysis for specificational sentences that makes sense of its relationship to each of these other readings.

2.2 OTHER TYPES OF SPECIFICATIONAL SENTENCES

Now that we have a better grip on the specificational reading in the paradigm case of pseudo-clefts, I would like to survey a wider range of specificational sentences. This survey will help us see how far an account of specification must extend, and what a successful theory must explain. I will argue in Chapter 3 that a particular theory, the question-answer analysis, best makes sense of these data.

The variation among specificational sentences that interests me most here concerns the subject phrase. In specificational pseudo-clefts, the pre-copular phrase is an explicit question, an expression headed by a *wh*-word. But linguists have recognized specificational sentences with other types of subjects for basically as long as they have recognized the category of specificational sentences. For example, Higgins (1979, p. 11) claimed that

(31) Nixon's peace plan is a bomb.

has an unfortunate specificational reading, on which the complement specifies or gives the content of Nixon's peace plan.⁵ This sentence has a possessive description rather than an explicit question in the subject position. Subsequent work has also recognized other kinds of expressions, such as anaphoric pronouns, as able to occupy the subject position of a specificational sentence. My primary goal in this section is to survey the variety of expressions that can give rise to a sentence with a specificational reading when they occupy this position.

2.2.1 Descriptions

Let's focus on descriptions first. By 'description', I mean any noun phrase headed by a determiner—what linguists now call a *determiner phrase* ('DP'). Which descriptions can occupy the subject position of a specificational

⁵The sentence also has a slangy predicational reading, on which what is said is that Nixon's peace plan is disastrous.

sentence? There are quite a variety. We already saw one class of descriptions that can do so in Chapter 1: the question-replacing descriptions, definite descriptions in which the head noun is a noun corresponding to a question word.⁶ But these are a relatively small family. In fact, definite descriptions seem generally able to occupy the subject position in a specificational sentence. Here are some examples:

- (32) a. The youngest student is Elizabeth.
 b. The claim he made was that he was innocent.
 c. The height of the building is 100 feet.
 d. The marital status she uses for tax purposes is single.
 e. The top priority is getting re-elected.

It should be clear that each of the sentences in (32) has only a specificational reading. In each case, inserting a colon between the copula and the complement would be appropriate. Thus these sentences have a specificational reading, according to the list analogy. And in each case, the complement phrase either cannot function semantically as a predicate (as in (32-a), (32-b) and (32-c)⁷), or it cannot sensibly be predicated of what the subject would denote in a predicational sentence (as in (32-d) or (32-e)), so the predicational reading is unavailable. I do not see any other potential readings for these sentences.

As I noted above, the directness test can only be applied to these sentences if we can find a question for which the specificational reading would give a direct answer, while the predicational reading would give an indirect answer. As a first pass, we can say that a suitable question can be formed by prefixing the subject with ‘what is’ (adjusting for tense and number as necessary). For example, (32-b) can be seen to be specificational because it directly answers

- (33) What was the claim he made?

⁶These nouns are listed in Table 1.1.

⁷If it seems unclear that in the case of (32-c), ‘100 feet’ cannot function semantically as a predicate, ask yourself what it would mean for ‘100 feet’ to be *true of* an object. The related expression ‘100 feet *tall*’ is semantically a predicate, but when this predicate is the complement of (32-c), I hear the sentence as having a (necessarily false) predicational reading, not a (possibly true) specificational one.

This way of forming a suitable question will not work in all cases, however. For one thing, a 'what' question is unsuitable for bringing out the specificational reading of a sentence whose complement refers to a person, such as (32-a). Consider the relationship of this sentence to the following question:

(34) What is the youngest student?

(32-a) does not seem like an appropriate answer to this question at all, much less a direct answer. Instead, this question can be directly answered by sentences like

(35) The youngest student is...

- a. very attentive.
- b. older than the teacher.
- c. Danish.

which are clearly predicational, rather than specificational. Rather than (34), the question which shows (32-a) to have a specificational reading is:

(36) Who is the youngest student?

because (32-a) directly answers it, while predications like those in (35) do not. Thus, when the complement of a copular sentence refers to a person, the 'who is...' question formed from its subject should be used in the directness test.

The ambiguity of 'what' can also obscure the specificational reading of sentences where the complement does not refer to a person. Consider, for example,

(37) What is the number of moons of Jupiter?

Usually, this question should be answered by a specificational sentence like

(38) The number of moons of Jupiter is four.

But in some contexts, it is more appropriate to answer (37) with a sentence which is most naturally interpreted predicationally, like

(39) The number of moons of Jupiter is even.

For example, this answer is appropriate when the questioner introduces her question (37) by saying: “Some numbers are odd, and some numbers are even.” This shows that passing the ‘what is...?’ version of the directness test is not always sufficient to show that a sentence has a specificational reading; we must also observe that the sentence has no predicational reading. This may not be possible in many cases when the subject is a description.

To avoid this problem, when a copular sentence has a description as its subject, it is frequently helpful to use a ‘which’ question instead. To form this question, use the schema

(40) Which *N* is *S*?

where *S* is the full subject of the sentence to be tested, and *N* is the head noun of that subject. Such ‘which’ questions cannot be answered by predicational sentences of the form

(41) *S* is *P*.

(where *S* is held constant between (40) and (41)), so if the copular sentence in question answers the ‘which’ question directly, it is generally clear that it has a specificational reading. For example,

(42) Which student is the youngest student?

(43) Which number is the number of moons of Jupiter?

can be used to test (32-a) and (38), respectively.⁸

Keeping these refinements of the directness test in mind, we can see that other forms of definite noun phrase, including possessives, plurals,

⁸The reason this works is that ‘which’, unlike ‘what’, always has to be paired with a noun; this forces a reading of the question on which its answers have the form ‘*S* is ...’ where the gap must be filled by an expression denoting the same kind of thing as *S* itself, thus ruling out predicational answers. Thus, the same result can be achieved with ‘what’ if we pair it with the same noun *N* as in (40): “What *N* is *S*?” will also rule out predicational answers. The problem is just that an *unadorned* ‘what’ can also introduce a variable over the predicate in a predicational sentence, as in (37) and (39). For more on the relationship between ‘which’ and the other *wh*-words, see Chapter 4.

and universal quantifications, may also occupy the subject position of a specificational sentence:

- (44)
- a. My youngest daughter is Elizabeth.
 - b. Every inmate's claim was that he was innocent.
 - c. (There is a new building on the waterfront.) Its height is 100 feet.
 - d. Your marital status is single, at least for tax purposes.
 - e. Their top priorities are getting re-elected, winning a majority in the House, and passing banking reform.

But definiteness is not a requirement. Indefinite descriptions can occupy this position as well:

- (45)
- a. One place to get a nice Chianti is the trattoria on Columbus.
 - b. A book that rewards careful study is Plato's *Republic*.
 - c. A few good routes to campus are Telegraph, Euclid, and College.
 - d. Some of my best students were Gareth, Valeria, and Angela.
 - e. Some great subjects to start with are mathematics and biology.

Again, in the examples in (44) and (45), the predicational reading does not seem to be available, either because the complement cannot function semantically as a predicate, or because interpreting it as a predicate applying to the subject is non-sensical. Nor are there other readings. The preferred reading of each sentence is the specificational one, which may be confirmed using both the list analogy and the directness test.

Note that in cases where an indefinite descriptions is the subject, the 'which' version of the directness test is not applicable. This is because asking a 'which' question of the form in (40) seems to require a description with a definite determiner in the place of *S*. Consider the 'which' question for (45-a), for example:

- (46) *Which place is one place to get a nice Chianti?

This question, if not ungrammatical, is at least extremely odd. For the sentences in (45), however, the 'what' version of the directness test is sufficient.

There seem to be few restrictions on the sort of description that can serve in the subject position of a specificational sentence. In the examples above, the descriptions employ determiners and nouns of several types. The determiners can be definite or indefinite, and singular or plural; and they may be articles ('the' or 'a'), possessives ('my', 'your', 'its'), or even quantifiers ('one', 'every', 'some', 'a few'). The head nouns can be abstract ('height', 'marital status') or concrete ('book'), and the whole noun phrase can be simple ('height') or complex ('place to get a nice Chianti').

I have been unable to find examples of descriptions that can never give rise to a specificational reading when they appear in the subject position of a copular sentence. It is true that some descriptions make it more difficult to get this interpretation than others. One class of descriptions which are difficult to interpret this way is simple, concrete indefinites, like 'a dog'. In the subject position of a non-predicational copular sentence, they tend to sound like awkward Yoda-speak:

(47) ??*A dog is Caesar.

Notably, this awkwardness does not attach to more complex indefinites, even when the head noun is concrete, as (45-b) shows. But even these simple indefinites seem to give rise to specificational sentences when appropriately contextualized:

- (48) a. I want to adopt a pet, but I don't know what kind of animal I want. What have you got available?
 b. ?Well, a dog is Caesar. He's very gentle and would make a great pet. (A cat we've got is Francisco, and a rabbit is Horace, but they're both a bit more ornery.)

Given the context in (48-a), the use of (47) in (48-b) sounds mostly felicitous to me, though somewhat strained. Its acceptability seems to depend on hearing the speaker as using 'a dog' to contrast the option of adopting Caesar with that of adopting another kind of animal. I do not have a good explanation for why these simple indefinites are so hard to get a specificational reading from, especially given that complex indefinites seem to work just fine. But if examples like (48) are sometimes felicitous, it seems like even simple concrete indefinites can in principle serve as the subject phrase of a specificational sentence.

Descriptions headed by strongly quantificational determiners, such as ‘both’ or ‘few’, are also often difficult to interpret specificationally:

- (49) a.??*Both authors are George.
b.??*Few votes are ‘nay’.

Other theorists, like Mikkelsen (2005, Cf. pp. 112–113), seem to think that such descriptions can never be specificational subjects. Still, like (47), descriptions like these seem mostly felicitous to me as specificational subjects when appropriately contextualized:

- (50) a. I’ve been getting two sorts of anonymous letters. In one sort, the author writes beautiful, flowing prose. In the other sort, the author’s thoughts are barely coherent and quite disconnected. I recently found out who the authors of these mysterious letters were. I expected them to be different, but as it turns out, both authors are George.
b. It appears that the budget will be nearly-unanimously approved. The roll-call is still ongoing; but so far, most of the votes have been ‘aye’, with a handful of legislators abstaining. Few votes have been ‘nay’.

I suggest we interpret these examples as follows. When strongly quantificational determiners are felicitously used in the subject of a specificational sentence, they can be thought of as quantifying over ‘potential specifications’. To see what I mean by this, consider the analogy with the case of predication. In predicational sentences with a quantificational noun phrase in the subject, of the form

- (51) $QF \text{ is}_{\text{pred}} G$

the quantifier Q may be thought of as quantifying over potential predications; it ‘summarizes’ how many of the potential predications of F s as being G are true. For example, someone who says

- (52) Most e-mails are not worth reading.

is saying something true just in case most of the potential predications in which ‘ x is not worth reading’ is predicated of an email are true.

I am suggesting that we use the notion of a ‘potential specification’ to extend the concept of specification to sentences with quantificational determiners in the subject in an analogous way. To see why this is helpful, notice that the speaker in (50-a) could also have said:

- (53) The author of the first kind of letters was George, and the author of the second kind was George, too.

Here it is clear that both of the embedded copular clauses are specificational. Likewise, the speaker in (50-b) is describing a situation that could potentially be described like:

- (54) The vote by legislator 1 was ‘aye’, and the vote by legislator 2 was ‘aye’, ... the vote by legislator n was ‘nay’, ...

where few of the embedded clauses have ‘nay’ in the complement. Again, it is clear that these embedded clauses are specificational. The speaker of (50-a) or (50-b) is, in effect, ‘summarizing’ these alternative descriptions: she is saying that both ways of specifying the authors will specify them to be George, or that few ways of specifying the votes of individual legislators will specify them to be ‘nay’.

By extending the notion of specification in this way, we are able to make sense of the felt distinction between quantified specifications, like those in (50), and quantified predications, like these:

- (55) a. Both authors are famous.
b. Few votes are cast on Sundays.

The fact that the examples in (50) are possible indicates that we need to keep track of this distinction, even if quantified specifications are comparatively rare.

In the absence of other candidate counterexamples, I tentatively conclude that *all* descriptions are capable of occupying the subject position in a specificational sentence, at least if we admit ‘quantified specifications’ as a type of specificational sentence. If we don’t, more work is required to demarcate those descriptions which can be specificational subjects from those which cannot. Our observations here suggest that drawing that line will involve distinguishing quantificational and non-quantificational determiners, and perhaps certain kinds of indefinites from others. But on the basis of examples like those in (48) and (50), I will assume for now

that all descriptions can be specificational subjects, and there is no pressing need to make such a distinction.

2.2.2 *Anaphora*

Anaphoric pronouns form another important class of expression that can serve in the subject position of a specificational sentence. For example, (56-b) has a specificational reading, and the anaphoric pronoun 'it' as its subject:

- (56) a. How far away is Sacramento?
b. It is about 80 miles.

The list analogy confirms this reading; a colon would be appropriate after the copula in (56-b). But in this case, the directness test is more useful. It is clear that (56-b) directly answers the question asked in (56-a), in a way that, say,

- (57) It is too far to drive tonight.

does not. (Sentence (57) is predicational, rather than specificational.) Given the context provided by (56-a), (56-b) clearly means something similar to

- (58) The distance to Sacramento is about 80 miles.

which belongs to the same class as the specificational sentences in (32).

It may not seem obvious that the 'it' in (56-b) is an anaphor. There are two alternative hypotheses that initially appear plausible. One is that pronouns in the subject position of specificational sentences are simply semantically vacuous. This seems credible because 'About 80 miles' (without the initial 'it is') gives exactly the same answer to (56-a) as (56-b). In that case, 'it' would not denote anything at all in (56-b). The other hypothesis is that these pronouns are instead cataphors, which inherit their denotation from complement phrase, rather than the preceding question. In that case, 'it' in (56-b) would refer to the distance specified by 'about 80 miles'.

These alternative hypotheses are unlikely, however, given some facts about agreement.

- (59) Who is at the door?
 a. It is Samantha and Dave.
 b. *They are Samantha and Dave.
- (60) Which guests were the last two to leave?
 a. ??It was Samantha and Dave.
 b. They were Samantha and Dave.
- (61) Who is the lead in *North by Northwest*?
 a. It is Cary Grant.
 b. ?He is Cary Grant.

Each of the answers given to the questions in (59)–(61) is a specificational sentence with a pronoun in the pre-copular position. What’s important here is that this pronoun seems to be required to agree with the presuppositions of the preceding question, rather than with any features of the complement, which gives the answer. In (59), the question does not presuppose that the answer should be plural, and answering with the singular ‘it’ is grammatical, while answering with the plural ‘they’ is not. The opposite pattern occurs in (60), where the question *does* presuppose that the answer should be plural, even though the complement (‘Samantha and Dave’) is the same. In (61), ‘it’ is preferable to ‘he’, especially when the question is not interpreted as presupposing that the lead is male. If the pronoun was a cataphor referring to Cary Grant, ‘it’ would be ungrammatical, and ‘he’ would be required.

The fact that agreement is required at all, and that examples like (56-b) can be paraphrased using a meaningful definite description for a subject, makes it unlikely that the pronouns in the subject positions of these specificational sentences are completely vacuous. On the other hand, since they are required to agree with the features of the antecedent question, rather than with the features of the subsequent answer, these pronouns are properly regarded as anaphora rather than cataphora.

The range of anaphoric pronouns that may serve as the subject of a specificational sentence is somewhat unclear. In English, the neuter pronouns ‘it’ and ‘they’ clearly work as specificational subjects, as shown by examples like (59-a) and (60-b). But it is not clear whether the gendered pronouns ‘he’ and ‘she’ can do the same. The issue here is whether we should classify sentences like (61-b), where a gendered pronoun appears before the copula and a name or other non-predicate is the complement,

as having a specificational reading. Although such sentences are not predicational, they do not unequivocally pass the tests I have been using to bring out the specificational reading. In addition, such sentences have a good claim to be identity statements or equatives; so our judgment about these cases depends on how the class of specificational sentences relates to the class of equatives. I discuss this issue further below.

Anaphoric uses of ‘this’ and ‘that’ are also pretty clearly acceptable as the subject of a specificational clause:

- (62) There was only one thing left to do, and {this/that} was to call the doctor.

Here, either demonstrative seems acceptable to me (though ‘that’ is clearly preferable), and is most naturally read as anaphoric on ‘one thing left to do’. The list analogy reveals the second conjunct to be specificational. Moreover, either version of the sentence may be paraphrased as

- (63) The only thing left to do was to call the doctor.

and this sentence is a specificational sentence of the same type as those in (32).

Non-anaphoric uses of demonstratives, however, are another unclear case. In his original taxonomy, Higgins classified sentences like

- (64) That is Francis.

as ‘identificational’ when the demonstrative is used deictically. According to Higgins, identificational sentences are distinct from both specificational sentences and identity statements or equatives (Higgins, 1979, Ch. 5). Whether identificational sentences form a distinct class of copular sentence, or whether they should be subsumed into some other class, is debated; see the discussion in Mikkelsen (2011).

2.2.3 *Other types of subject*

There are a few other sorts of expression which may be able to serve as specificational subjects. Certain abstract nouns can appear as the subject of a specificational sentence:

- (65) a. Piety is being loved by all the gods.

- b. Validity is never having true premises and a false conclusion simultaneously.
- c. Racism is actions, practices or beliefs that consider different races to be inherently inferior to others.⁹

Similarly, infinitives and gerunds can also take the subject position:

- (66) a. To be wise is to know what is good.
- b. Having courage is knowing what to fear.

The intuitive tests render fairly clear verdicts here. The list analogy predicts that the examples in (65) and (66) all have a specificational reading. A colon would be appropriate in these sentences, as would adding further items to the list, as in

- (67) Piety is: being loved by all the gods, making timely sacrifices, and praying twice daily.
- (68) To be wise is: to know what is good, and to act in accordance with that knowledge.

Likewise, each could serve as a direct answer to a corresponding ‘what is...?’ question, and it is reasonably clear that on this interpretation, the questions do not permit predicational sentences as answers, though I will not try to argue that here.

As with sentences with gendered pronouns for subjects, however, sentences like those in (65) and (66) have a good claim to be equatives, or ‘generalized identities’. So a more considered judgment about these sentences must be postponed until after the discussion of equatives in Chapter 3.

Still, these examples raise an interesting point. Neither abstract nouns nor infinitives and gerunds count as descriptions in the sense used above, since they do not have the form of a noun phrase headed by a separate determiner. The feature that these expressions share with questions and many descriptions, however, is that they are *nominalizations*. They are expressions which can syntactically occupy the same kinds of positions as proper names, but which are derived from expressions in other syntactic categories, such as adjectives or verbs. (One such position, of course, is

⁹A longer version of this sentence began the Wikipedia article on racism (accessed at <https://en.wikipedia.org/wiki/Racism> on May 16, 2014).

the subject position of a copular sentence.) In some cases, English has a special form for the nominalization of another expression, as ‘justice’ is the nominalized form of the adjective ‘just’. But we also have productive devices for nominalizing: verbs are nominalized by forming their infinitives or gerunds; adjectives and concrete nouns are commonly nominalized with suffices like -ity, -hood, -ness, and -ism.

If we take the category of nominalizations to encompass all the examples of specificational subjects we have looked at so far, we can ask: can expressions which are *not* nominalizations serve as specificational subjects?

The answer appears to be negative. I have been unable to find examples of well-formed copular sentences which clearly have a specificational reading, but which have anything other than nominalizations in the subject position, such as adjectives, adverbs, or conjugated verbs:

- (69) a. Blue is ...
 b. Dreary is ...
 c. Older than me is ...
- (70) a. Quietly is ...
 b. At noon in Paris is ...
 c. Forevermore is ...
- (71) a. *Run is ...
 b. *Will consider is ...
 c. *Had felt worse is ...

There seems to be no way of supplying a complement in these examples that will yield a specificational sentence.¹⁰ In some cases, a well-formed copular sentence is possible, but it will only have a reading which is not

¹⁰In conversation, John MacFarlane has suggested that examples like “Quietly is without making a sound” or “Biweekly is every two weeks” are specificational sentences with adverbial subjects. These examples do not seem clearly grammatical to me; but insofar as they are, I read them as akin to, or elliptical for, sentences like “‘Quietly’ means *without making a sound*”, in which the adverb appears quoted. Since quoting an adverb forms a nominal expression, it seems to me that reading such sentences specifically still somehow requires nominalizing the adverb. Thus, if they do indeed have a specificational reading, they won’t be counterexamples to my claim that specificational subjects are nominalizations. For a fascinating meditation on some further themes in this general direction, see Sellars (1985).

clearly specificational, such as the predicational reading, or a ‘pure’ equative reading:

- (72) Blue is a color.
 (73) Honest is honest.

Interestingly, proper names also do not seem to permit specification. There are no clearly specificational sentences of the form:

- (74) a. Ralph and Amelia are ...
 b. Cicero is ...

though there are predicational and equative sentences of this form:

- (75) a. Ralph and Amelia are at the dance.
 b. Cicero is Tully.

This observation permits me to summarize the conclusion of this section: *nominalizations can serve as specificational subjects; other expressions cannot*. I have not, of course, offered a theory of what ‘nominalizations’ are. Still, this choice of terminology is appropriate, because our intuitive concept of nominalization lines up fairly neatly with the observations we have made about which expressions can serve as specificational subjects. Nominalizations include questions in the pre-copular position of pseudo-clefts; all descriptions; and at least some anaphora, abstract nouns, infinitives, and gerunds. They exclude adjectives, adverbs, conjugated verbs, proper names, and a host of expressions in other syntactic categories.

Thus, although there are a great variety of nominalizations, it seems like there is still an interesting distinction between expressions which can, and expressions which cannot, serve as the subject of a specificational sentence. By contrast, I have found no interesting distinctions between the kinds of expression which can and cannot serve as the *complement*. The examples above reveal a great variety of expressions in this position. There are isolated examples of expressions which cannot work there, such as bare articles like ‘a’, or particles like ‘of’. But I have no suggestions for how to categorize specificational complements more systematically. In a sense, this is not surprising, given the intuitive tests: specificational complements encompass anything we might care to make

a list of, or answer a question with; and that is a very wide category indeed.

This concludes my survey of the variety of specificational sentences. I turn now to developing and defending a semantic analysis which makes sense of these data.

Chapter 3

Semantic analyses of specification

We are now in a position to evaluate some existing semantic analyses of specificational sentences. There are three prominent proposals in the literature about the semantics of specificational sentences: the inversion analysis, the equative analysis, and the question-answer analysis. I will assess each of these analyses in this chapter on intuitive, empirical, and theoretical grounds, arguing that the question-answer analysis provides the best prospects for a philosophical understanding of specification.

To assess each analysis, I will rely on criteria that have emerged from the survey of specificational sentences in Chapter 2. The criterion I consider most important is how well each analysis captures the pre-theoretical characterization of specification discussed in Section 2.1: namely, that the subject of a specificational sentence *expresses a constraint*, while its complement *says what satisfies* that constraint. This characterization of specification is brought out by the list analogy and the directness test. It tells us what unifies specificational sentences at a semantic level, despite the significant syntactic variation in both their subject and complement positions. And it distinguishes the specificational interpretation of a sentence from other readings it may have, such as a predicational or issue-predicational reading. Any semantic analysis of specification should thus try to represent this characterization in theoretical terms.

To do so, a successful analysis must capture two different ideas. First of all, it should capture the asymmetry in meaning between the subject and complement of a specificational sentence. This is an asymmetry of semantic function, or *role*: it is the difference between ‘expressing’ a constraint and ‘saying what satisfies’ it. On the other hand, a successful analysis should also capture a certain symmetry in the meaning of the

subject and the complement, namely, the fact that they are concerned with the *same* constraint. The subject and complement are in this sense ‘about the same thing’; their meaning is connected by sharing the same *subject matter*, though they play different roles with respect to that subject matter.

The idea that specificational sentences involve both an asymmetry of role and a symmetry of subject matter is recognized throughout the literature, though it has not been expressed in consistent terminology. Akmajian, for example, says that the subject of a specificational sentence “contains what is essentially a semantic variable, a semantic ‘gap’ that must be ‘filled’ or specified by the [complement]” (Akmajian, 1970, p. 19). Higgins (1979), Mikkelsen (2005, 2011) echo this terminology of ‘introducing a variable’ and ‘specifying a value’. Higgins also says that a specificational sentence “functions rather like a list”, where the subject “constitutes the heading” of the list, and the complement is an item on it (Higgins, 1979, p. 8). And Heycock and Kroch say that the subject “has the same denotation as” the complement, but differs in “information packaging”, with the subject being less informative than the complement (Heycock & Kroch, 1999, pp. 388,394). Each of these ways of speaking posits something shared by the subject and complement (a variable, a gap, a list, a denotation), as well as something that differs between them.

I shall continue to say that the subject and complement share a subject matter, but differ in semantic role. Compared to the other choices, this terminology seems less theoretically-loaded as well as more general. But whatever terminology we use, the challenge for any analysis of specification is to represent both what is common, and what differs, between the meanings of the subject and the complement. What makes this difficult is the wide variety of expressions found in specificational sentences. A successful analysis must work for all the syntactic and semantic types we have found in both the subject and the complement positions.

Each of the three prominent analyses goes some way toward capturing the intuitive characterization, but also faces some problems. I will argue that the inversion analysis neatly captures the asymmetry of role, but fails to capture the symmetry of subject matter, in specificational sentences. The equative analysis, on the other hand, easily captures the symmetry of subject matter, but fails to capture the asymmetry of role. These two analyses thus have complementary virtues and vices; each captures an aspect of the intuitive characterization that the other does not.

This argument points toward the question-answer analysis as a perspicuous way of unifying the virtues of the other two into a single account. Compared to either inversion or equation, questions and answers provide a clearer model of how two expressions can share the same subject matter while differing in semantic role.

The question-answer analysis has recently been challenged on empirical grounds, however. I will argue that it survives this challenge, so it is the best contender for a satisfactory view.

To be clear, I am not claiming that the inversion and equative analyses are empirically inadequate, or that they cannot be made adequate. They probably can. Rather, I am claiming that evolving either of them toward a satisfactory account will involve making it more like a version of the question-answer analysis, because only the question-answer analysis currently has the resources to capture both the symmetry of subject matter and asymmetry of role in specificational sentences.

3.1 THE INVERSION ANALYSIS

The inversion analysis proposes that specificational sentences are *inverted predications*. They are predicational copular sentences in which the predicative expression occupies the grammatical subject position, rather than the complement position. This idea has mostly been pursued as a syntactic hypothesis, in literature that starts with Williams (1983).¹ Here, I focus on a semantic proposal which supports and complements this syntactic hypothesis, defended by Mikkelsen (2005). According to this proposal, the subject of a specificational sentence functions *semantically* as a predicate. That is, in a specificational sentence, the subject purports to describe, or be true of, the complement.

Two important ideas motivate the inversion analysis. The first is that we can understand the semantic roles of the expressions in the subject and complement positions of a specificational sentence by examining the roles they play in non-specificational copular sentences. For example, sentence (1-a) is specificational, like the examples in (32). But it is reasonable to expect that ‘the lead in *North by Northwest*’ and ‘Cary Grant’ make the same contributions in (1-a) as they make in (1-b), which is not speci-

¹See citations in Mikkelsen (2005, Ch. 1) and Heycock and Kroch (1999) for other work in this line.

fictional. The inversion analysis thus hopes to understand the semantic roles of these expressions in (1-a) in terms of the roles they play in (1-b).

- (1) a. The lead in *North by Northwest* is Cary Grant. (specificational)
 b. Cary Grant is the lead in *North by Northwest*. (non-spec.)

The second motivating idea is that definite descriptions in the complement position of copular sentences can function semantically as predicates. This idea dates back at least to Strawson (1950), and is formalized by Partee (1987). In type-theoretic terms, the idea is that the type of the subject in a sentence like (1-b) is e , while the type of the complement is $\langle e, t \rangle$. Most significantly, this means that the complement is not a singular term, and that the sentence need not be analyzed as an identity statement.²

By putting this idea together with the first, we arrive at the following proposal. In a specificational sentence like (1-a), where a definite description appears as the subject and a proper name occurs as the complement, these expressions retain the types they have in predicational sentences like (1-b), despite appearing on the other side of the copula. The definite description in this sentence has type $\langle e, t \rangle$, while the proper name has type e . That is, in a specificational sentence, the subject functions semantically as a predicate of the complement—an inversion of the relationship observed in predicational sentences. This proposal is summarized in Table 3.1.

Table 3.1: The inversion analysis of specification

	Subject	Complement
Predication	e	$\langle e, t \rangle$
Specification	$\langle e, t \rangle$	e

²A variety of empirical tests support the claim that definite descriptions can function as predicates in the complement position of a copular sentence. According to these tests, definite descriptions in this position pattern with clear cases of predicates, such as adjectives, rather than with paradigmatically e -type expressions, such as proper names. For example, post-copular definite descriptions can be coordinated with adjectives and other predicates, while names cannot: consider “He is tall, handsome, and {the lead actor/*Cary Grant}.” For a more extensive discussion of these tests for predicativity, see Rieppel (2013b).

3.1.1 *Support for the inversion analysis: pronominalization of specificational subjects*

How does the inversion analysis make it plausible that a definite description in the subject of a specificational sentence has type $\langle e, t \rangle$, rather than type e ? Mikkelsen (2005) argues directly for this claim using pronominalization data. She observes that there is a systematic difference in the way definite descriptions pronominalize in specificational and predicational clauses. In English, the difference appears with descriptions that can refer to humans. When such a description appears as the subject of a predicational clause, it must pronominalize with a gendered pronoun, such as ‘he’ or ‘she’³:

- (2) a. The tallest girl in the class is Swedish, isn’t {she/*it/*that}?
 b. As for the tallest girl in the class, {she/*it/*that} is Swedish.
 c. What nationality is the tallest girl in the class?
 {She/*It/*That} is Swedish.

But when the same description appears as the subject of a specificational clause, it can pronominalize with neutral anaphors, such as ‘it’ and ‘that’:

- (3) a. The tallest girl in the class is Molly, isn’t {it/?she}?
 b. As for the tallest girl in the class, {it/that/?she} is Molly.
 c. Who is the tallest girl in the class?
 {It/That/?She} is Molly.

Mikkelsen argues that this contrast in the way descriptions pronominalize in these positions points to a difference in the semantic type of the subject expression in specificational and predicational clauses. She points out that ‘it’ and ‘that’ cannot be used to pronominalize e -type expressions which refer to humans, such as proper names:

- (4) a. Cary Grant is tall, isn’t {he/*it/*that}?
 b. As for Cary Grant, {he/*it/*that} is tall.

The pronominalization facts for predicational clauses in (2) parallel those for proper names in (4). Thus, they are consistent with treating descriptions in the subject of a predicational clause as having type e . But this

³The examples in (2) and (3) appear in Mikkelsen (2005, Ch. 5).

implies that such descriptions must *not* have type e when they appear as the subject of a specificational clause. If they did, they would obligatorily pronominalize with gendered pronouns, since neutral anaphors cannot be used to pronominalize expressions which refer to humans. In fact, however, a gendered pronoun is not obligatory; such descriptions can pronominalize with neutral anaphors, as the examples in (3) show.

On the other hand, neutral anaphors can be used to pronominalize $\langle e, t \rangle$ -type expressions:

- (5) Cary Grant is tall. {It/That} is a good thing to be in Hollywood, if you want to work as a leading man.

Thus, the pronominalization facts illustrated in (3) are consistent with the hypothesis that the subject of a specificational sentence is of type $\langle e, t \rangle$, but not the hypothesis that it is of type e , as the inversion analysis predicts.⁴

3.1.2 *Assessing the inversion analysis*

So the inversion analysis appears to be motivated by independently-plausible ideas, as well as supported by some empirical data. How well does it capture the characterization of specification which emerged from our earlier discussion? Can the inversion analysis explain the sense in which specificational sentences exhibit both a symmetry of subject matter and an asymmetry of semantic role?

The inversion analysis does a fair job of representing the asymmetry of semantic role between the subject and complement of a specificational sentence. The contrast between the roles of ‘expressing a constraint’ and ‘saying what satisfies it’ is encoded in the theory as the contrast between

⁴As noted, we can only contrast how the subjects of specificational and predicational sentences pronominalize using English descriptions that can refer to a human. This is because English only uses gendered pronouns to pronominalize human-denoting expressions. But Mikkelsen also shows that in Danish, which has grammatical gender, the contrast extends to descriptions which refer to inanimate objects. In Danish, such descriptions also pronominalize like $\langle e, t \rangle$ -type expressions rather than like e -type expressions when they appear as the subject of a specificational clause. This indicates that the contrast is robust and not a mere fluke of how English pronominalizes human-denoting descriptions. For more detailed discussion of the data, see Mikkelsen, 2005, § 5.3.

$\langle e, t \rangle$ -type subjects and e -type complements. This is a natural way of representing the asymmetry. In general, a semantic type system is meant to encode differences of semantic role, such as the difference between referential, predicative, and quantificational expressions. Moreover, type $\langle e, t \rangle$ is a natural choice for a type representing ‘constraints’ on values denoted by expressions of type e . An object satisfies a constraint when the constraint is *true of* it, and does not satisfy it otherwise. This idea has a fruitful history in logic and computer science.

Still, this representation of the asymmetry looks somewhat limited in scope, given the variety of specificational sentences that we surveyed in Section 2.2. First of all, it’s not clear how well it extends to specificational sentences where the complement is not of type e , like

- (6) a. The claim he made was that he was innocent.
 b. The top priority is getting re-elected.

The natural way to deal with such examples is to generalize the inversion analysis along the lines in Table 3.2: for any type α , in a specificational sentence whose complement has type α , the subject has type $\langle \alpha, t \rangle$. This extends the idea that the subject expresses a constraint, while the complement says what satisfies it, to higher-type complements.

Table 3.2: An inversion analysis for higher-type complements

	Subject	Complement
Specification	$\langle \alpha, t \rangle$	α

This proposal has a technical cost. It requires at least as many types for definite descriptions as there are types for complements in specificational sentences. For example, consider a version of (4) that has a definite description for a subject:

- (7) The thing I bought for mother was some flowers.

On the standard analysis, the complement ‘some flowers’ is a generalized quantifier, with type $\langle \langle e, t \rangle, t \rangle$. Thus, according to Table 3.2, the definite description in the subject needs to have type $\langle \langle \langle e, t \rangle, t \rangle, t \rangle$. Though definite descriptions are generally recognized to take on at least three

different types, this is not among them.⁵ So this proposal implies that higher-type definite descriptions would be needed in specificational sentences. Admitting such higher-type descriptions would likely require constraining their availability elsewhere in semantic theory, and it is not clear how easy this would be to accomplish.

Another issue with generalizing the inversion analysis this way arises for specificational sentences with $\langle e, t \rangle$ -type complements. According to the generalized inversion analysis, such sentences would have a subject of type $\langle \langle e, t \rangle, t \rangle$. This assignment of types, however, collides with the standard analysis of sentences containing generalized quantifiers in subject position.

- (8) a. Most students are smart.
b. What Amelia is is smart.

If we accept both the standard analysis of generalized quantifiers and the inversion analysis, then both of the sentences in (8) have a subject of type $\langle \langle e, t \rangle, t \rangle$ and complement of type $\langle e, t \rangle$. Yet (8-b) is specificational, while (8-a) is not. The inversion analysis thus fails to distinguish specificational and non-specificational sentences in this case.

Even if we leave these technical issues aside, the inversion analysis faces some empirical problems. Perhaps the most pressing problem is that inverting a predicational sentence does not always yield a specificational sentence. Consider an adjective like 'handsome', as used in a sentence like (9-a). This sentence is an ordinary predicational sentence, with a subject of type e and complement of type $\langle e, t \rangle$. But inverting this sentence around the copula, as in (9-b), does not yield a specificational sentence.

- (9) a. Cary Grant is handsome.
b. *Handsome is Cary Grant.

Note that the problem here is not simply that adjectives are syntactically barred from the subject position of copular sentences:

- (10) a. Handsome is what I want to be.
b. Honest is honest.

⁵Following Partee (1987), the standard types are e , $\langle e, t \rangle$, and $\langle \langle e, t \rangle, t \rangle$.

- c. Happy is the man that findeth wisdom.⁶

The problem with (9-b) instead seems to be that ‘Cary Grant’ cannot be the complement of a sentence which has such an adjective as its subject. Since proper names can be specificational complements in general, the problem must be that adjectives like ‘handsome’ cannot be specificational subjects, at least not when paired with a complement that is a proper name.⁷

This point is in line with our earlier observation that adjectives, conjugated verb phrases, and other non-nominalizations cannot serve as the subject of a specificational sentence. Under the inversion analysis, it is not clear why this should be so, since many such expressions are standardly analyzed as predicates of individuals. Thus, even if we grant the inversion analysis’ proposal that specificational sentences are inverted predications, that cannot be the whole story. There’s something special about nominalizations that allows them to serve as specificational subjects, while adjectives and other $\langle e, t \rangle$ -type expressions cannot.⁸

The inversion analysis also has trouble accounting for the fact that specificational sentences support some substitution inferences. Consider the inference in (11):

- (11) Ms. Pruett praised the tallest girl in the class.
 The tallest girl in the class is Molly.
 So, Ms. Pruett praised Molly.

This looks like a good substitution inference, with a specificational sentence as its second premise. The problem is that the inversion analysis can’t on its own explain *why* this inference is good, rather than fallacious: according to the inversion analysis, ‘the tallest girl in the class’ as it appears in the second premise is a predicate, not a singular term. Thus, according to the inversion analysis, the inference has a logical form like:

⁶Proverbs 3:13, King James Bible.

⁷Similar observations are made by Heycock and Kroch (1999).

⁸Mikkelsen herself recognizes a parallel gap here; she says “a necessary, but not sufficient, condition for a DP to be the subject of a specificational clause is that the DP can...occur in type $\langle e, t \rangle$ ” (Mikkelsen, 2005, p. 108). But so far as I can see, she does not offer an explanation of why some $\langle e, t \rangle$ -type expressions cannot be specificational subjects, whether they are DPs or not.

- (12) $S(p, t)$
 $T(m)$
 So, $S(p, m)$

which is clearly invalid as it stands. To validate the inference, the inversion analysis needs some other principle that connects the denotation of ‘the tallest girl in the class’ when it occurs as a predicate in the second premise with its occurrence as a singular term in the first premise.⁹

These last two problems are symptoms of a more general one. Although the inversion analysis captures the asymmetry of role in the subject and complement of a specificational sentence, it has little to say about their symmetry of subject matter. From my perspective here, this is the most serious problem with the inversion analysis.

There is a clear and intuitive sense in which a name and a definite description can be ‘about the same thing’. This fact explains why the second premise in (11) supports the substitution inference: substituting ‘Molly’ for ‘the tallest girl in the class’ is legitimate because the specificational premise guarantees that these expressions are about the same thing. On the other hand, it is equally clear and intuitive that a name and an adjective cannot be ‘about the same thing’ in this same sense. This seems like the reason why (9-b) does *not* work as a specificational sentence, despite the fact that adjectives can be subjects, and proper names can be complements, of copular clauses in general. The fact that the inversion analysis has trouble accounting for these examples suggests that it does not sufficiently distinguish when a subject and complement expression can be ‘about the same thing’, and when they cannot. That is, it does not adequately represent the symmetry of subject matter in specificational sentences.

Actually, we have already seen a similar problem in example (23) of Chapter 1, repeated here as (13). We observed there that ‘number’ and ‘way’ are not interchangeable:

- (13) a. #Four is the way of moons Jupiter has.

⁹In this case, an appropriate type-shifting principle can probably do the job. But Mikkelsen (2005) proposes no such principle, so it is hard to tell whether the inversion analysis can use type-shifting to account for such inferences in general, including inferences where the specificational premise has a higher-type complement. I discuss one such case below, in connection with the equative analysis; see example (31).

- b. #Quietly is the number Mary entered.

Consider a version of (13-b), inverted to form a clear case of an (attempted) specificational sentence. The result is still semantically unacceptable:

- (14) #The number Mary entered was quietly.

By contrast,

- (15) The way Mary entered was quietly.

is a perfectly fine specificational sentence. But there is no reason to suspect that ‘the *number* Mary entered’ is a defective specificational subject, while ‘the *way* Mary entered’ is not. Rather, in the language of this chapter, the problem in (14) is that ‘quietly’ cannot say what satisfies any constraint that ‘the number Mary entered’ could express. These two expressions cannot share a subject matter.

Again, I am not claiming that the inversion analysis cannot respond to these difficulties. But I think they do show that the analysis needs to be extended, and they indicate that the extension should be toward representing the symmetry of subject matter in the subject and complement of specificational sentences. This is a point on which the equative analysis fares better, so let’s consider that next.

3.2 THE EQUATIVE ANALYSIS

The original proponent of an equative analysis was Frege. Frege, of course, held that sentences like

- (16) The number of moons of Jupiter is four.

are identity statements. In such sentences, the copula expresses the identity relation, and both the subject and the complement are singular terms.

As we have seen, there are some reasons to doubt Frege’s analysis. In general, specificational sentences can have complements which are not singular terms. Once we recognize that (16) is a specificational sentence, and that ‘four’ is not a singular term in many of the other contexts where it appears, it’s natural to suppose that ‘four’ does not occur as a singular term in (16) either. This sentence is a specificational sentence, but it may

not be an identity statement, because it may have a higher-type complement. At any rate, it is obvious that Frege’s analysis of (16) will not extend to clear cases of specificational sentences with higher-type complements: he takes the identity relation to be distinctive of objects, and the sign of identity to only appropriately occur between singular terms.¹⁰

There are two possible responses to this problem. One response holds that specifications with higher-type complements undermine Frege’s original thought, and show that a completely different analysis of specification is needed. Another response tries to hang on to Frege’s idea that specifications somehow ‘identify’ their subject and complement, by generalizing it to account for specifications with higher-type complements. The equative analysis is a version of the second response.

For a proponent of the equative analysis, specificational sentences are one species of equatives. Recall from Section 2.1.3 that equatives are copular sentences that state ‘generalized identities’. So according to the equative analysis, a specificational sentence *equates* the denotations of its subject and its complement, even when neither is a singular term. In type-theoretic terms, this implies that for any type α , when the subject of a specificational sentence has type α , the complement must have type α too. This proposal is summarized in Table 3.3.

Table 3.3: The equative analysis

	Subject	Complement
Specification	α	α

Recent versions of the equative analysis appear in Heycock and Kroch (1999) and Heller (2005). I will focus on the version defended by Heycock and Kroch, because their proposal specifically addresses specificational sentences with higher-type complements. Heller’s proposal is also quite interesting, and I will have more to say about it later. But because she focuses on specificational sentences with *e*-type complements, her proposal suffers from the same lack of generality as Frege’s.

¹⁰Frege says, for example: “the relation of equality (*Gleichheit*) by which I understand complete coincidence, identity, can only be thought of as holding for objects, not concepts” (Frege, 1891/1997a, p. 175).

3.2.1 *Support for the equative analysis: pure equatives and connectivity effects*

Heycock and Kroch motivate their version of the equative analysis by observing that English contains ‘true equatives’ like those in the following examples.¹¹

- (17) a. Your attitude toward Jones is my attitude toward Davies.
 b. Your opinion of Edinburgh is my opinion of Philadelphia.
- (18) a. When it comes down to it, honest is honest.
 b. In the end, long is long.
 c. You can dress is up if you like, but in the end being dishonest is just being dishonest.

The sentences in (17) help show that equatives are not plausibly interpreted as a species of predicational sentences, even when the subject and complement are different. In these examples, neither expression is ‘more referential’ or ‘more predicative’ than the other, and neither is predicated of the other. The same is true of the tautologous (embedded) equatives in (18). These examples show that there are clear cases of equatives with higher-type expressions, such as adjectives, on both sides of the copula. Heycock and Kroch conclude from these examples that we need a category of equatives which includes sentences with higher-type complements. So the only question is whether specificational sentences are among them.

To show that specificational sentences are indeed equatives, they appeal to the following examples (Heycock & Kroch, 1999, pp. 379–380):

- (19) John is the one thing I have always wanted a man to be (that is, he’s honest).
- (20) a. The one thing I have always wanted a man to be is honest.
 b. *The one thing I have always wanted a man to be is John.

Sentence (19) shows that ‘the one thing I have always wanted a man to be’ is most naturally read as predicative, or type $\langle e, t \rangle$. But when this

¹¹These examples are from Heycock and Kroch (1999, pp. 373–374). See also examples (25) and (10) in Chapter 2, which are plausible cases of equatives.

predicative phrase is used as the subject of a specificational sentence, it only allows other predicative expressions as complements. This is shown by the contrast between sentences (20-a) and (20-b): an $\langle e, t \rangle$ -type complement like ‘honest’ is grammatical, but an e -type complement like ‘John’ is not. The pattern in (20) therefore suggests that the subject and complement of a specificational sentence must have the same type, as the equative analysis predicts.¹²

Once we take the subject and complement to have the same semantic type, an equative analysis of specificational sentences seems like the only viable option. For example, a sentence like (20-a) is naturally read as saying that ‘the one thing I have always wanted a man to be’ and ‘honest’ denote the same property. That is, it seems best interpreted as equating the denotations of its subject and complement. Other possible relations between same-type denotations (such as non-identity or a generalized quantifier) do not seem to fit the meaning of the sentence.

Here are the details of how Heycock and Kroch implement their equative analysis. They focus on specificational pseudo-clefts, and assume that these sentences have an equative syntax. They also assume that the subject of such a sentence is a free relative, and claim that this free relative has the form of a definite description $\iota y[f(y)]$, where the type of the variable y “ranges over all the semantic types that free relatives can denote” (Heycock & Kroch, 1999, p. 383). They define the ι operator as follows (Heycock & Kroch, 1999, p. 383):

(21) $\iota y[f(y)]$ denotes a iff $f(a)$ (and a uniquely satisfies f)

I am abstracting here from Heycock and Kroch’s initial formulation of the uniqueness condition, which they eventually remove in favor of thinking of uniqueness as a pragmatic presupposition.

Thus, according to their analysis, in a specificational pseudo-cleft like

(22) What Mary did was run the marathon.

the subject is a free relative with the form

¹²Note that the inversion analysis predicts the opposite pattern in this case. If specificational sentences are inverted predications, then (20-b) should be grammatical rather than (20-a).

(23) $\iota y[\text{Mary did } y]$

y here has type $\langle e, t \rangle$; it is a predicate of Mary. ‘Did’ constrains y to be an action in the past, which I will simply represent as ‘ y -ed’. So writing (23) in the form need to apply the definition of ι , we have

(24) $\iota y[\lambda x. [\text{Mary } x\text{-ed}](y)]$

The whole sentence then has a form that is more perspicuously represented as:

(25) $\iota y[\lambda x. [\text{Mary } x\text{-ed}](y)] = \text{run the marathon}$

where ‘run the marathon’ has type $\langle e, t \rangle$, and ‘=’ is a (polymorphic) sign of equation.¹³

An equative sentence like (25) is true just in case the two expressions in the equation have the same denotation. We can thus compute its truth conditions by computing the denotation conditions for the subject, given a denotation for the complement. Suppose that the complement ‘run the marathon’ denotes the property of running the marathon. Call this property rm for short, keeping in mind that since ‘run the marathon’ has type $\langle e, t \rangle$, rm is a function from individuals to truth values. So (25) is true just in case the subject phrase *also* denotes rm .

We can now apply the definition of the ι operator to reduce this as follows:

(26) the subject denotes rm , iff
 $\lambda x. [\text{Mary } x\text{-ed}](rm)$ (and rm is unique), iff (by def. of ι)
 Mary rm -ed (and rm is unique) (by β -reduction)

So the sentence is true just in case Mary performed the action of running the marathon in the past—that is, if she ran the marathon. It presupposes that running the marathon is *all* she did in the relevant context.

Notice that by applying the definition of ι in this way, we effectively transform the pseudo-cleft (22) into a simpler sentence, one that would

¹³Heycock and Kroch treat equative syntax as a type of small clause, and argue that it is independent of the copula (Heycock & Kroch, 1999, §3.4). On their analysis, we should not think of the ‘=’ sign in (25) as expressed by the copula in (22); rather, the equation is implicit in the clause type, and the copula is semantically vacuous. Again, this detail is not important for my purposes here.

read like:

- (27) Mary ran the marathon.

Heycock and Kroch call this transformation *ι*-reduction, and they go on to show that, by treating *ι*-reduction as an obligatory operation for interpreting a specificational pseudo-cleft, they can explain why these sentences exhibit some syntactic features known as *connectivity effects*. I won't go into too much detail here, because explaining connectivity is mainly a problem for the theory of syntax and outside the scope of my discussion. But I will examine it briefly, since it provides important empirical support for Heycock and Kroch's view.¹⁴

Here's the idea. Some specificational pseudo-clefts exhibit a puzzling syntactic property:

- (28) a. What Vinny didn't buy was any wine.
b. What Narcissus likes is staring at himself.

These sentences each contain an element in the complement that normally must occur in a configuration with some other expression. In (28-a), this element is the negative polarity item 'any', which normally must occur in the scope of a negation (or other downward-entailing context). In (28-b) it is the reflexive pronoun 'himself', which normally must occur in the scope of a human-denoting expression. But in specificational pseudo-clefts, these expressions are licensed in the complement even though standard syntactic analyses say that those configurations do not obtain. That is the connectivity effect: an NPI or reflexive pronoun is grammatical in the complement, even though it is predicted not to be.

Heycock and Kroch explain this by observing that *ι*-reduction transforms a pseudo-cleft with a connectivity effect into an equivalent sentence where the normally-required configurations are present.

- (29) a. Vinny didn't buy any wine.

¹⁴I should point out that, for reasons that are too complicated to review here, Heycock and Kroch think of *ι*-reduction as a *syntactic* transformation. But their discussion makes clear that they also think of *ι*-reduction as part of the *interpretation* of a sentence; calling it 'syntactic' is, for them, compatible with thinking of *ι*-reduction as a step in a *semantic* derivation of the truth conditions of a specificational pseudo-cleft. That is how I shall think of it here.

- b. Narcissus likes staring at himself.

ι -reducing (28-a) produces (29-a), where the NPI ‘any wine’ occurs in the scope of the negation in ‘didn’t’. Likewise, ι -reducing (28-b) produces (29-b), where ‘himself’ is c-commanded by ‘Narcissus’. Basically, Heycock and Kroch’s idea is that ‘any’ and ‘himself’ are licensed in the pseudo-clefts in (28) because they are licensed in these reduced forms, which are obligatorily produced during interpretation of the pseudo-clefts. Thus, the equative analysis, when paired with obligatory ι -reduction, accounts for an important class of connectivity effects.

3.2.2 *Assessing the equative analysis*

So the equative analysis is supported by empirical data, and has some important explanatory advantages. I think it also has an undeniable pre-theoretical appeal: specificational pseudo-clefts in particular ‘feel like’ equations of some kind. When I say

- (30) What I bought for mother was some flowers.

what I say seems true only if I am using ‘what I bought for mother’ and ‘some flowers’ as different descriptions of the very same thing.

In my terminology, the right way to articulate this intuition is that the subject and complement of a specificational sentence exhibit symmetry of subject matter. The equative analysis represents this in a natural way: for the two parts of a specificational sentence to have the same subject matter is for them to have the same *denotation*. The theoretical concept of denotation plays many roles, but one role it often takes on is providing a precise counterpart to the idea that linguistic expressions are ‘about’ things in the world. If two expressions have symmetry of subject matter when they are ‘about the same thing’, then equality of denotation is a good candidate for representing symmetry of subject matter. So the equative analysis captures one part of our characterization of specification quite well.

One consequence of this is that the equative analysis does not face some of the worries that face the inversion analysis. For example, the equative analysis has a straightforward story about why (9-b) is ungrammatical:

(9-b) *Handsome is Cary Grant.

The problem here is that ‘handsome’ has type $\langle e, t \rangle$, and ‘Cary Grant’ has type e , and neither of these expressions can take on other types. Because of the type mismatch, this sentence cannot be an equative sentence. Thus, it cannot be a specificational sentence, even though both the subject and complement can appear in other copular sentences at the same positions.

Likewise, the equative analysis has no trouble accounting for the inference in (11):

- (11) Ms. Pruett praised the tallest girl in the class.
 The tallest girl in the class is Molly.
 So, Ms. Pruett praised Molly.

According to the equative analysis, ‘the tallest girl in the class’ has type e in the second premise, so this premise is just an identity statement. Thus, the equative analysis can validate this inference just by appealing to the logic of identity.

With higher-type complements, the picture is more complex, because it is less clear that higher-type expressions with the same denotation are intersubstitutable *salva veritate*, even in extensional contexts.¹⁵ Still, it appears that some higher-type substitutions make good inferences:

- (31) Oscar came to campus the (same) way Ernie came to campus.
 The way Ernie came to campus was via Telegraph.
 So, Oscar came to campus via Telegraph.

Here, the second premise is a specificational sentence with a higher-type complement. Substituting this complement for the occurrence of the subject in the first premise, in a manner analogous to the inference in (11), yields a conclusion that seems to follow. If this is indeed a good inference, the equative analysis could validate it just by appealing to a higher-order analogue of the laws of identity. This seems like a virtue, although I don’t want to pursue this here.

¹⁵One reason this is not clear is that higher-type expressions of the same *semantic* type often do not have the same *syntactic* type, and substituting an expression of one syntactic type for another usually yields nonsense, even if they have the same denotation.

Still, the equative analysis faces both technical and empirical problems. On the technical side, Heycock and Kroch's analysis faces problems similar to those facing the inversion analysis. First of all, the analysis needs to be extended to specificational sentences other than pseudo-clefts before it can account for all the types of specificational subject that we observed in Section 2.2. Officially, Heycock and Kroch's assumption that specificational subjects are descriptions of the form $\iota y[f(y)]$ only applies to free relatives in the subject position of a pseudo-cleft, though it is straightforward to extend this assumption to overt definite descriptions as well. It is less obvious how to extend their analysis to specificational sentences with other types of definites or with indefinites as subjects, like those in (44) and (45). But presumably, these cases could be handled by defining analogues of the ι operator to represent the different determiners in these kinds of subject.

Like the inversion analysis, the equative analysis also may require more types of definite description than are standardly recognized. The problem is less aggravated in this case, because the equative analysis does not hold that the subject has a *higher* type than the complement. But there are still some potentially awkward examples. Non-intersective modifiers in the complement provide one kind of case:

- (32) (Some mothers are harsh, some are overbearing.) But the kind of mother mine is is good.

On the specificational reading, this sentence says what kind of mother the speaker has; it says he has a *good* mother. Since 'good' is a non-intersective modifier, the best analysis of 'good' in this reading of the sentence might assign it a type like $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$. The equative analysis would then assign this non-standard type to the definite description in the subject position as well. Again, the problem is that admitting such higher types of definite descriptions likely requires constraining them in other contexts, and it's not clear how easily this can be done.

On the empirical side, the most pressing problem for the equative analysis is that it has no obvious explanation of the pronominalization data which supported the inversion analysis. Recall that human-denoting descriptions may pronominalize with neutral anaphors in English when they occur as the subject of a specificational sentence (but generally not elsewhere), as we saw in (3):

- (3) a. The tallest girl in the class is Molly, isn't {it/?she}?
 b. As for the tallest girl in the class, {it/that/?she} is Molly.
 c. Who is the tallest girl in the class?
 {It/That/?She} is Molly.

Mikkelsen (2005, §5.1) argues that anaphors must agree in type with their antecedent, and that the antecedents of the anaphors in these sentences are the subjects of the specificational clauses. On the equative analysis, however, the subjects of these specificational clauses should be *e*-type expressions that refer to a human, Molly. So the equative analysis either needs to dispute Mikkelsen's argument that a pronoun reflects the type of its antecedent, or show that expressions which refer to humans actually *can* pronominalize with neutral anaphors like 'it' and 'that'. As far as I can see, neither Heycock and Kroch (1999) nor Heller (2005) (who is aware of Mikkelsen's argument) pursue either option.

But the most serious problem with the equative analysis is that it does not adequately account for the asymmetry of role in specificational sentences. For all we have said so far, the equative analysis cannot distinguish specificational sentences like those in (33) from 'pure' equatives like those in (34).

- (33) a. The lead in *North by Northwest* is Cary Grant.
 b. What Mary did was run the marathon.
- (34) a. Cary Grant is Cary Grant.
 b. What Mary did is what I'm going to do next month.

According to the equative analysis, both types of sentences are equatives, so a sentence in either group is true if and only if its subject and complement have the same denotation. Yet the specificational sentences in (33), in some obvious and intuitive sense, exhibit an asymmetry that the 'pure' equatives in (34) do not. In the 'pure' equatives, we cannot construe one expression as expressing a constraint, and the other as saying what satisfies it, any more than the other way around. But in specificational sentences, these different roles always belong to the subject and complement respectively. This asymmetry of role is what makes specificational sentences special, and distinguishes them from equatives more generally. Since the equative analysis is an analysis of *specification*, it is

obliged to give an account of the asymmetry.¹⁶

Heycock and Kroch recognize this problem, but their solution is hard to make sense of. Basically, they suggest that the syntax of a specificational pseudo-cleft instructs an interpreter to interpret its subject and complement in different ways. Rather than computing denotations for the subject and complement independently and then comparing the results, the interpreter should compute the denotation of the complement and then *assign* it as the denotation of the subject. This suggestion recovers an asymmetry in how the subject and complement are interpreted in specificational sentences. Here is how they put it:

In [specificational] pseudoclefts, the free relative [i.e., the subject] expresses the ground, an open proposition that the speaker assumes to be a salient member of the belief set that the hearer is using to interpret the discourse; and the focus [i.e., the complement] of the pseudocleft expresses the value of the variable that the speaker intends the hearer to add to his or her belief set. ...

If we accept the above treatment of the information structure of pseudoclefts, we see quite easily how the equative syntax of a pseudocleft encodes its information-packaging effect; but the analysis of pseudoclefts as equative sentences takes on a somewhat different significance. In logic, $a = b$ and $b = a$ are indistinguishable; but in natural language, the two arguments of an equative sentence typically differ in informativeness. It is this difference in information packaging that has been wrongly described as a subject-predicate asymmetry by proponents of the syntactic inversion analysis of copular sentences. *The asymmetry is better captured by reconceiving "equal-*

¹⁶One might think that the contrast between (33) and (34) is purely a pragmatic contrast, not a semantic one, and so the equative analysis is not obliged to account for it. Against this, I would point out that this sort of contrast is quite stable across contexts. No particular contextual setup is required to feel the contrast; it seems to arise from the default interpretation of these sentences. Thus, the contrast seems to reflect a general fact about what these *sentences* mean, not a particular fact about what *speakers* mean *by* them, or about how they are interpreted in particular contexts. For that reason, I think of the relationship as a semantic one, rather than a purely pragmatic one. I discuss this contrast further, as well as how the category of specificational sentences relates to the distinction between semantics and pragmatics, in Section 4.3.

ity” as an instruction to assign a value to a variable, that is, as an instruction to assign to the precopular phrase the denotation of the post-copular one. (Heycock & Kroch, 1999, 393–394, my emphasis)

This solution seems to make room for the asymmetry of role in specificational sentences only by giving up the idea that they are semantically equative in any substantial sense. It is not clear how much a ‘pure’ equative has in common with a sentence that instructs an interpreter to assign a value to a variable. A pure equative presumably does not contain such an instruction, since that would introduce exactly the asymmetry which is absent in pure equatives. So by reconceiving the ‘equality’ in specificational sentences as an asymmetric assignment instruction, Heycock and Kroch are in danger of divorcing this notion from the notion of semantic equality that applies to pure equatives.

Their suggestion does preserve the idea that interpreting a specificational sentence as true involves assigning the same denotation to the subject as the complement. But this is not sufficient to say that a specificational sentence is interpreted equatively, or that it belongs to the semantic category of equatives. For Heycock and Kroch, how an expression is interpreted depends on *both* the denotations of its constituents and its information packaging. In particular, two copular sentences whose subjects and complements all have the same denotation, such as (33-a) and (34-a), can still be interpreted differently when they differ in information packaging. In fact, that is the whole point: Heycock and Kroch’s explanation of connectivity effects requires that interpreting a specificational pseudo-cleft involves ι -reducing the sentence to a *non*-equative form. So when they say that the equative analysis of specification “takes on a somewhat different significance” on this proposal, they seem to be severely understating the point. If their suggestion is taken seriously, it becomes totally unclear that we have any reason to think specificational sentences are interpreted equatively, or to analyze them as *semantically* equative, at all.

The very first sentence in the passage quoted above indicates that something has gone wrong here. There, Heycock and Kroch suggest that the information packaging in a specificational pseudo-cleft takes the following form. The subject expresses an open proposition, while the complement expresses a value which satisfies that open proposition. But notice that the general semantic type of an open proposition is $\langle \alpha, t \rangle$,

while the type of an expression denoting a value which satisfies that open proposition is α . This agrees exactly with the ‘generalized inversion analysis’ I suggested in Table 3.2. So just at the point when they try to capture the asymmetry of role in specificational sentences, Heycock and Kroch’s analysis looks like it treats specifications as inverted predications rather than equatives. This is a result that they take great pains to reject earlier in their discussion. Thus, I conclude that Heycock and Kroch have not offered any satisfactory account of the asymmetry.

Heller’s (2005) version of the equative analysis, on the other hand, offers a much more interesting approach to the asymmetry. She proposes that specificational sentences have *rising discriminability*. This means that in a specificational sentence, the complement is somehow ‘more discriminating’ than the subject. Rising discriminability distinguishes specifications from pure equatives (as well as predications). In a pure equative, the subject and complement have equal discriminability.

To make sense of this idea, Heller demonstrates that there is a systematic pattern in how speakers choose between expressions which refer to the same individual. Speakers generally prefer proper names to definite descriptions, which they in turn prefer to free relatives. Consider, for example, the following sentences¹⁷:

- (35) a. Giacomo brought the lasagna.
b. The neighbor brought the lasagna.

In a context where Giacomo is the neighbor and both the speaker and hearer know who ‘Giacomo’ refers to, (35-a) is preferred to (35-b) as a way of saying that he brought the lasagna. Assuming that speakers prefer to use expressions which will make it clearest to their listeners which object they are referring to, this preference indicates that proper names are more discriminating than definite descriptions. This seems like the right result: though many people might satisfy the description ‘the neighbor’ in this context, only one (contextually relevant) person is named ‘Giacomo’, so by using the latter, the speaker makes it clearest who she is talking about.

If that is right, then we can explain the difference between a pure equative like (36-a) and a specificational sentence like (36-b):

¹⁷These sentences are based on an example in Heller (2005, §4.1.1).

- (36) a. Giacomo is Giacomo.
 b. The neighbor is Giacomo.

The difference is that, in the pure equative, the subject and complement have equal discriminability, whereas in the specificational sentence, the complement is more discriminating than the subject. The asymmetry of role found in specificational sentences is thus a matter of their subjects and complements differing in discriminability.

This idea strikes me as being on the right track. But rather than developing a complete theory of discriminability, Heller stops short, and so her proposal does not provide a fully satisfying account of the asymmetry of role in specificational sentences. She eventually retreats from the claim that proper names are more discriminating than definite descriptions, holding instead that “discriminability is not fixed” and suggesting that the relative discriminability of two expressions is highly sensitive to contextual factors, which she only partly explores (Heller, 2005, p. 182). More importantly, as I noted above, Heller is focused on specificational sentences with *e*-type complements. She does not extend her observations about discriminability to higher-type expressions, instead leaving that for further research (Heller, 2005, p. 181). It is not clear, however, that her concept of discriminability can be extended to accommodate such cases.

Heller says that a speaker chooses among referring expressions based on “the familiarity of the interlocutors with the entity” to which he is referring, and she explains discriminability in terms of a systematic pattern in how we make such choices (Heller, 2005, p. 132). The trouble is that we find exactly parallel choices among expressions which are not referring expressions, but instead play some other semantic role. For example, compare (37-a) to (37-b):

- (37) a. Audrey is coming to campus by bicycle.
 b. Audrey is coming to campus how she usually does.

In a context where the audience knows what it means to come by bicycle but does not necessarily know Audrey’s transportation habits, a speaker will generally prefer (37-a) to (37-b) as a means of saying how Audrey is coming to campus. This preference is just like the general preference for (35-a) over (35-b), except that the speaker is choosing between the adverbial phrases ‘how she usually does’ and ‘by bicycle’. It is likewise

reflected by the asymmetry in a related specificational sentence:

(38) How Audrey usually comes to campus is by bicycle.

In this case, though, the pattern is not well-described by saying that the speaker chooses based on the interlocutors' familiarity with the entity he is referring to, since it is not plausible that he is using these adverbial phrases to *refer* to anything at all. Even if we grant that he refers to something with these expressions, what would it mean for an interlocutor to be more or less familiar with the referent of an adverbial phrase? We would have to make sense of this obscure idea before it could explain the asymmetry in (38).

So although I think we need something *like* Heller's concept of discriminability, it is not general enough in its present form to account for the asymmetry of role in specificational sentences. I conclude that the equative analysis has no satisfying explanation of the asymmetry, and a different account is needed.

3.3 THE QUESTION-ANSWER ANALYSIS

Here is a summary of where we are at. Different empirical and theoretical considerations support the inversion analysis and the equative analysis. The inversion analysis is supported by data about how specificational subjects pronominalize. It accounts for the asymmetry of role in specificational sentences by assigning the subject and complement different semantic types, but it does not account for their symmetry of subject matter. As a consequence, the inversion analysis has trouble explaining why not all inverted predications are specificational sentences, and why some simple substitution inferences are valid.

The equative analysis is supported by data about connectivity effects. It accounts for the symmetry of subject matter in specificational sentences by assigning the subject and complement the same denotation, and it can explain the cases that the inversion analysis can't. But it has no explanation of the pronominalization data which support the inversion analysis, and it provides no satisfying account of the asymmetry of role in specificational sentences.

Thus, the inversion analysis and the equative analysis have complementary virtues and vices. Each accounts for one important set of data

that the other does not, and each captures one part of our characterization of specification that the other does not. It is time to see whether we can unify their virtues into a single account.

Unfortunately, we can't simply combine the inversion analysis and the equative analysis into an account with the virtues of both. At the empirical level, we have seen that the inversion analysis and the equative analysis each have trouble accounting for some data that support the other. At the theoretical level, the two analyses take incompatible approaches to representing the symmetry of subject matter and asymmetry of role in specificational sentences. Both analyses are presented in the framework of a typed denotational semantics, where it is assumed that differences in semantic type imply differences in denotation. Within this framework, we cannot simultaneously assign the subject and complement of a specificational sentence the same denotation and different types. Thus, we cannot combine the equative approach to the symmetry with the inversion approach to the asymmetry.

A unified analysis must therefore take a different tack at both levels. I will argue that the question-answer analysis is the best candidate to do so. According to the question-answer analysis, the subject of a specificational sentence functions like a question, and the complement functions like its answer. Versions of this analysis have been offered by Ross (1972), den Dikken, Meinunger, and Wilder (2000), Schlenker (2003), and Romero (2005). I will focus here on arguing for the question-answer analysis in informal terms, using the empirical and theoretical considerations we have already seen.

3.3.1 *Support for the question-answer analysis*

To motivate the question-answer analysis, recall that our paradigm cases of specificational sentences are pseudo-clefts. The subject of a pseudo-cleft begins with a *wh*-word, just like a question. Intuitively, a specificational pseudo-cleft reads like a sentence containing both a question and its answer, as we saw in Chapter 1. Moreover, for every specificational pseudo-cleft, there is a corresponding question-answer dialogue. There is a clear parallel in meaning between a specificational pseudo-cleft like

(39) What I bought for mother was some flowers.

and its corresponding question-answer dialogue, like

- (40) a. What did you buy for mother?
 b. Some flowers.

This intuitive parallel underlies the directness test, which I employed above as a way of identifying the specificational reading of pseudo-clefts which also have a predicational reading.

There are also some empirical observations which support the question-answer analysis. Let's first consider some syntactic observations, starting with one made by Ross (1972). In the question-answer dialogue in (41), the first two answers not full sentences, but in some sense *elided*. Full-sentence answers like (41-c) are also possible:

- (41) What did you do then?
 a. Call.
 b. Call the grocer.
 c. I called the grocer.

Ross observed that there is a parallel phenomenon in the complement position of specificational pseudo-clefts.

- (42) a. What I did then was call.
 b. What I did then was call the grocer.
 c. What I did then was I called the grocer.

Notice that the possible complements in these pseudo-clefts exactly mirror the possible answers for the corresponding question in the dialogue in (41). In general, as den Dikken, Meinunger, and Wilder (2000) and Schlenker (2003) observe, it appears that the principles of ellipsis which apply to answers in question-answer dialogues mirror those which apply to the complement of a specificational pseudo-cleft, except that such ellipsis seems to be more often obligatory in pseudo-clefts. This is evidence that specificational pseudo-clefts function like question-answer pairs, at least at a syntactic level.

This syntactic parallel also suggests a natural explanation of connectivity effects. If specificational complements are syntactically analyzed as elided full-sentence answers, connectivity effects can be explained by the elided material. For example, consider again sentence (28-b):

- (28-b) What Narcissus likes is staring at himself.

Recall that the connectivity effect in this sentence is that the reflexive pronoun ‘himself’ is grammatical in the complement, even though it does not appear to be in a configuration which licenses reflexives. But if the complement is treated as an elided sentence, like

(43) What Narcissus likes is ~~Narcissus~~-likes staring at himself.

then ‘himself’ occurs in the proper configuration to be licensed by the elided ‘Narcissus’.

Den Dikken, Meinunger, and Wilder (2000) and Schlenker (2003) develop syntactic analyses of specificational pseudo-clefts to explain connectivity effects in this manner. They also extend their analyses to explain connectivity in specificational sentences with other types of subject. I cannot evaluate here whether all specificational sentences should be analyzed syntactically as question-answer pairs, or whether this approach will ultimately provide a satisfying explanation of connectivity effects. Still, it appears that the question-answer analysis has a natural path to such an explanation, which means that it plausibly offers the same explanatory advantages as the equative analysis with respect to connectivity data.

Several semantic observations also support the question-answer analysis. First of all, the analysis is consistent with the pronominalization data which support the inversion analysis. Recall that Mikkelsen (2005) showed that human-denoting descriptions in the subject position of a specificational sentence can pronominalize with neutral anaphors like ‘it’ and ‘that’. She argued that this is not consistent with such descriptions having type e , and instead concluded that they have type $\langle e, t \rangle$.¹⁸ But in fact, these data only seem to require that we not assign type e to the subject of a specificational sentence. As Mikkelsen (2005, p. 61) herself notes, treating the subject of a specificational sentence as a question is compatible with the pronominalization data. This is because questions, like predicative expressions, pronominalize with neutral anaphors.

We can see this by considering issue-predicational sentences. As I observed in Section 2.1.3, issue-predicational sentences are sentences where the complement is predicated of what the subject denotes, but the subject denotes a question or issue. When the subject of an issue-predicational

¹⁸See the discussion above of the examples in (3).

sentence is anaphoric on the previous expression of a question, it will be a neutral anaphor, as the following examples show.

- (44) a. How many people are expected for lunch?
 That is a good question. (I wish I knew the answer.)
 b. Who we should invite is still being discussed, and I expect *it*
 will not be an easy question to reach agreement on.

The answer in the dialogue in (44-a) is an issue-predicational sentence. It predicates ‘being a good question’ of what its subject denotes. But its subject is the neutral anaphor ‘that’, which is only plausibly interpreted as having the prior question as its antecedent. Similarly, the second conjunct of (44-b) is an issue-predicational sentence with ‘it’ as its subject, which is only plausibly interpreted as anaphoric on the question expressed by ‘who we should invite’ in the first conjunct. Note that the first conjunct is also an issue-predicational sentence, which makes it clear that ‘who we should invite’ denotes a question there, rather than a person or group of people.¹⁹

We can also see that questions pronominalize with neutral anaphors by looking at verbs like ‘to wonder’, which are standardly analyzed as taking questions as arguments.

- (45) Silas wondered who Bertram would marry,
 a. ... but concluded that *it* was an inscrutable question.
 b. ... and concluded that {*it*/??*he/??*she} would be Leslie.

In both of the examples in (45), the question ‘who Bertram would marry’ is the argument of ‘wondered’ in the first conjunct. When the sentence is continued with a second conjunct containing an issue-predicational clause, as in (45-a), a neutral pronoun is grammatical, and the only plausible antecedent for the pronoun is the question in the first conjunct. When the second conjunct instead contains a specificational clause, as in (45-b), a neutral anaphor is still grammatical, but gendered pronouns are very marginal, even when the specificational complement refers to a person. This is consistent with Mikkelsen’s observation that it is preferable to pronominalize a specificational subject with a neutral anaphor,

¹⁹: See the discussion of example (26) in Section 2.1.3 above.

even when a gendered pronoun is possible.²⁰

The example in (45-b) also suggests a further point. When the subject of a specificational clause is an anaphor, it is required to agree in type with its antecedent. In all of Mikkelsen's examples, the antecedent is a description; and since descriptions can have type $\langle e, t \rangle$, they point toward treating specificational subjects as predicates. But there do not seem to be any plausible predicative antecedents for 'it' in the first conjunct, unless we construe the question 'who Bertram would marry' as a predicate. If there are reasons for assigning this question a type other than $\langle e, t \rangle$, then in cases like (45-b), the question-answer analysis may actually explain the pronominalization behavior of specificational subjects *better* than the inversion analysis does.²¹

Another semantic observation which supports the question-answer analysis is that descriptions can sometimes make the same semantic contributions as questions. This is an important point for the question-answer analysis to establish, since it is clear that many specificational sentences have a description in the subject position. Unless descriptions can be analyzed as behaving semantically like questions, the question-answer analysis will fall flat as a general proposal about the semantics of specification.

In fact, though, the claim that descriptions can express questions is not controversial. One way to see this is to look at how descriptions interact with verbs like 'to tell' and 'to find out'. These verbs are called *concealed question verbs* because in addition to explicit questions, they can take descriptions as arguments:

- (46) a. Tell me what your favorite book is.
b. Tell me your favorite book.
- (47) a. Edward found out where the treasure was.
b. Edward found out the location of the treasure.
- (48) a. Charlotte knows when the meeting is.
b. Charlotte knows the time of the meeting.

²⁰See the judgments in example (3) above. Note that Mikkelsen interprets the possibility of gendered pronouns in the subject position as indicating an ambiguity between specificational and equative readings of the clause in question.

²¹Schlenker (2003) makes similar observations, employing an anaphoric relation between the subject of a specificational sentence and the argument of 'to wonder' to argue that specificational subjects should be construed as questions.

Because they also accept explicit questions as semantic arguments, these verbs are standardly analyzed as requiring a question in their object positions. Thus, the fact that these verbs also allow descriptions in this position indicates that some descriptions can make semantic contributions akin to explicit questions. (We made similar observations in Chapter 1.) Such a description is called a ‘concealed question’ in the literature.

Other recent work in linguistics shows that descriptions are usefully analyzed as behaving like questions in a variety of contexts. For example, Barker (2016) shows how a natural semantics for questions can explain a distributional difference in the head nouns of descriptions which can function as concealed questions. AnderBois (2010) develops a semantics that treats indefinite descriptions like questions in order to explain licensing conditions on sluicing. And Romero (2005) shows how an ambiguity between two readings of concealed questions, first observed by Heim (1979), is also present for descriptions in the subject position of a specificational sentence. She develops a semantics which treats specificational subjects as concealed questions in order to generate both readings.

Thus, empirical observations about both the syntax and the semantics of specificational sentences support treating them like question-answer pairs. In particular, the question-answer analysis seems to have a plausible story about both sets of data that supported the equative and inversion analyses. Since neither of the other analyses had an explanation for both sets of data, the question-answer analysis looks like the most promising candidate for a unified empirical account.

3.3.2 *An empirical challenge*

Before I can conclude that the question-answer analysis is empirically satisfactory, however, I need to respond to some objections raised by Caponigro and Heller (2007). Their objections are the strongest empirical claims against the question-answer analysis that I am aware of. They argue that, despite the observations we’ve made above, specificational subjects cannot be construed as questions, either syntactically or semantically. I accept their claims about syntax, but I think their semantic claims fail. Let’s look at both arguments in turn.

On the syntactic side, Caponigro and Heller (2007) argue that the subjects of specificational pseudo-clefts belong to the category of *free relatives* rather than the category of *interrogatives*. They appeal to data from lan-

guages that morphologically distinguish these categories (namely Macedonian, Hungarian, Wolof, and Hebrew), and show that, in each of these languages, free relatives can appear as the pre-copular phrase in a specificational sentence, while interrogatives cannot. (Caponigro & Heller, 2007, § 4.1)

English does not morphologically distinguish these categories; some English phrases that begin with *wh*-words work as both interrogatives and free relatives. But Caponigro and Heller (2007) demonstrate that there is a distributional difference between these categories, so there is still reason to treat them as syntactically distinct. In English, interrogatives that begin with ‘which’ or ‘how much’ cannot also function as free relatives. This provides a test of whether an environment accepts free relatives or interrogatives. They observe that verbs like ‘wonder’ accept the full range of interrogatives²²:

- (49) Interrogatives
- a. I wonder *where* she has lunch.
 - b. I wonder *what* John is reading.
 - c. I wonder *who* gave you the flowers.
 - d. I wonder *which* book John is reading.
 - e. I wonder *how much* Sue weighs.

But ‘which’ and ‘how much’ constructions cannot be used in contexts where only free relatives, and not interrogatives, are grammatical, such as the object positions of verbs like ‘to read’ or ‘to meet’:

- (50) Free relatives
- a. I have lunch *where* she has lunch.
 - b. I read *what* John is reading.
 - c. ??I met *who* gave you the flowers.
 - d. *I read *which* book John is reading.
 - e. *I weigh *how much* Sue weighs.

They then point out that the distribution of *wh*-phrases in the subject position of specificational pseudo-clefts patterns with free relatives, not

²²The examples in (49), (50), and (51) are from Caponigro and Heller (2007, §7.4.2). Interestingly, they do not consider phrases that begin with ‘whether’, ‘how’, or ‘how many’; in all of these cases, I think, the distributional difference is less clear.

interrogatives:

- (51) Specificational pseudo-clefts
- a. *Where* she has lunch is at the cafeteria.
 - b. *What* John is reading is *Ulysses*.
 - c. ??*Who* gave you the flowers was your advisor.
 - d. **Which* book John is reading is *Ulysses*.
 - e. **How much* Sue weighs is 130 pounds.

So, they conclude, the subject of a specificational pseudo-cleft is a free relative, not an interrogative.

I agree that there is a distributional difference here, and I am happy to assume that, as far as syntax is concerned, the subject of a specificational pseudo-cleft is a free relative rather than an interrogative. In fact, making this assumption may strengthen a point I made in Chapter 1. I observed there that many ‘questions’ can be intersubstituted with descriptions but that ‘which’-questions were a notable exception. If the subject position of a specificational pseudo-cleft only allows free relatives, not interrogatives, that explains the exception for the case I am interested in. Moreover, Caponigro (2003, p. 10) *defines* the category of free relatives by observing that free relatives can be intersubstituted with descriptions without change of truth conditions. He argues that this is part of what distinguishes free relatives from interrogatives, which cannot always be intersubstituted with descriptions. If Caponigro is right about this, I can then strengthen my claim: a specificational sentence with a ‘question’ (i.e., free relative) in the subject position is *always* equivalent to one with a description in the subject position.

In light of this syntactic point, I shall henceforth be more careful with my terminology. As I will use the terms, ‘free relative’ and ‘interrogative’ are strictly *syntactic* categories, which have different distributions. I will reserve the term ‘question’ for a *semantic* category. Interrogatives paradigmatically express, or are interpreted as, questions. The issue here is whether free relatives, descriptions, and other kinds of specificational subjects can also express or be interpreted as questions, and if they can, in what sense they can.

Caponigro and Heller (2007, § 7.5.2) think that specificational subjects cannot be interpreted as questions, but their argument here is much weaker. They argue for two claims. First, they argue that free relatives

cannot be complements of concealed-question verbs. Second, they argue that some descriptions which can occur as specificational subjects cannot be complements of concealed-question verbs. Together, these claims show that there is a distributional difference between the arguments of concealed-question verbs and the subject position in specificational sentences. They take this to show that, since the arguments of concealed-question verbs are interpreted as questions, specificational subjects cannot be.

To support the claim that free relatives cannot be complements of concealed question verbs, Caponigro and Heller again appeal to data from languages that distinguish free relatives and interrogatives. They show that in Hebrew, Wolof, and Hungarian, free relatives cannot appear as the complement of concealed question verbs, while interrogatives can.

Their second claim is that some descriptions can be specificational subjects, but cannot be concealed questions. To support this claim, they appeal to data from English. I reproduce their examples (41)–(43) in full (Caponigro & Heller, 2007, p. 258):

- (52) a. The president of the United States is G.W. Bush.
 b. Tell me the president of the United States.
 c. The boy who ran over my pet snake was John.
 d.*??Tell me the boy who ran over my pet snake.
- (53) a. The capital of France is Paris.
 b. Tell me the capital of France.
 c. The city I live in is Paris.
 d. ??Tell me the city you live in.
- (54) a. The candy Jill wants to buy is jelly beans.
 b. Tell me the candy Jill wants to buy.
 c. The money that was stolen was Swiss Francs.
 d. *Tell me the money that was stolen.

Each of these examples purports to contrast two types of definite description. The first type is headed by what they call a *functional* noun, like ‘president’ or ‘capital’; the other is headed by a non-functional noun, like ‘boy’ or ‘city’. They propose that while both types of description can be specificational subjects (as we see in the a- and c-examples), only descriptions headed by functional nouns can be interpreted as questions, and serve as arguments of concealed-questions verbs (as we see in the b-

and d-examples).

Unfortunately, I simply disagree with their grammaticality judgments here, so I do not see the purported contrasts. I do find (52-d) *slightly* awkward on its own, but it seems completely acceptable in a context like this:

- (55) Tell me the boy who ran over my pet snake, and who put him up to it, or else it's detention for you!

Thus, I do not see that (52-d) contrasts with (52-b). Likewise, (53-d) is totally acceptable to my ear—I can imagine using this sentence in a police interview, for example—so there is no contrast with (53-b).

As for (54-c), I cannot get a specificational reading of this sentence; it seems ungrammatical when read specificationally (though it's better when read predicationally). Thus, it's no surprise that (54-d) would be ungrammatical too, if 'the money that was stolen' is interpreted as a question. It seems to me that the problem with both (54-c) and (54-d) is that it is just not clear what question 'the money that was stolen' is supposed to express. When the intended question is clearer, both concealed question verbs and specificational sentences admit descriptions headed by 'money' just fine:

- (56) First we need a clear budget. You tell me the money we need to start up, and then I'll take care of finding investors.

- (57) The money we need to start up is two hundred thousand dollars.

So again, I see no reason to think that descriptions headed by 'money' contrast with descriptions headed by 'candy' with respect to their behavior as specificational subjects and as concealed questions.

But perhaps you disagree with my judgments, or perhaps Caponigro and Heller could come up with better examples. Even if that is so, there is a more significant problem with the argument: it is simply not probative for the issue at hand. The issue here is precisely whether differences in *syntactic* distribution should lead us to conclude that there is a *semantic* difference between specificational subjects and the arguments of concealed question verbs. Even if there are some specificational subjects that cannot serve as arguments of concealed question verbs, it is clear that many can. In order to claim that there is a semantic difference between the two positions, Caponigro and Heller need to show that when

a given phrase can appear in both positions, it is nevertheless interpreted differently.

Caponigro and Heller (2007, §7.5.3) make a final argument for this claim, using ‘what the capital of France is’ as an example of an expression that can appear both as a specificational subject and as an argument of a concealed question verb. They claim that this phrase admits of different answers in the two contexts, which shows that it is interpreted differently. When it appears as an argument of ‘to tell’, ‘Paris’ is a felicitous answer, while ‘beautiful’ is not:

- (58) Tell me what the capital of France is.
 a. Paris.
 b. #Beautiful.

But the opposite pattern occurs when it appears as a specificational subject:

- (59) What the capital of France is is ...
 a. *Paris.
 b. beautiful.

Since ‘what the capital of France is’ is interpreted differently in the two positions, Caponigro and Heller conclude that the subject position of a specificational sentence is not a concealed question environment, so treating a free relative in this position as expressing a question is at best stipulative.

This still seems like quite weak evidence against the question-answer analysis. First of all, as we observed above²³, ‘what’-questions are ambiguous, so it is not surprising in general if string-identical ‘what’-questions can be interpreted differently in different contexts. If we look at ‘where’-questions, which are much less ambiguous, we see that appropriate answers in the two kinds of case do not come apart:

- (60) a. Tell me where Paris is.
 Paris is {in France/in Europe/northeast of Le Mans... }.
 b. Where Paris is is {in France/in Europe/northeast of Le Mans... }.

²³See the discussion of example (3) in Chapter 2.

This parallel needs to be explained, and suggests that analyzing specificational subjects as questions is not merely stipulative.

Second of all, when the context disambiguates between the different senses of the 'what'-question, the contrast between (58) and (59) vanishes. In a context which makes clear that it is the *qualities* of the capital of France which are being asked for, 'beautiful' would be a perfectly felicitous answer to (58), and 'Paris' would be infelicitous. Likewise, in a context which makes clear that it is the *name* of the capital of France which is wanted, 'Paris' can be felicitous in the complement of (59), while 'beautiful' would not be.²⁴ While I agree that 'what the capital of France is' is by default interpreted differently in the two constructions, the difference seems to lie in *which* question it is interpreted as expressing, not in *whether* it is interpreted as a question at all. So I do not think Caponigro and Heller have successfully challenged the claim that the subject of a specificational sentence functions semantically as a question, or that it is interpreted as one.

3.3.3 *Symmetry of subject matter and asymmetry of role*

I now want to argue that the question-answer analysis can succeed where the inversion analysis and the equative analysis both fail. The question-answer analysis can account for the fact that the subject and complement of a specificational sentence simultaneously exhibit symmetry of subject matter and asymmetry of role. Combined with its empirical strengths, this makes the question-answer analysis the best candidate for a unified and satisfactory analysis of specification, according to the criteria guiding my discussion here.

At the level of complete sentences in discourse, there is a clear and intuitive sense in which questions and their answers exhibit both symmetry of subject matter and asymmetry of role. When someone asks a question, and someone else in the conversation answers it, the question and the answer are 'about the same thing' in the sense that they address the same topic. At the same time, it is obvious that questions and their answers play different roles in a conversation. They convey different

²⁴I admit, though, that it would be unusual for a speaker to use a 'what' pseudo-cleft in such a context at all. More likely, she would simply use 'the capital of France' as the subject in (59). Still, using the free relative version of this sentence to specify the name of the capital of France does not strike me as *impossible*.

information to the audience, and are governed by different norms. So questions and answers provide a model of how two sentences can be about the same thing when they are used in a particular context, yet differ in their role, impact, and meaning. The question-answer analysis of specification extends this model to the subsentential level, in order to account for the relationship between the subject and complement in specificational sentences. Specificational subjects and complements exhibit symmetry of subject matter and asymmetry of role for the same reasons that sentential questions and answers do.

To see that this is so, let's look more closely at how sentential questions and answers are interpreted. In what sense do they exhibit symmetry of subject matter and asymmetry of role? Consider first their symmetry of subject matter. Here is a brief question-answer dialogue:

- (61) a. Where did Marcus eat?
b. Marcus ate at the Ethiopian place on Grand Avenue.

In what sense are (61-a) and (61-b) 'about the same thing'? To count as the answer to a certain question, a claim must be *relevant* to that question. (61-b) makes a claim that is obviously relevant to the question asked by (61-a). It is a felicitous answer to (61-a) because both the question and the claim are 'about' where Marcus ate. Compare this to the claim made by (62):

- (62) Bill picked up the dry-cleaning today.

It is equally obvious that (62) is not relevant to the question expressed by (61-a), unless the context is rather peculiar. It is not about where Marcus ate, but something else. For that reason, asserting it in response to (61-a) generally would not answer the question; it would change the subject. So an answer is 'about the same thing' as a question, and shares a subject matter with it, in the sense that it is relevant to the question.

On the other hand, the question expressed by (61-a) and the answer expressed by (61-b) differ in several respects, which reflect their different roles and purposes in the conversation. For one thing, the question and the answer have different conditions of correct use. Most importantly, they differ in the epistemic requirements they place on their speakers. The questioner does not need to know where Marcus ate in order to correctly ask his question, for he can ask this question in order to acquire

such knowledge. But to correctly answer the question, the speaker *does* need to know where Marcus ate, since she is aiming to inform the questioner, not to mislead him. This pragmatic difference points to a semantic one: assuming the speakers correctly use the sentences expressing the question and the answer, the answer supplies information that the question does not. The question does not tell the audience where Marcus ate, but the answer does. Thus, the question and the answer differ both with respect to what their speakers must know in order to use them, and in what their expressions convey to the audience.

These same features also characterize questions and answers when they are expressed by subsentential expressions. One clear case involves interrogatives and ‘that’-clauses embedded under verbs like ‘to discover’:

- (63) a. Audrey discovered where Marcus ate.
 b. Audrey discovered that Marcus ate at the Ethiopian place on Grand Avenue.

The relationship between the interrogative ‘where Marcus ate’ in (63-a) and the clause ‘that Marcus ate at the Ethiopian place on Grand Avenue’ in (63-b) exactly parallels the relationship between the sentences in (61). The answer expressed by the ‘that’-clause is relevant to, and shares a subject matter with, the question expressed by the embedded interrogative. But using the ‘that’-clause in (63-b) as opposed to the interrogative requires the speaker to know more, and conveys more to the audience.

One way to observe the relevance of the ‘that’-clause in (63-b) to the embedded interrogative is to contrast it with an irrelevant one in sentences where the interrogative is explicitly topicalized:

- (64) a. As for where Marcus ate, Audrey discovered that Marcus ate at the Ethiopian place on Grand Avenue.
 b. #As for where Marcus ate, Audrey discovered that Bill picked up the dry-cleaning today.

The fact that (64-a) is felicitous, while (64-b) is not, indicates that the ‘that’-clause in (63-b) is relevant to the question expressed by the embedded interrogative in (63-a).

The asymmetry between the ‘that’-clause and the embedded interrogative in (63) also clearly parallels the asymmetry in (61). Notice that a

speaker who asserts (63-a) does not himself need to know where Marcus ate, whereas a speaker who asserts (63-b) does. Likewise, the speaker of (63-a) does not inform his audience where Marcus ate, whereas the speaker of (63-b) does. I shall put the point like this: an embedded interrogative *leaves open* the question it expresses. I mean by this that using it in discourse raises the question, but it does not convey the answer, and does not require the speaker to know the answer. The interrogative in (63-a) leaves open where Marcus ate. The 'that'-clause in (63-b) does not leave open where Marcus ate; I shall say it *closes off* this question.

Thus, the relationship of relevance, and the differences in the knowledge required by the speaker and conveyed to the audience, characterize the symmetry and asymmetry between questions and answers. This is so both when they are expressed by complete sentences and when they are expressed by subsentential expressions. These are the central features of how questions and answers are interpreted, and they are features which most semantic analyses of questions and answers attempt to capture.

Semantic analyses of questions and answers generally represent these features as follows. A question presents a set of *alternatives*, or possible answers, while an answer to that question selects among these alternatives. The alternatives represent the subject matter shared by the question and its answers. A claim is relevant to a question, and counts as an answer to it, just in case it selects among the alternatives presented by the question; if it does not select among those alternatives, it is about something else.

The asymmetry between questions and answers is represented as the difference between *presenting* alternatives and *selecting among* them. By selecting among the alternatives which a question presents, an answer closes off some alternatives that the question left open, so it conveys more information to the audience than the question itself does. Assuming that a speaker generally should not convey information which he doesn't know to be true, this implies that a speaker who expresses an answer to a question is epistemically constrained in a way that the speaker who expresses the question is not.

A typical analysis of questions would thus describe the examples in (63) like this. The embedded interrogative in (63-a) presents a set of alternatives: the sentence is true if Audrey discovered that Marcus ate at the Ethiopian place on Grand Avenue, or at the ballpark, or in his bedroom, or There are indefinitely many such alternatives, any of which would

be one in which Audrey discovered where Marcus ate; (63-a) is true just in case one of these alternatives obtains. By contrast, the ‘that’-clause in (63-b) selects one of these alternatives, excluding any others which are incompatible with it. This sentence is true just in case Audrey discovered that Marcus ate at the Ethiopian place on Grand Avenue, which implies she did *not* discover that he ate in his bedroom (unless ‘the Ethiopian place on Grand Avenue’ can be interpreted as a veiled reference to Marcus’ quarters).

How does this picture apply to specificational sentences? The crucial observation we need is this: the kinds of expressions which can be specificational subjects and complements systematically exhibit the same features as other subsentential expressions of questions and answers, such as embedded interrogatives and ‘that’-clauses. Like embedded interrogatives, specificational subjects leave open a certain question; like ‘that’-clauses, their corresponding complements are relevant to that question, but do not leave it open. This observation holds of these expressions quite generally, both when they appear in specificational sentences and outside them. Thus, it makes sense to apply a semantic analysis of questions and answers to them. We can analyze specificational subjects as presenting a set of alternatives; and we can analyze specificational complements as selecting among such alternatives.

Let’s look first at free relatives, which appear as subjects in specificational pseudo-clefts. We observed in Section 3.3.2 that free relatives differ syntactically from interrogatives. But semantically, there is a clear parallel between them: a free relative leaves open the same question that the corresponding embedded interrogative does, and that would be asked by the corresponding sentential interrogative. Here are some examples of interrogatives, parallel to the example in (63-a):

- (65) a. Nora discovered what James read.
 b. Nora discovered what James is.
 c. Nora discovered how James traveled.

And here are some examples in which the corresponding free relatives appear:

- (66) a. Nora read what James read.
 b. Nora is what James is.
 c. Nora traveled how James traveled.

Each of the free relatives in (66) leaves open the same question as the corresponding interrogative in (65). The free relative in (66-a), for example, is like the interrogative in (65-a) in that it leaves open the question expressed by “What did James read?” A speaker who uses this free relative in (66-a) raises this question to salience. But she does not need to know the answer to this question in order to make her assertion, and her assertion does not convey the answer to her audience. Similarly, the free relative in (66-b) leaves open “What is James?” and the free relative in (66-c) leaves open “How did James travel?”. Free relatives are just like interrogatives with respect to the epistemic constraints they place on the speaker and the information they convey to the audience when used to make an assertion.

To see this more clearly, it helps to compare the sentences in (66) with parallel examples where the questions left open by the free relatives are closed off:

- (67) a. James read *Ulysses*, and Nora read it, too.
 b. James is an elephant, and Nora is, too.
 c. James traveled on a barge, and Nora did too.

Each of these sentences, like the sentences in (66), links something true of Nora to something true of James, but it closes off what its counterpart in (66) leaves open. For example, (67-a), like (66-a), tells the audience that James and Nora read the same thing. But it also tells us something more: it tells us *what* James read, namely, *Ulysses*. A speaker who asserts this sentence must be able to directly answer “What did James read?”, while a speaker who asserts (66-a) need not. The sentences in (67-b) and (67-c) similarly contrast with (66-b) and (66-c).

Each sentence in (67) closes off the question left open by the free relative in its counterpart in (66) because its first clause expresses a relevant answer to that question. We can confirm this by noticing that the sentences in (67) remain felicitous when the question they answer is explicitly topicalized:

- (68) a. As for what James read, James read *Ulysses*, and Nora read it, too.
 b. As for what James is, James is an elephant, and Nora is, too.
 c. As for how James traveled, James traveled on a barge, and Nora did too.

Thus, the same features characterize the relationship between the free relatives in (66) and the initial clauses in (67) as we found between embedded interrogatives and 'that'-clauses in examples like (63), and between interrogative and declarative sentences in examples like (61). They are related as questions and answers. In each case, the latter expressions are relevant to the questions that the former leave open, but correctly using them requires the speaker to know more, and conveys more to the audience.

We can therefore capture the relationships in (66) and (67) with the same sort of analysis. Like interrogatives, the free relatives in (66) present a set of alternatives. (66-a) is true just in case James read *Ulysses*, or the evening *Post*, or a new translation of Homer's poetry, or . . . , and Nora read it, too. (66-b) is true just in case James is an elephant, or happy, or extremely infectious, or . . . , and Nora is that, too. And (66-c) is true just in case James traveled on a barge, or along the Rhine, or without luggage, or . . . , and Nora did so, too.

The sentences in (67) differ from their counterparts in (66) by selecting among these alternatives. In each case, we can isolate the selection made by the closed-off sentence to a single expression. In (67-a), it is '*Ulysses*' which tells us what James read. Likewise, in (67-b), 'an elephant' tells us what James is, and in (67-c), 'on a barge' tells us how he traveled. These expressions are the foci of answerhood in (67), the expressions which tell us something above and beyond what the sentences in (66) do.

Suppose we now pair the free relatives from (66) with the answering expressions from (67) around a copula:

- (69) a. What James read was *Ulysses*.
 b. What James is is an elephant.
 c. How James traveled was on a barge.

The resulting sentences, of course, are specificational pseudo-clefts, and they mirror the question-answer relationships between the sentences in (66) and (67). Their complements express relevant answers to the questions left open by their subjects, just as the initial clauses in (67) do for the questions left open by the free relatives in (66). And their complements close off what their subjects leave open, just as the sentences in (67) close off what those in (66) leave open. In light of this parallel, we may analyze these specificational pseudo-clefts as follows. The free relative in

the subject position presents a set of alternatives, just as it does in the corresponding sentence in (66). The expression in the complement position selects among these alternatives, just as it does in the corresponding sentence in (67).

This is how the question-answer analysis explains the distinctive semantic features of specificational sentences. It assimilates the symmetry and asymmetry between the expressions *within* specificational sentences to the symmetry and asymmetry these expressions exhibit when used *outside* specificational sentences, in examples like (66) and (67). Since the subject and the complement deal with the same set of alternatives, they exhibit symmetry of subject matter. But because the subject *presents* these alternatives while the complement *selects among* them, they exhibit asymmetry of role. Once we recognize that the symmetry and asymmetry in specificational sentences is an instance of a broader phenomenon, this explanation seems simple, natural, and what a principle of compositionality would lead us to expect.

The same considerations apply, with only minor adjustments, to specificational sentences which are not pseudo-clefts. Consider versions of the examples in (66), (67), and (69) that exchange free relatives for definite descriptions:

- (70) a. Nora read the book James read.
- b. Nora is the kind of animal James is.
- c. Nora traveled the way James traveled.
- (71) a. James read *Ulysses*, and Nora read it, too.
- b. James is an elephant, and Nora is, too.
- c. James traveled on a barge, and Nora did too.
- (72) a. The book James read was *Ulysses*.
- b. The kind of animal James is is an elephant.
- c. The way James traveled was on a barge.

The relationships in these examples are just the same as those in (66)–(69).²⁵ The definite description in each sentence in (70) leaves open a certain question. The initial clause in its counterpart in (71) is relevant to this question, and supplies an answer to it. Placing the definite description and the focus of the answer on either side of a copula yields

²⁵I have repeated the examples from (67) in (71) for clarity, though they are unaltered.

a specificational sentence in (72). The symmetry of subject matter and asymmetry of role in that specificational sentence mirrors the question-answer relationship of the sentences in (70) and (71).

Perhaps it does not seem obvious that there are any questions involved here, now that the *wh*-words ‘what’ and ‘how’ are gone from view.²⁶ But the important point is that each of the definite descriptions in (70) expresses a question in the sense that it is interpreted as presenting a set of alternatives. It thus leaves a certain question open, just like the free relatives in (66) do, and the sentences in (71) do not. Which questions do they leave open? As a first approximation, we can try the same questions that the free relatives leave open, the questions expressed by “What did James read?”, “What is James?” and “How did James travel?” But this is not quite right, because the sentences in (70) do not leave these questions quite as open as the sentences in (66) do. The speaker who asserts (70-a), for example, must at least know that what James read was a *book*, and her assertion conveys this fact to her audience. So we have to be a little more refined: the questions left open by the descriptions in (70) are more like “Which book did James read?”, “Which kind of animal is James?” and “Which way did James travel?”

The question-answer analysis offers a natural way to understand this difference. The definite descriptions in (70) present a different set of alternatives than the free relatives in (66) do. (70-a) is true just in case James read *Ulysses*, or *Ivanhoe*, or *Great Expectations*, or . . . , and Nora read it, too. But it cannot be true in cases where James and Nora merely both read yesterday’s newspaper or a short poem, while (66-a) can. Similarly, (70-b) is true just in case James is an elephant, or a tree frog, or a musk ox, or . . . , and Nora is that, too. But it cannot be true in cases where Nora and James are merely both happy or both extremely infectious, while (66-b) can. There does not seem to be any real barrier, then, to analyzing the symmetry and asymmetry in the specificational sentences in (72) as the question-answer analysis does: their definite-description subjects present a set of alternatives, and their complement phrases select among these alternatives.

One may still feel some hesitation. In a sentence like (70-a), doesn’t the definite description ‘the book James read’ *refer* to a certain book? It

²⁶Remember, though, that definite descriptions are standardly analyzed as expressing questions in some contexts. See the discussion of examples (46)–(46) in Section 3.3.1.

seems to be doing just the same work that the title '*Ulysses*' would do in the same position, namely, picking out an individual thing. It does not then express or leave open a question, and it should not be analyzed as presenting a set of alternatives, either in (70-a) or in (72-a). The only 'alternative' it could present is a single object, namely, the book James read.

But this picture is too simple, as the history of debate about definite descriptions makes clear. Unlike the title '*Ulysses*', the description merely 'denotes' the book (to use Russell's terminology). A speaker may correctly assert (70-a), and her audience may perfectly well understand what she has said, without either of them knowing which particular book it was that James and Nora read. That is not true of a sentence which refers to the book by its title, such as (71-a). Picking out a particular individual book thus cannot be an essential part of the description's contribution toward what is said and understood. I prefer Strawson's way of making the point in this case: the description does not refer, though the speaker of (70-a) may be using it to refer. What she says does not determine which particular book Nora read, but by describing it as the book James read, she may be giving her audience enough information to discern which book it was, based on their knowledge of James' reading history. In using the definite description, she is offering her audience a refined set of alternatives, leaving it to them to choose among them if they can. When they can, their choice can be made explicit by a specificational sentence like (72-a). Which book was it? The book James read was (not *Ivanhoe*, not *Great Expectations*, but) *Ulysses*. Such a sentence presents and selects from the same alternatives that the speaker of (70-a) offers to her audience, and that they implicitly select from in interpreting her assertion.

I have now described how the question-answer analysis applies to specificational sentences whose subjects are free relatives or definite descriptions. To summarize: the subjects of these sentences leave open a certain question, and may be analyzed as presenting a set of alternatives. Their complements close off that question, by selecting among the alternatives the subjects present. This analysis is plausible because it treats the two expressions in a specificational sentence as having the same relationship which they exhibit outside it. The features of this relationship are the semantic features which characterize the relationship between expressions of questions and answers generally: their symmetry of subject matter, thought of in terms of relevance; and their asymmetry

of role, thought of in terms of the different information the expressions convey when used, and the different epistemic requirements a speaker must meet to use them correctly.

These two cases—specificational sentences with free relatives or definite descriptions for subjects—have received the most attention in the literature. The inversion and equative analyses target these cases, and have plausible stories to tell about them. I argued above that these stories were not entirely satisfying, because they cannot simultaneously account for the symmetry of subject matter and asymmetry of role in specificational sentences. Thus, the fact that the question-answer analysis can account for both of these features within the scope of a single account is enough to recommend it over the others.

3.3.4 *Extending the analysis*

Once we recognize specificational sentences as presenting and selecting among alternatives, though, the question-answer analysis also extends naturally to some other cases. I want to briefly examine these cases to further demonstrate the virtues of the question-answer analysis, and to show how it can accommodate the insights of the other two.

To extend the question-answer analysis, we need one further observation about questions and answers. Questions can impose different sorts of constraints on *how* their answers must select among the alternatives they present. To make this vivid, imagine facing each of the following questions on a history exam:

- (73) a. Which Roman emperor invaded Britain?
b. Which Roman emperors invaded Britain?
- (74) a. Who was one Roman emperor that invaded Britain?
b. Who were some Roman emperors that invaded Britain?

Correctly answering each of these questions requires naming Roman emperors who invaded Britain. In that respect, the questions are similar, and intuitively concern the same subject matter. Thus, we can analyze them as presenting the same alternatives—namely, that Claudius was a Roman emperor who invaded Britain; that Septimius Severus was a Roman emperor who invaded Britain; and so on. But some of these questions are more demanding than others. The questions differ with respect to

how many such emperors their answers must give, and whether answering correctly requires listing all of them. I shall say they make different *requests*. (73-a) and (73-b) both request answers which select all the alternatives which are true, while (74-a) and (74-b) can be correctly answered by providing an incomplete list of true alternatives. Within each of these pairs, the two questions also differ with respect to how many alternatives they request their answers to select. The first question requests a single alternative, while the second question gives no upper bound on how many selections are appropriate.²⁷

These same distinctions can be applied to the questions left open by specificational subjects. Consider how the question-answer analysis would describe the following specificational sentences:

- (75) a. The Roman emperor who invaded Britain was Claudius.
 b. The Roman emperors who invaded Britain were Claudius and Septimius Severus.
- (76) a. One Roman emperor who invaded Britain was Claudius.
 b. Some Roman emperors who invaded Britain were Claudius and Septimius Severus.

Like the specificational sentences we looked at above, their subjects leave open a question to which their complements provide an answer. But the questions left open by the definite descriptions in (75) differ from those left open by the indefinite descriptions in (76), in the same way that the questions in (73) differ from those in (74): they request an exhaustive set of alternatives. At the same time, the questions left open by the singular subjects in the (75-a) and (76-a) differ from the questions left open by the plural subjects in (75-b) and (76-b). They request a single alternative as opposed to several.

The complements of these specificational sentences make selections which purport to satisfy these different requests. For that reason, speak-

²⁷I should point out that this way of describing the questions in (73)–(74) is not compatible with every semantic analysis of questions. In particular, it is incompatible with analyses that treat the alternatives presented by any question as mutually exclusive and jointly exhaustive, as in the partition semantics for questions developed by Groenendijk and Stokhof (1997). But analyses which allow more than one of the alternatives presented by a question to be true can represent these questions as presenting a common set of alternatives, and differing only in their requests. Such an analysis is developed by Belnap and Steel (1976), from whom the terminology of ‘requests’ derives.

ers who assert these sentences must know different things, and they convey different information to their audiences. A speaker who asserts (75-a) must know that Claudius was the *only* Roman emperor who invaded Britain, because that is what her assertion conveys to her audience. But a speaker who asserts (76-a) does not need to know this, and claims only that Claudius was a Roman emperor who invaded Britain, without excluding the possibility that others did the same. Similarly, a speaker who asserts (76-b) conveys that several Roman emperors invaded Britain, but without claiming to have provided an exhaustive list of them. Thus, by appealing to the concept of a question's request, the question-answer analysis can accommodate specificational sentences whose subjects are indefinite descriptions.

By distinguishing a question's request from the alternatives it presents, we can also see what made the inversion and equative analyses tempting. The questions expressed by (73)–(74) and the subjects of (75)–(76) all present a set of alternatives with a common form, namely:

x is a Roman emperor \wedge x invaded Britain

The presented alternatives are those claims which are instances of this logical form, the propositions which result from giving an appropriate value to x . In other words, we may think of these questions as presenting a set of alternatives by displaying what each alternative must claim to be *true of* some individual. This is an open proposition of type $\langle e, t \rangle$ —that is, a predicate of exactly the sort that the inversion analysis would identify as expressed by the subjects in (75)–(76). Seen through the lens of the question-answer analysis, the kernel of truth in the inversion analysis is that it identifies the predicate which determines the alternatives presented and selected among in a specificational sentence.

We can also see more clearly why an equative analysis of a sentence like (75-a) looks appealing, but ultimately comes up short. Recall that the equative analysis explains the symmetry of subject matter in specificational sentences that have definite description subjects by assigning the subject and complement in such sentences the same denotation. This is a plausible strategy when the question left open by the subject is one that requests a unique true alternative, like the question in (73-a). In that case, it is possible to take both the question and the answer as concerned with exactly one alternative, assign this alternative as their common denotation, and leave the difference between them to be explained in terms

of some other concept, like that of information packaging, discriminability, or Frege's notion of *Sinn*. This strategy seems less plausible, though, once we recognize that some questions permit non-exhaustive selections, and selections of multiple alternatives.²⁸ The equative analysis pays too little attention to cases like (76-b), where symmetry of subject matter is not easily construed as concern for just one particular alternative.

So in addition to accounting for their symmetry of subject matter and asymmetry of role, the question-answer analysis affords us a wider perspective on specificational sentences. From this perspective, we can account for some specificational sentences which previous literature has often ignored, namely, specificational sentences whose subjects are indefinite and plural descriptions. We can also accommodate the insights of the other two analyses, seeing what is right about each. I conclude that the question-answer analysis is an excellent candidate for a general, unified account of specification. Going forward, I shall take it that the question-answer analysis as I have described it here is broadly correct.

To achieve this wider perspective, we had to distinguish between the alternatives presented by a question and the request that it makes about how to select among those alternatives. The examples in (75)–(76) suggest that this distinction is closely related to the two parts of a description when it appears as a specificational subject: the determiner or quantifier is related to the request, and the noun phrase it modifies is related to the alternatives presented. The next chapter will examine this relationship in more depth, offering an account of how these two parts of a description function which is in line with the basic insight I have argued for here: a description is interpreted as expressing, or leaving open, a question.

²⁸For example, to analyze specificational sentences which make multiple selections, such as (75-b), proponents of the equative analysis like Heycock and Kroch (1999) typically follow Link (1983), treating both the plural description 'the Roman emperors who invaded Britain' and the entire complement phrase as denoting a single plural entity, namely, the entity which is the sum of Claudius and Septimius Severus. This seems at best non-perspicuous, and at worst like it gives the wrong truth conditions: if (75-b) is true, then Britain was invaded by two Roman emperors, not one emperor-sum.

Part II

Investigation

Chapter 4

Investigation

I have argued that a specificational sentence pairs a question with its answer. Its subject is interpreted as presenting a set of alternatives, while its complement is interpreted as selecting among those alternatives. But this conclusion raises a puzzle, which I now want to address in more detail.

We have seen that *descriptions* are one important class of specificational subjects.¹ Here is one example:

- (1) Two of the British students are Audrey and Bert.

The puzzle stems from the fact that descriptions like ‘two of the British students’ could in most contexts be analyzed using the resources of ordinary first-order logic. Outside of specificational sentences, this description occurs in nominal position:

- (2) a. John met two of the British students.
b. Two of the British students won medals in the competition.

The contribution of the description to the truth-conditions of such sentences can be given by a first-order schema like this:

- (3) $\exists x \exists y (Sx \wedge Sy \wedge x \neq y \wedge \phi(x, y))$

Here, Sx translates ‘ x is a British student’, while $\phi(x, y)$ schematically stands for a translation of the rest of the sentence in which the description occurs. In the translation of (2-a), for example, $\phi(x, y)$ will encode ‘John

¹See the discussion in Section 2.2.

met x and John met y '. Thus, all we need to adequately represent the logical contribution of the description 'two of the British students' are a pair of existential quantifiers, the predicate S , a relation of identity, and the usual sentential connectives.

The puzzle is, if such descriptions normally just express first-order quantification over individuals, why think they express questions in the subject of a specificational sentence? To make this sharp, consider that we might just analyze (1) as an instance of (3), as follows:

$$(4) \quad \exists x \exists y (Sx \wedge Sy \wedge x \neq y \wedge x = a \wedge y = b)$$

Where in this formula can we find the question presented by 'two of the British students'? It doesn't seem to be there anywhere: the formula just asserts the existence of two distinct British students who satisfy a certain open formula, namely, $x = a \wedge y = b$. But if that is an adequate representation of the truth conditions of (1), where do questions and answers enter the picture? On the basis of this example, it might seem that the question-answer analysis is an unnecessarily complicated strategy for representing the semantics of specification. Theoretical simplicity, someone might say, demands that we stick with the first-order analysis if we can.

The puzzle is only puzzling because it assumes that the question-answer analysis and the first-order analysis are incompatible alternatives. This assumption is natural given a certain understanding of first-order logic, and in particular first-order quantification. Philosophers are used to thinking of first-order quantifiers as generalizing over a domain of objects, and we are not used to thinking of them as having anything to do with asking and answering questions. If we hold that descriptions express first-order quantification over a domain of individuals, there seems to be little room to say that they present questions, and vice versa.

But that is not the only way of understanding first-order quantification, and rather than assuming it is correct, I want to solve the puzzle by arguing that first-order analyses and question-answer analyses of specificational sentences are compatible. To do so, I will articulate a conception of first-order quantification that aligns better with the question-answer analysis than the usual referential conception does. This 'investigatory' conception of quantification is interesting because it illuminates the special role that specificational sentences play in our language. It also has

interesting consequences for the claim that numbers are objects, which I will discuss in the next chapter.

Roughly, the idea behind the investigatory conception is that the meaning of first-order quantifiers is importantly connected with practices of asking and answering questions—practices of inquiry, or as I shall call it, *investigation*. An investigation is an activity with a structure imposed by a question. It is undertaken, or performed, or carried out, for the sake of answering a certain question; finding the answer to that question is the end, or aim, of the activity.

This may sound rather esoteric and abstract, but I mean to be talking about a process that is humdrum and familiar. A clear model of the kind of activity I am talking about is the act of solving an elementary algebra problem. Consider, for example, a simple polynomial equation of the form:

$$x^2 - 6x + 9 = 0 \tag{4.1}$$

In a beginning algebra course, an equation like (4.1) is typically used to give a problem or exercise. To solve this problem, a student might proceed by writing:

$$(x - 3)(x - 3) = 0 \tag{4.2}$$

$$x - 3 = 0 \tag{4.3}$$

$$x = 3 \tag{4.4}$$

The process of moving from the first equation to the last is an investigation. The first equation gives a problem because it invokes an implicit question: which number makes the equation true? The last equation answers this question, by giving such a number as the value of x . This solves the problem and concludes the investigation. The intervening lines represent the steps the student takes with the aim of finding that answer.

The relevance of this example becomes clear when we recognize that (4.4) functions as a specificational sentence in the derivation above. Unlike the other equations in that derivation, it gives, or specifies the value of x ; it says which number makes equation (4.1) true. Indeed, we might translate what this equation says into English by means of a specificational sentence like

- (5) The number x such that $x^2 - 6x + 9 = 0$ is three.

This suggests a hypothesis: we can understand specificational sentences in terms of the role they play in practices of investigation. A sentence counts as specificational because it plays this role, because it gives the value of a variable and thereby completes an investigation.

I am not the first to suggest this hypothesis. When linguists give an intuitive characterization of the category of specificational sentences in natural language, they often spontaneously resort to a comparison with algebraic equations like (4.4). Mikkelsen (2011, pp. 1809–1810) says, for example, that in a specificational sentence, “the subject phrase introduces a variable... and the predicate complement provides the value for that variable”. Similar ideas appear in Akmajian (1970), Higgins (1979), and Heycock and Kroch (1999). In these discussions, though, it is simply taken for granted that we know what it means to ‘give the value of a variable’. But this is far from clear: why does (4.4) count as giving the value of x , for example, while (4.2) does not? I aim to fill that gap here, by describing ‘giving the value of a variable’ as a role in practices of investigation, and stating criteria for a statement to play that role.

Here, then, is the plan. I will begin by describing the structure of investigations in more detail, and asking what it means to ‘give the value of a variable’, using elementary algebra problems as a model. I will then show how to give a semantics for first-order languages which is stated in terms of this practice. This semantics will vindicate the idea that specificational sentences occupy a special role in practices of asking and answering questions, and will solve the puzzle I laid out above: since first-order languages can be understood in terms of investigations, and investigations are structured by questions, the first-order and question-answer analyses of a specificational sentence are compatible. I will conclude by returning to specification in natural language, showing that by thinking of specificational sentences in terms of their role in practices of investigation, we can better understand some of their features.

4.1 THE STRUCTURE OF INVESTIGATIONS

4.1.1 *Two perspectives on investigations*

An investigation is an activity with a certain teleological and epistemological structure. This structure is imposed by a question. At the start of an investigation, the question is understood, but its answer is not known.

By the end, the answer has become known. To move from the one state to the other, one or more investigators must usually take a series of steps. These steps make progress toward the discovery of the answer. The specific nature of these steps depends on the subject matter of the question: when the question is an algebra problem, the intervening steps will involve manipulating equations; when the question is about the physical world, they might involve running an experiment; when the question is what I ate for breakfast, they will involve consulting my memory or the dishes in my sink. But in all cases, the basic structure of the activity is the same: once the question is understood, the investigators use suitable means to try to find an answer. They conclude when they have succeeded in this aim.

One way to look at an investigation, then, is as an attempt to answer a question; but the example of solving algebra problems shows that there is sometimes another perspective we can adopt. As the student solves the problem above, she is attempting to find a certain *number*. The equation on line (4.4) solves the problem because it gives the number which satisfies equation (4.1), namely, the number 3. The process of solving an algebra problem is thus not just an attempt to answer a certain question; it is a process of finding the value of a variable which ranges over a certain class of numbers.

Understanding the relationship between these two perspectives is important, because they are intimately linked. It is helpful in this connection to draw a distinction between two types of questions which may structure an investigation. As noted in Chapter 3, literature on the logic of questions characterizes questions as presenting a class of alternatives, or possible answers; the distinction concerns how they do so. Belnap and Steel (1976) call the two types of questions *whether*-questions and *which*-questions, and characterize them as follows. A *whether*-question presents an explicit, finite list of alternatives. Here are some examples:

- (6) Is Alec in Germany?
 - a. (Yes,) Alec is in Germany.
 - b. (No,) Alec is not in Germany.
- (7) Is Alec in Germany, or France?
 - a. Alec is in Germany.
 - b. Alec is in France.

I mention this class only to distinguish it from which-questions and put it aside. Although a whether-question can still structure an investigation, it is not usually helpful to view such an investigation from the second perspective, as an attempt to find and give the value of a variable.

Unlike a whether-question, a which-question presents an indefinitely large, perhaps infinite, class of alternatives. In general, the particular alternatives cannot be explicitly listed. Instead, a which-question presents its alternatives by giving their common form, which Belnap and Steel (1976) call the *matrix* of the question. The matrix is a statement containing an ‘unknown’. In natural language, such an unknown is expressed by a ‘wh’-word; in formal analyses, we represent it with an algebraic variable. Because it contains a variable, the matrix gives the form of an unknown truth. What’s unknown is which value, or values, of the variable will make the matrix true. A which-question thus asks, which statements of this form are true?

An elementary algebra problem is a prototypical which-question. Its matrix is the polynomial equation used to state the problem. The alternatives presented by the question are the statements that result from assigning a value to the variable in the matrix. For example, the alternatives presented by (4.1) are:

$$\text{that } 0^2 - 6 \cdot 0 + 9 = 0$$

$$\text{that } 1^2 - 6 \cdot 1 + 9 = 0$$

$$\text{that } 2^2 - 6 \cdot 2 + 9 = 0$$

and so on. Similarly, most questions expressed by ‘wh’-words (except ‘whether’) in natural language are which-questions, and their alternatives can be represented using a matrix containing an algebraic variable. Here are some examples of such questions, together with some examples of the alternatives they present and a matrix giving their common form:

- (8) Who is in Germany?
- a. x is in Germany.
 - b. {Alec is in Germany, Fiedler is in Germany, Liz is in Germany, ... }
- (9) Where is Alec?
- a. Alec is in x .

- b. {Alec is in Germany, Alec is in France, Alec is in Spain, ... }

Because a which-question presents its alternatives by giving their general form, rather than listing particular alternatives explicitly, an investigation structured by a which-question can be viewed from the second perspective: it is an attempt to find and give the value (or values) of a variable. One can give the answer to such a question *by* giving the values of a variable. To answer a question is to claim that some of the alternatives it presents are true. Giving values for the variable in a which-question suffices to determine such a claim, because the question itself gives its general form; to answer the question, all that's needed is to indicate which particular instances of this general form are true. Thus, saying "George" in response to (8) suffices to claim that George is in Germany, that the person in Germany is George. Similarly, writing down the equation $x = 3$ in the context of the derivation (4.1)–(4.4) suffices to claim that $3^2 - 6 \cdot 3 + 9 = 0$, that the number x such that $x^2 - 6 \cdot x + 9 = 0$ is 3.

Thus, we may look at an investigation structured by a which-question from two perspectives, or levels: it is both an attempt to find the answer to a certain question, and an attempt to find the value of a variable. Giving the value of a variable is a *means* of answering the question. This duality of perspectives will be important later on, because it is at the center of the solution to the puzzle I described above: it is because we can see an investigation both as a search for the value of a variable and as an attempt to answer a which-question that the first-order and question-answer analyses are compatible, and even mutually supporting. But before we get there, it is necessary to reflect in more detail on what it means to 'give', or specify, the value of a variable, since that is the crucial step in answering a which-question and completing the investigation which it structures. This will further explicate the hypothesis that a specificational sentence gives the value of a variable and therefore plays a distinctive role in a practice of investigation.

4.1.2 *Giving the value of a variable*

Let's turn our attention, then, to the language and practice of elementary algebra, since that is where our standards for what counts as giving the value of a variable are sharpest. Consider again the derivation in (4.1)–

(4.4).

$$x^2 - 6x + 9 = 0 \quad (4.1)$$

$$(x - 3)(x - 3) = 0 \quad (4.2)$$

$$x - 3 = 0 \quad (4.3)$$

$$x = 3 \quad (4.4)$$

Notice that (4.4) is the *only* line in this derivation that counts as giving the value of x , and answering the problem posed by (4.1). A student who stops at line (4.2) or line (4.3) at best provides a partial answer, and a student who merely copies down (4.1) does not give any answer at all. On the other hand, once a student derives the equation on line (4.4), no more work can be done: this equation fully answers the question, and any further manipulations would carry the student further away from a complete answer to the question.

The status of this equation with respect to the problem posed by (4.1) is therefore somewhat special, because it alone gives the value of x . The other equations in the derivation are like the one used to state the problem, in that they do not *give* the value of x , though they do describe or constrain it. What confers this special status on (4.4)? Why does it count as giving the value of a variable and concluding the investigation, while the other equations in the derivation do not? This is the question we have to answer if we want to understand specificational sentences by analogy with such equations.

As a preliminary step toward answering this question, we should notice that in the derivation in (4.1)–(4.4), each line is materially equivalent to the others. Each is true if, and only if, the value of x is three. Thus, to understand why only the final equation counts as giving the value of x , we need to look beyond its material truth conditions. What else, then, could explain its unique meaning and role?

There are four different features of the equation in this final line which are crucial for it to count as giving the value of x and as answering the question structuring an investigation. These are: that the variable is in the same *scope* as the variable introduced by the problem statement; that the value given is in the *range* of the variable; that the equation represents a *complete* solution; and that the equation is in a *canonical form*. These are not features of the expressions in the equation themselves; they are features of how those expressions are used and understood in this particu-

lar context. Each of them is individually necessary for this equation to be understood as giving the value of x . Together, they provide a reasonably complete account of why this equation gives a solution to the problem, but the earlier equations do not. Let me describe each in turn.

It is easiest to see what I mean by ‘scope’ if we ask: why does the equation $x = 3$, as it is used on line (4.4), give the solution to the problem given on line (4.1), as opposed to some other problem? The obvious answer is that the equations on these two lines of the derivation employ the same variable. But what is it for the variable to be ‘the same’ in these two lines? This is obviously not a purely typographical matter. The *symbol* ‘ x ’ can be used to express a variable in many different problems. I might introduce a new problem to you with an equation like

$$x^2 - 4x + 4 = 0 \tag{4.5}$$

without causing any confusion. Students of elementary algebra have to learn that the symbol ‘ x ’ is being used differently here than in (4.1), and that its use in (4.4) is connected with its use in (4.1) but not in (4.5). This is what I mean by scope: a variable’s scope includes just those statements in which the expression for the variable is used in the same way.²

In the language of elementary algebra, the scope of a variable always extends exactly as far as the statements belonging to a single investigation, and no further. This is because an algebraic variable originates in the matrix of a which-question. To understand a statement containing a variable requires recognizing it as part of an attempt to answer that question by finding the variable’s value. A variable’s scope thus opens with a statement of a problem, using one or more equations; it closes with a statement of its solution, which gives the value (or values) of that variable; and it includes whatever statements occur between the problem statement and the solution statement that are part of the problem-solver’s effort to move from the former to the latter.

²I avoid saying that a variable’s scope includes just those statements in which it has the same value or denotation. This is true in a sense, but potentially misleading. One can understand that two occurrences of ‘ x ’ are within the same scope without knowing their value or denotation. This is just the position that a student is in, for example, when she has derived an equation on the way to a solution, but without having arrived at the solution itself yet. It would be misleading to describe this student as knowing that the occurrences of ‘ x ’ in her partial derivation have the same value, or stand for the same thing, because this invites the inference that she knows what value they stand for—which is something she won’t know until she solves the problem.

Furthermore, there is no sense in asking for the value of a variable apart from some scope or other. A variable only has a value relative to the question which introduces it. To say that an equation gives the value of a variable is therefore to connect that equation with an attempt to answer a particular problem or question. It is to place the use of that equation in the context of some investigation or other. For an equation to count as giving the value of a variable, the variable in that equation must occur in the same scope as in the problem to which the equation gives the solution.

Of course, the problem which introduces a variable must actually *have* a solution if its value is to be given. This depends on the second feature of how the variable is understood in the context of the problem: its range. Take, for example, the problem given by the following equation:

$$x^2 + 5x + 6 = 0 \quad (4.6)$$

If x is taken to range over the natural numbers in this problem, then it has no solution, as you can see once it has been factored:

$$(x + 2)(x + 3) = 0 \quad (4.7)$$

The roots of (4.6) are -2 and -3 . They are both negative integers, so they are not among the natural numbers. So whether or not this problem has solutions depends on whether x ranges over just the natural numbers, or over some more inclusive set, like the integers. Thus, it is only possible to give a value for x in this case when x is understood as ranging over some set that includes at least -2 or -3 . For an equation to count as giving the value of a variable, that variable must be understood as having a certain range, and the value given must be within that range.

This example raises another important point. Because it gives the value of x , the equation (4.4) is a *complete* solution to the problem in (4.1). In that problem, the equation only has one root, and so giving that root necessarily means completely solving the problem: after that value has been given, no more work remains to be done. But many polynomial equations, unlike (4.1), have more than one root (at least on some ways of understanding the range of the variable). For this reason, it is in general too simple to ask under what conditions a statement gives *the* value of a variable. When x is understood to range over the integers in (4.6), an equation like

$$x = -2$$

could be said to give ‘a’ value of x , but not ‘the’ value of x , since no value in that range uniquely satisfies the equation. For an algebraic statement to count as a solution statement, it must give at least *one* value for the problem’s variable. But in general, there is a further question as to *how many* such values must be given before a complete solution which marks the end of an investigation has been obtained.

For equations with multiple distinct roots, we need to distinguish among several possible understandings of the problem in order to answer this question. Because it has multiple roots, an equation like (4.6) actually underdetermines which question is being asked. If the question is “What are *all* the numbers x such that $x^2 + 5x + 6 = 0$?”, then a complete solution must give both roots. A student solving this problem would have to write something like

$$x = -2 \text{ or } x = -3 \quad (4.8)$$

to completely answer this question and achieve the end of her investigation. On the other hand, if the question is “What is *one* number x such that $x^2 + 5x + 6 = 0$?” then

$$x = -3 \quad (4.9)$$

gives a complete solution to the problem. Even though it does not exhaustively list the roots of the equation, the problem-solver can stop at this point, since this equation provides everything the question asks for.

Belnap and Steel’s terminology (which I introduced in Section 3.3.4) is helpful for understanding what is going on here. As I noted above, an equation like (4.6) is the matrix of a which-question: it gives the general form of the alternatives presented by that question. But there is more to a question than just the alternatives it presents. Two questions may present the same set of alternatives, but make different *requests*. A question’s request determines how many alternatives an answer should select, and whether a correct answer should provide an exhaustive list of them. In general, an answer is not complete unless it fulfills this request.

The different understandings of the problem in (4.6) correspond to two different requests. We can express these requests in natural language using two different quantifiers, as the difference between “What is *one* number x ...?” and “What are *all* the numbers x ...?” In the first case, we are interpreting the problem as requesting just one true alternative, and hence just one value of x , so (4.9) counts as a complete solution statement.

In the second case, we are interpreting the problem as requesting an exhaustive set of true alternatives, and hence an exhaustive set of roots. Since there is more than one true alternative, a statement like (4.8) is required to complete the investigation structured by this problem.

This shows that we should distinguish between giving a value of a variable and giving the solution to a problem. When the problem requests more than one true alternative, it is necessary to give more than one value for the variable to give its solution and complete the investigation. The same distinction is needed when we look at the more general class of algebra problems given by systems of equations in multiple variables, such as the problem given by (4.10) and (4.11):

$$x^2 + y^2 = 25 \quad (4.10)$$

$$x - y = 1 \quad (4.11)$$

Here is one solution to this problem:

$$x = 4, y = 3 \quad (4.12)$$

This solution statement contains two equations as parts, each of which gives a different value for a different variable. Giving a value for a variable is necessary, but not sufficient, for completing an investigation. An investigation is only complete when we have given as many values as requested for each of the variables introduced by the problem.

The final criterion concerns which kinds of *expressions* are suitable for giving the value of a variable. I call these criteria of 'canonical form'. When an equation is suitable for giving the value of a variable, it must also be in canonical form. To see the importance of canonical form, consider again the final two lines from the derivation in (4.1)–(4.4):

$$x - 3 = 0$$

$$x = 3$$

As we noted earlier, these equations are materially equivalent. What then is the difference between them? Why is it that only the second of these equations counts as giving the value of x , and thereby giving the solution to the problem in (4.1)?

The point I want to emphasize here is that some standard or other is always required in a problem-solving practice to distinguish among

equations like these. Part of understanding a problem is grasping the criteria that distinguish those statements which can be used to give values for the problem's variables from those which, though otherwise equivalent, cannot. In the absence of such criteria, we have no grounds for distinguishing equations that are used to give problems from those that are used to give solutions, and so no possibility of a problem-solving activity at all. Thus, for an equation to count as giving a value of a variable, some such criteria must be in place, and the equation must fulfill them.

What are those criteria? Looking at the two equations above, a natural first thought is that the second equation is syntactically and arithmetically simpler than the first. This can't be the whole story about canonical form, though, because there are certain pairs of equations in which one equation is no simpler than the other, but only one counts as giving the value of x . Here are two examples:

$$\begin{array}{ll} 2x = 1 & -x = 2 \\ x = \frac{1}{2} & x = -2 \end{array}$$

At least in the classrooms where I was taught, only the second equation in each pair would have counted as a complete solution to a problem, even though it is no syntactically or arithmetically simpler than the first.

Usually, we can say that for an equation to be in canonical form, the variable should be arithmetically isolated, and purely numerical expressions should be fully computed. But even these generalizations are not exceptionless, and apart from them, whether an expression counts as being in canonical form is a fairly local matter. Consider, for example, the problem given by

$$18x = 8$$

Here are two possible ways of writing a solution to this problem:

$$x = \frac{4}{9} \qquad x = 0.\bar{4}$$

Neither of these equations is 'more canonical' than the other in any global sense. A teacher who assigns this problem as an exercise might accept one equation as a solution but not the other depending on his pedagogical purposes. For example, he might accept the second as a complete solution, but the first as only a partial solution, in a lesson emphasizing

decimal expansions of rational numbers. In a lesson emphasizing exact expressions of ratios, or where calculators are not permitted, the standard might be the other way around.

Criteria of canonical form, then, are inherently local and purpose-relative, even if there are some general rules that apply in most problem-solving settings. So what, if anything, unifies the various standards we might have? Is there any general account of what it is for an expression to be in canonical form?

Here is my answer: a standard of canonical form determines a distinguished class of expressions such that, when one of these expressions is used to give the value of a variable, no further question can be appropriate about *which* value is meant. This becomes clearer if we imagine an algebraic investigation formulated as an explicit question and a series of attempted answers:

$$(10) \quad \text{What is the number } x \text{ such that } 4x^2 - 4x + 1 = 0?$$

If you were to ask someone this question, it would obviously be unsatisfying to be told, in reply, that

$$(11) \quad \text{It is the number } x \text{ such that } 4x^2 - 4x + 1 = 0.$$

even though this answer is in a sense perfectly true. This ‘answer’ merely repeats the problem; it is unsatisfying because it invites the rejoinder, “But *which* number is that? That’s what I wanted to know.” It would only be slightly more satisfying if the reply was

$$(12) \quad \text{It is the number } x \text{ such that } (2x - 1)^2 = 0.$$

Here again, it can be appropriate to ask “But which number is that?”, and insofar as it is, the respondent has not fully answered the original question. But this question is no longer appropriate if the respondent says:

$$(13) \quad \text{It is the number } x \text{ such that } x = \frac{1}{2}.$$

A further ‘Which number is that?’ question in this case would indicate a misunderstanding on the part of the questioner. Our practices of giving numbers simply do not allow any further request here; the expression ‘ $x = \frac{1}{2}$ ’ makes it as clear which number is meant as anyone has any right

to demand. This equation, unlike those in the other answers, is in canonical form. Its being in canonical form just consists in the fact that no such demand is appropriate. The criteria for when such a demand is or is not appropriate may vary from one problem-solving context to the next. But the important point is that we do in fact have such criteria, and some criteria or other are always in play during an algebraic investigation, because it is by applying such criteria that we recognize that the problem has been solved, and the investigation is concluded.³

One further distinction concerning canonical form will be helpful later on. I have spoken of canonical form in terms of standards that apply to statements, or sentences. Very often, though, these sentence-level standards will be articulated into standards or rules concerning their subsentential parts. For example, in most classrooms,

$$x = \frac{2}{4}$$

is not in canonical form, while

$$x = \frac{1}{2}$$

is. But the only difference between these equations concerns the expression on the right-hand side, and it is most natural to say that the former is not in canonical form because the fraction is not fully reduced. The former equation fails to meet the standards of canonical form because it fails to present a canonical *expression* for the value of x . I will reserve ‘canonical form’ for standards governing complete sentences, and ‘canonical

³Dummett (1981, pp. 67–69) puts a similar thought to a slightly different purpose. As part of a series of criteria designed to distinguish genuine from spurious singular terms, he proposes that we can use such repeated ‘But which do you mean?’ questions to distinguish expressions of first-level generality from those of higher-level generality. The details of his proposal are not important here; the important point is just that Dummett also recognizes that such “requests for specification” are sometimes grammatically well-formed, but inappropriate because they express a misunderstanding. In such cases we must therefore implicitly grasp the criteria I am here calling criteria of canonical form: criteria by which we recognize that a request for specification has been completely fulfilled, and can’t be improved upon. Our grasp of such criteria also underlies the ‘directness’ test I proposed in Chapter 2 for distinguishing specificational from non-specificational readings of copular sentences.

expression' for standards governing subsentential expressions, particularly those that stand for values in the range of a variable. Our standards about which expressions for values are canonical are included in our standards of canonical form. But the two are in general distinct, since we might have standards of canonical form that go beyond the rules about which expressions a sentence may contain. For example, standards of canonical form might include rules about the order in which expressions should appear (such as 'Write the variable on the left').

Here is a summary of the above observations. Algebraic practice is a practice of investigation, of posing questions or problems and seeking their solutions. When we say that an equation like

$$x = 3$$

'gives the value of a variable', we are assigning it a certain role within that practice. To say that an equation gives a value for a variable is to say that it is part of a solution statement, which concludes an investigation by fulfilling its aim. It determines the answer to the question structuring the investigation, and so gives what was sought. This distinguishes it from equations which state problems, or mark intermediate steps between a problem and a solution, which do not give a value for the variable and cannot be part of a solution statement.

In order for an equation to play that role, it must be understood in a certain way. Its variable must be in the same scope as the variable introduced by the problem statement, and it must give a value within the range that the variable is understood to have. To count as *giving* a value, rather than just describing it, the equation must satisfy some criteria for being in canonical form, using an expression for the value which is not subject to a 'But which do you mean?' question by the standards of the problem-solving practice. When a polynomial used to give a problem has more than one root, or contains more than one variable, more than one such equation might be necessary to solve the problem and complete the investigation. How many such equations are necessary is a matter of how the question structuring the investigation is understood: how many variables its matrix contains, how many true alternatives it requests, and whether or not it requests an exhaustive list of true alternatives.

4.2 INVESTIGATORY SEMANTICS

4.2.1 *From investigations to truth conditions*

In the previous section, I described investigations as practices or activities which aim at answering a question. In the case of which-questions, they can be viewed as attempts to find and give the value of one or more variables. I then described four criteria for a statement to count as giving the value of a variable, and thus playing the role of concluding an investigation. I have not yet said anything about how investigations relate to truth, though. In order to solve the puzzle with which I began, we need to connect first-order analyses of the truth conditions of specificational sentences with practices of asking and answering questions—that is, investigations.

We can state the truth-conditional contributions of definite, indefinite, and numerically-quantified descriptions schematically using first-order quantifiers. When those descriptions appear as subjects of specificational sentences, we can analyze those sentences as instances of those schemas. On this strategy of analysis, a specificational subject will always introduce one or more existential quantifiers as the outermost connectives, and its complement will always translate into a conjunction of identity statements, in which an existentially-bound variables appears on one side, and a constant appears on the other. Thus, in the simplest case, a specificational sentence with an indefinite subject like

(14) A good book to read is *Waverley*.

has a first-order analysis like

$$\exists x(B(x) \wedge x = w)$$

where B is a predicate true of good books, and w is an individual constant denoting *Waverley*. Definite descriptions can be accommodated via Russell's theory of descriptions, simply by adding a uniqueness condition:

(15) The author is Scott.

may be analyzed as

$$\exists x(A(x) \wedge \forall y(A(y) \rightarrow y = x) \wedge x = s)$$

If, like many theorists, we are willing to analyze free relatives as definite descriptions, such that for example ‘where ...’ is synonymous with ‘the place ...’ and ‘when ...’ is synonymous with ‘the time ...’, we can extend this analysis to specificational pseudo-clefts. Finally, numerically-quantified descriptions can be analyzed by introducing further existential quantifiers and adding distinctness conditions:

- (16) Two of the characters are Edward Waverley and Baron Bradwardine.

may be analyzed as

$$\exists x \exists y (C(x) \wedge C(y) \wedge x \neq y \wedge x = e \wedge y = b)$$

Thus, this strategy of analysis can apparently accommodate the vast majority of specificational subjects surveyed in Chapter 2.

There are many refinements we might want to make on this basic semantic program. For example, we might want to distinguish the asserted content in a specificational sentence from what it presupposes or implicates, or require that the first-order representations be derived compositionally from the syntax and basic lexicon of their natural language counterparts. Even with such refinements, though, the basic strategy will remain the same: the goal is to analyze the meaning of specificational sentences by representing their content in a formal language equipped with the resources of (at least) first-order logic. My question is whether this *kind* of analysis is compatible with the question-answer analysis which I argued for in Chapter 3. So I will leave the refinements aside, and assume, for the moment, that the strategy of analysis just outlined is basically correct, and provides an accurate representation of the truth conditions of specificational sentences like (14), (15), and (16).

The trouble is that the usual semantics for first-order languages makes it difficult to see what connection there could be between such analyses and practices of investigation. According to the usual semantics, an existentially quantified statement is true just in case some object satisfies the formula to which the quantifier attaches. The truth conditions of first-order existential statements are usually described using a clause like this one:

$\exists x \phi(x)$ is true in a model M under assignment g iff there is an

object a in the domain of M such that $\phi(x)$ is true in M under $g^{x \rightarrow a}$, where $g^{x \rightarrow a}$ is just like g except that it maps x to a .

There is, apparently, no reference to asking and answering questions, or ‘finding’ or ‘giving’ the value of a variable, in this explanation. This is what gives rise to the puzzle. We saw in the last section how to connect answering questions with finding and giving the value of a variable; but we have not yet connected giving the value of a variable with first-order quantification.

My response, therefore, is to describe an alternative semantics for first-order languages which explicitly draws that connection. Fortunately, such a semantics is ready to hand, in the form of the *game-theoretical semantics* for first-order logic developed by Jaakko Hintikka, in work that begins with Hintikka (1973).⁴ This semantics is formally equivalent to the usual semantics, so it raises no new questions about the adequacy of first-order analyses: a first-order sentence used to present the truth conditions of an English specificational sentence will be true in all the same models under both the usual and the game-theoretical semantics. But under the game-theoretical semantics, quantification is described in terms of practices of investigation, of finding and giving individual objects as the values of variables. I will therefore present the game-theoretical semantics, and show how, under this semantics, the apparent tension between the question-answer and first-order analyses of specificational sentences disappears.

4.2.2 *Game-theoretical semantics*

I will first present the game semantics in its original setting, as a definition of truth for a first-order language. In the next section, I will turn to showing that the question-answer analysis of specification is compatible with a first-order analysis, when we think of the first-order quantifiers as they are interpreted by the game semantics.

The game semantics defines truth for sentences in a first-order language L in terms of a certain game played using sentences of the language. Here is how the game is played. We need two players, an initial sentence S of L , and an interpretation or model M of the non-logical constants in L . The two players are called the Verifier and the Falsifier. They

⁴For a recent overview, see Hintikka and Sandu (2011).

have opposing goals: the Verifier is trying to produce a sentence which is verified by M , and the Falsifier is trying to produce a sentence which is falsified by M . For simplicity, we shall assume that L contains only the propositional connectives \neg , \wedge , and \vee , and the quantifiers \exists and \forall . (Thus, every connective in L except negation has a dual.)

The play proceeds in rounds. Each round destructures the initial sentence S into a new sentence. The game stops, and one of the two players wins, when a closed atomic sentence is reached. Before that point, the move made in each round is determined by the main connective of S , in accordance with the following rules (Hintikka, 1974, p. 156):

- G. \exists .** If S is of the form $\exists xS'$, the move is made by the Verifier. She goes to the domain of individuals in M , chooses an individual, and gives it a name ' a '. The name is substituted for ' x ' in S' . The resulting sentence is called ' $S'(a/x)$ '. The game continues in the next round with respect to $S'(a/x)$.
- G. \forall .** If S is of the form $\forall xS'$, the same thing happens as in G. \exists , except that the individual is chosen and named by the Falsifier.
- G. \vee .** If S is of the form $(S' \vee S'')$, the Verifier chooses one of the disjuncts S' or S'' , and the game continues with respect to that disjunct.
- G. \wedge .** Similarly, if S is of the form $(S' \wedge S'')$, the Falsifier chooses one of the conjuncts, and the game continues with respect to that conjunct.
- G. \neg .** If S is of the form $\neg S'$, the two players switch roles, and the game continues with respect to S' .

These rules define the game $G(S, M)$. Here are the conditions for winning $G(S, M)$. After the last move, which produces a closed atomic sentence S' , we consult the model M . If S' is true in M , then the player who was initially the Verifier wins; otherwise, the initial Falsifier wins. Thus, we are assuming a definition of truth-in-a-model for *atomic* sentences, in order to state the winning conditions for the game. The usual denotational semantics will do: ' $F(a)$ ' is true in M iff the object denoted by ' a ' in M is in the extension of ' F ' in M ; ' $R(a, b)$ ' is true in M iff the pair of objects denoted by ' a ' and ' b ' is in the extension of ' R ' in M ; and so on.

We use the game to obtain a definition of truth-in-a-model for all sentences of the language L :

Game-theoretical truth. A sentence S of L is true in a model M if and only if the Verifier has a *winning strategy* in the game played with respect to S and M .

That is, S is true in M if and only if, at every round, the Verifier can choose her moves such that the game will end with some atomic sentence or other which is true in M , no matter what the Falsifier does on other rounds.

This definition has a number of important features that we might expect for a definition of first-order truth. The semantics is classical: the Verifier has a winning strategy in exactly one of $G(S, M)$ and $G(\neg S, M)$, so exactly one of S and $\neg S$ is true in M .⁵ More importantly, the game semantics is in a certain sense conservative. It can be proven that game-theoretical truth agrees with the standard Tarski-style definition of truth for a first-order language. That is, any sentence in a first-order language is true in a model M according to one semantics if and only if it is true in M according to the other.⁶ Thus, choosing between a game semantics and the standard definition of truth-in-a-model is not a matter of giving sentences of L different truth conditions, but only a matter of choosing how to *describe* those truth conditions.

The differences between the two styles of semantics concern the auxiliary concepts they make use of. When we choose between the two styles of semantics, we are choosing between two different sets of auxiliary concepts, based on how well they align with our particular theoretical purposes. For example, the Tarski-style recursive definition makes use of the concept of satisfaction of a formula by a sequence of objects. The concept of satisfaction is used to give the recursive clauses for the truth of quantified sentences, since their immediate subformulas are open formulas, rather than sentences with a truth value. The game semantics does not need the concept of satisfaction, because the game is structured so that every round is played with respect to a closed sentence. On the other hand, the most important auxiliary concept in the game semantics is the concept of a *strategy*, which does not appear in the Tarski-style semantics.

⁵This is because the game is a finite, zero-sum game of perfect information, so the Gale-Stewart theorem implies that it is determined; whenever the Verifier has a winning strategy, the Falsifier lacks one, and vice versa. See Hodges (2013).

⁶For a proof sketch, see Hintikka and Kulas (1985, pp. 6–7).

The appeal to strategies is important partly because it is necessary for the formal adequacy of the game-theoretical definition of truth. Imagine we had defined truth without the notion of strategy, just in terms of whether the Verifier does in fact win the game played with a sentence S on a model M . To see that this would be an inadequate definition, it suffices to notice that in any particular play of the game, the Verifier might win not because she succeeds in verifying the sentence, but just because the Falsifier fails to make optimal moves.

For example, consider the sentence

$$\forall x \exists y P(y, x)$$

and a model M_N which interprets P as the predecessor relation on the domain of natural numbers. The Falsifier moves first in the game; suppose she chooses 3 for x . Then in the second round, starting from $\exists y P(y, 3)$, the Verifier can choose 2 for y , producing $P(2, 3)$. This results in a win for the Verifier, because 2 does precede 3 in M_N . But that does not show that the sentence is true in this model; it only shows that the Falsifier played badly. If the Falsifier had instead chosen 0 for x , the Verifier would not have been able to win, since there then is nothing she can choose for y to produce a true atomic sentence: no natural number precedes 0. This shows that the Verifier does not have a winning strategy in the game with respect to this sentence and M_N , despite being able to win in some (indeed, almost all) possible plays of that game. The game-theoretical definition of truth requires that the Verifier can *always* choose an appropriate number for y , no matter what the Falsifier chooses for x . Since she cannot, the sentence is not true in M_N , as expected.

4.2.3 *Quantifier moves as investigations*

The concept of strategy is also important for another reason: it draws our attention to the fact that the game may be thought of as an activity governed by practical rationality. Playing the game is a matter of making choices at each round, within the confines of the rules. Formally, we may think of those choices as given by an arbitrary mapping from one game state to another (legal) game state. But informally, it is appropriate to call such a mapping a strategy just because some mappings are better than others with respect to a player's goal of winning the game. Thus, the concept of strategy highlights the teleological structure of playing the

game. It licenses a perspective from which we think of the players as choosing some actions for the sake of others, and ultimately for the sake of winning the game.

In the case of the quantifier rules, the actions connected with making a move in the game have the structure of an investigation. Consider again the rules $G.\exists$ and $G.\forall$. A player must do several things in order to make a move in the game in accordance with one of these rules: she must search for an object, choose a suitable one, and give it a name. How should she do this? Obviously, she should not merely select an object at random. Since she is aiming to win, she must think about which object she *should* choose. This is a strategic problem: she must ask herself which object will best enable her to win, considering what her opponent will do on subsequent rounds.

Because of how the conditions for winning are defined, this strategic problem is just like an algebra problem, in that it induces an investigation structured by a which-question. Consider the case for the Verifier.⁷ The Verifier wins if she can force the game to end with a true atomic sentence. So when she is facing a statement of the form

$$\exists x\phi(x)$$

she should find and name an object that satisfies $\phi(x)$, since the sentence which ends the game will be a subformula of this formula. Thus, $\phi(x)$ acts like the matrix of a which-question structuring her investigation: to make a strategically optimal move, she should attempt to answer the question, which objects are ϕ ? She finds and gives a value for x as a means to answering this which-question, which in turn is a means to winning the game.

It is easiest to see this in the context of a specific example. Suppose I am playing as the Verifier in the game with respect to a model M whose domain consists of the pieces on a chessboard in front of me. The sentence I am given is

$$\exists x(P(x) \wedge B(x))$$

⁷To keep things simple in what follows, I will just focus on the Verifier's moves in accordance with $G.\exists$. All the same considerations apply symmetrically to the Falsifier's moves in accordance with $G.\forall$; we just have to view him as undertaking investigations in which he tries to find an object that makes a certain matrix false, i.e., one that makes its *negation* true. This symmetry follows from the fact that the game semantics is classical, so the two quantifiers are duals: $\forall x\phi(x)$ is equivalent to $\neg\exists x\neg\phi(x)$.

where P is interpreted as a predicate of the pawns, and B is interpreted as a predicate of the black pieces; the sentence thus says that some pawns in the domain are black. It is my turn to select an object for x , in accordance with $G.\exists$. Which object should I choose, given that I am the Verifier and that I am aiming to win the game? The Falsifier will make the next choice, which will determine whether the game ends with an atomic sentence of the form $P(x)$ or one of the form $B(x)$. So the question I must answer is, what is an object x that is both $P(x)$ and $B(x)$? That is, which object or objects are both pawns and black pieces?

To answer this question, I must undertake an investigation: I must try to find such an object, to give as the value of x . My method of searching in this case will obviously be different than the one I use to solve algebra problems. Instead of manipulating equations, I will have to survey the pieces on the chessboard. But the structure of my search is largely the same. I start from an understanding of the problem: I am seeking at least one object among those in a given range that satisfies a certain condition. I take certain steps to look for such an object. If those steps lead me to recognize an object as satisfying the condition, I conclude my search successfully. In this case, I must use empirical rather than mathematical means to search for and recognize the object I am seeking. But that is obviously inessential: my investigation would have this same structure if the model M offered a mathematical, rather than empirical, interpretation of the predicates P and B .

To further cement the parallel between this kind of investigation and the activity of solving an algebra problem, it is helpful to see how the four criteria for 'giving the value of a variable' apply to the moves made by the Verifier in accordance with $G.\exists$. First, scope. The Verifier's choice of a value for x obviously occurs in the context of a particular problem or question. When she is facing a statement of the form $\exists x\phi(x)$, this is the question, which objects are ϕ ? How she should choose, and whether her choice is strategically correct, depends on the particular formula $\phi(x)$. The strategic problem she faces in the game played on a different formula $\psi(x)$ is a different problem, even if the variable is typographically the same, just as two different polynomials give different algebra problems, even if they are both expressed using ' x '.

Next, range. The Verifier is constrained by the rule $G.\exists$ to choose an object from the domain of M to be the value of the variable bound by the quantifier. An understanding of this range is necessary for correct play

in the game. In the chessboard case, if I try to choose, say, a pear from the adjacent fruit bowl to be the value of x , I am not making a legal move, but revealing a misunderstanding: I have not grasped that my choice must be from the pieces on the board. As in an algebra problem, the range affects whether or not there are solutions to the Verifier's strategic problem, and how many solutions there are. There might be no black pawns on the board, or one, or several; this will affect what I can choose, and thus whether or not I have a winning strategy in the game.

Next, completeness. The question structuring the Verifier's investigation for a single quantifier move always requests just a single alternative. It is like the algebraic question 'What is *one* number x such that...?'; it does not require giving more than one alternative or an exhaustive list. She completely answers this question, and the game moves on to the next round, once she has given any one value for the variable bound by the quantifier. This just follows from the way the rule $G.\exists$ is written. Thus we have no need to distinguish, here, between giving a value for a variable and giving a complete solution statement.⁸

Finally, canonical form. Notice that the rule $G.\exists$ requires the Verifier to give a *name* for the object she selects. In effect, this part of the rule states a standard of canonical form for the Verifier's moves. It is not enough for the Verifier to select an object; she must make it known *which* object she has selected to be the value of the quantified variable, using a canonical expression for it.

To really see the significance of this requirement, we have to reflect on what counts as a 'name' in a first-order language L . Since L is a first-order language, it is equipped with a category of terms, which in general can be divided into three classes: individual variables, individual constants, and terms constructed out of these by means of function symbols or other operators. Names are obviously terms, since the players substitute names for individual variables when making a quantifier move. But which terms count as names? Do the names include just the individual constants, such as ' a ' or ' 1 '? Or could they include functional terms, like ' $f(a)$ ' or ' $1 + 1$ '? What if the language includes other term-forming op-

⁸It is interesting to consider how to generalize the game semantics to languages containing other types of quantifiers, such as numerical quantifiers, by adding rules to allow or require selecting multiple objects. But since most of the quantifiers I am interested in here are definable in terms of the standard first-order quantifiers, I leave this question for further research.

erators, like a definite description operator? Does $'(\iota x)\phi(x)'$ count as a name, or not?

There is no general *syntactic* answer to these questions. Whether or not a given term counts as a name is a matter of how and for what purpose the language is being used. In the context of the game semantics, the category of names is used in opposition to the category of individual variables. A variable represents a choice that a player has yet to make; a name represents the outcome of a choice. The purpose of these choices is to move the game closer to its final state, where a winner can be determined by consulting the model. Thus, how we conceive of names is linked with how we conceive of the models or interpretations of the language L with respect to which the game is played.

From this we may deduce a few things about the category of names. First, the notion of model we are working with requires that a model directly determines a truth value for atomic sentences, so that at the end of the game, the model is sufficient to determine who has won. Assuming the usual denotational semantics for atomic sentences, all that really matters is that names are given a *fixed* interpretation in the model, which is to say that at the end of the game, the model itself contains all the information needed to determine which individual in the domain each name picks out, without any further input from the players. This is the paradigm of a canonical expression: what's important is that at the end of the game, no question can arise about *which* object a player selected, since such questions would prevent determination of whether the final sentence is true or false. In the context of the game semantics, a 'name' is thus simply a canonical expression.

In fact, we can go further than this. Since the scope properties of the quantifiers depend on the order in which players' choices are made, it is important that names have a fixed interpretation in the model as soon as they are introduced, not just at the end of the game.⁹ Players

⁹This is not to say that every name must have a fixed interpretation *before* it is introduced by a player's choice. We must allow for some objects in the domain to lack names (for example, when the language is countable but the domain is uncountable). Such objects can still be chosen and given names by players. In such cases, we should picture the player's choice of a name for an object as simultaneously *extending* the language and the model: the effect of the player's choice is to add the name to the language and the mapping to the model. This depends in turn on the players' background ability to indicate which object is meant *without* using a name.

must not be able to ‘put off’ choices by using expressions whose interpretation depends on choices to be made in subsequent rounds. This excludes any terms which are semantically variable, such as open functional terms, from serving as names. Likewise, it means that terms containing quantifier-like operators, such as a description operator, must be interpreted as having a fixed, determinate interpretation when they are introduced by a player’s choice during a quantifier move.

To see this, suppose the language L contains both a unary function symbol ‘ s ’ and a binary relation symbol ‘ $<$ ’, and consider the game for

$$\exists x \forall y (y < x)$$

on the model N , which interprets ‘ $<$ ’ as the usual ordering relation on the domain of natural numbers, and ‘ s ’ as the usual successor function. This sentence is of course false in N , since no natural number is greater than every natural number. But if the Verifier is permitted to use an expression like ‘ $s(y)$ ’ as the name she substitutes for x , she will have a winning strategy and the sentence will turn out to be true: there is no number that the Falsifier could choose for y to make

$$\forall y (y < s(y))$$

false. The problem here is that the expression ‘ $s(y)$ ’ does not have a fixed interpretation *at the time the Verifier makes her choice*, since its interpretation depends on the value of y ; there is no determinate answer to the question of which particular number she is choosing, so ‘ $s(y)$ ’ cannot be a canonical expression or name. Similarly, we cannot in general admit terms like ‘ $(\iota z) \forall w (w < z)$ ’ or ‘the successor of the Falsifier’s next choice’ as names. Even if we treat these as syntactically closed terms, the relevant *semantic* feature of a name is that it is a canonical expression, so that no question can arise about which individual it is associated with at any point in the game.¹⁰

¹⁰It is even alright if we include among names expressions which are *syntactically* variables, so long as we think of the quantifier moves in the game as fixing their interpretations. Imagine, for example, that in the game played on the chessboard model we make moves by hanging tags on the chess pieces, on which we have inscribed the variable whose interpretation the choice fixes. Thus, I hang a tag labeled ‘ x ’ on the pawn in front of the black queen, say, to mark that pawn as my choice for the value of x . What’s important is not the syntactic category of the label, but that each syntactic variable in the final formula corresponds to a tag which has been hung by the end of the game, and that the tags were hung in a certain order without changing.

Within those bounds, there's plenty of room for different standards about which expressions count as names, which will in turn depend on how we conceive of the language and its models. The important point is that some such standard governs each quantifier move in the game, that there is a distinction between those expressions in the language which signify a determinate choice of individual and those which do not. For until such a choice is expressed, the move is not complete and the game cannot continue; and if the choice is not determinate, the game will not assign the correct truth conditions to the original sentence.

Thus, each quantifier move satisfies all four criteria for giving the value of a variable. The strategic problem a player is faced with when she makes a quantifier move is to find a value for a variable which makes a certain formula true. She must find at least one such value within a certain range, the domain of the model. To complete her move, she must give that value by means of a canonical expression. Her action therefore has just the same structure as the act of solving an elementary algebra problem: it is an investigation, an attempt to answer a which-question.

4.2.4 *Solving the puzzle*

It's time to return to the opening puzzle of this chapter: how are the first-order and question-answer analyses of a specificational sentence related? Are they competing theoretical alternatives, or are they somehow compatible? I will now argue that they are best seen as compatible, by showing that the question and answer in a specificational sentence can be found in the first-order formula which analyzes it, if we interpret that formula using the game semantics.

This is easiest to see if we reflect a little more on the general strategy of analysis I outlined above. The first-order formula which analyses a specificational sentence, which I will call a *specification-formula*, takes the form of a statement with one or more leading existential quantifiers. The open formula beneath those quantifiers will consist of a conjunction of two main parts, corresponding to the subject and the complement of the natural language sentence. Let us call the part corresponding to the complement the *specifier clause*. It consists of a conjunction of identity statements, each of which has one of the existentially-bound variables on one side, and a constant term on the other. Let us call the part corresponding to the subject the *question clause*. The question clause will itself

be a conjunction of two parts, which following Belnap and Steel (1976) I will call its *matrix* and its *request*. The matrix clause represents the contribution of the noun phrase in the subject. The request represents the contribution of the determiner in the subject; it expresses the uniqueness and distinctness conditions (if any) imposed by the natural language determiner. Thus, the general first-order form of a specificational sentence will be:

$$\exists x_1 \dots \exists x_n ((M(x_1, \dots, x_n) \wedge R(x_1, \dots, x_n)) \wedge S(x_1, \dots, x_n))$$

Here, M schematizes the matrix, R the request, and S the specifier clause.

For example, consider again the specificational sentences (14)–(16), together with their first-order analyses:

- (17) a. A good book to read is *Waverley*.
 b. $\exists x (B(x) \wedge x = w)$
- (18) a. The author is Scott.
 b. $\exists x ((A(x) \wedge \forall y (A(y) \rightarrow y = x)) \wedge x = s)$
- (19) a. Two of the characters are Edward Waverley and Baron Bradwardine.
 b. $\exists x \exists y ((C(x) \wedge C(y) \wedge x \neq y) \wedge x = e \wedge y = b)$

In (17-b), the specifier clause is $x = w$, and the question clause contains only the matrix $B(x)$; the request is null or empty, since the indefinite article carries no uniqueness or distinctness condition. In (18-b), the request contains the uniqueness or exhaustivity condition $\forall y (A(y) \rightarrow y = x)$, in order to capture the additional contribution of the definite article. In (19-b), which has multiple leading existential quantifiers, the specifier clause contains a conjunct for both existentially bound variables, and the request $x \neq y$ captures the distinctness (but non-exhaustivity) condition of the English numerical quantifier ‘two’.

There are a few important things to notice about analyses of this form. First, consider the specifier clause, which corresponds to the complement of the natural language sentence. To capture the distinction between specification and predication, the specifier clause relies on the same device as we use in the language of elementary algebra: it gives the values of the variables introduced by the subject using equations in canonical form. It would not do to allow an arbitrary first-order formula in the

specifier clause. If in (18-b) we replace $x = s$ with an arbitrary formula $H(x)$, the result is suitable for analyzing *predicational* sentences like

(20) The author is much hated.

but not *specificational* sentences like (18-a). We cannot even allow arbitrary formulae whose predicates are restricted to the identity predicate, for that would not rule out expressions like $x = x^2$. This again is suitable to analyze *predicational* sentences like

(21) The number is its own square.

so if we allow it, we fail to capture the distinction between specification and predication. It is important for the adequacy of the analysis as an analysis of *specification* that we be able to think of the terms that appear on the right hand side of the equations in the specifier clause as names or canonical expressions, so that the clause does not merely constrain the values of the variables, but actually *gives* or specifies them.

The second important point is that the first-order form of a *specificational* sentence trivially entails a formula in which the specifier clause has been dropped, i.e., a formula of the form

$$\exists x_1 \dots \exists x_n (M(x_1, \dots, x_n) \wedge R(x_1, \dots, x_n))$$

consisting of just the existential closure of the question clause. Let us call such a formula the *corresponding question-formula* of a *specification-formula*.

In what sense does such a formula ‘correspond to a question’? This is where the game semantics is most revealing. An existential formula is true when the Verifier has a winning strategy in the associated game, which will begin with a series of quantifier moves made by the Verifier. In order to find such a strategy, the Verifier must undertake an investigation (or a series of nested investigations) structured by the question: which objects are M and R ? Or more colloquially, since R plays the role of a determiner: what are R objects that are M ? The Verifier will have such a strategy, and the sentence will be true, if and only if there is an answer to that question.

This question is exactly the one left open by the subject of the original *specificational* sentence. A few examples suffice to illustrate this. Intu-

itively, the questions left open by the specificational subjects in (17-a)–(19-a) are:

- (22) What is a good book to read?
- (23) Who is the author?
- (24) Who are two of the characters?

These questions are colloquial expressions of the following more cumbersome questions, which are exactly the corresponding questions for (17-b)–(19-b):

- (25) a. What is (at least one example of) an individual that is a good book to read?
b. $\exists xB(x)$
- (26) a. What is (at least one example of) the unique individual that is an author?
b. $\exists x(A(x) \wedge \forall y(A(y) \rightarrow y = x))$
- (27) a. What are two (examples of) distinct individuals that are characters?
b. $\exists x\exists y(C(x) \wedge C(y) \wedge x \neq y)$

A specification-formula entails its corresponding question-formula. Under the game semantics, that means that if the Verifier has a winning strategy in the game associated with a specification-formula, she also has one in the game associated with its corresponding question-formula. Of course, many first-order formulae will have this property, not all of which will be specification-formulae. The interesting property of the specification-formula is that its truth not only guarantees that there *is* a winning strategy for the Verifier in the game for the question-formula, but it also *encodes* or *displays* such a strategy, or rather the part of the strategy governing the Verifier's initial quantifier moves. She can just choose for each variable x_i the individual given for x_i in the specifier clause. As we just observed, the specifier clause gives this individual by means of a canonical expression, in an equation in canonical form. So it tells the Verifier everything she needs to know to make her initial moves in the game for the corresponding question-formula: it tells her how to answer the strategic question of which objects are M and R .

Notice, then, what we have achieved: we have rendered the question-answer analysis of specification compatible with the first-order analysis using the game semantics. Analyzing a specificational sentence in a first-order language yields a specificational formula. This formula is true just in case its specifier clause answers the strategic question faced by the Verifier in the game for the corresponding question-formula. This is the question we intuitively associate with the specificational sentence's subject. The specifier clause is derived from its complement. So, a specificational sentence is true just in case its complement answers the question presented by its subject. At least, that is so to the extent that we take the analysis to successfully capture and explicate the intuitive truth conditions of the sentence. This solves the puzzle: a specificational sentence is true just in case a certain first-order formula is true, and that formula is true just in case one part of it answers the question contained in the other part. There is no tension between the first-order and question-answer analyses.

None of this, I hope, seems particularly surprising. The value of the discussion above does not consist in presenting a novel analysis of specification, or a new technical approach. The logical tools we have used are in fact rather simplistic, and the analyses are the sort one would be trained to produce in an elementary logic course. Greater sophistication is possible in both respects, though more sophisticated approaches will still have these basic tools at their core. Instead, the value lies in having made explicit, step by step, the connection between these tools and practices of asking and answering questions. The intuition that a specificational sentence 'gives the value of a variable', or some values of some variables, is commonly accepted, but it is not usually further explored. I have tried, here, to articulate what lies behind this intuition, by describing what it means to 'give' the value of a variable within an investigation, and then describing the truth conditions of specificational sentences in terms of investigations, by means of the game semantics. In logic, as in algebra, a variable is best seen as an *interrogative* device, an expression whose role and purpose derives from our activities of formulating which-questions and finding their answers. To say that a specificational sentence gives a value for a variable is to assign it a special role within that activity. This is the role of connecting a question with its answer, of connecting the start of an investigation to its end.

4.3 BACK TO NATURAL LANGUAGE

With these insights in hand, we are in a better position to explain some of the properties of natural language specifications. I would like to conclude this chapter by revisiting some of those properties, and seeing what light we can shed on them by viewing specificational sentences as playing an analogous role in natural language as their counterparts play in elementary algebra.

One of the features of specificational sentences that supported the question-answer analysis was their sensitivity to prior questions in discourse.¹¹ This is clearest in the case of specificational sentences that have anaphoric pronouns or demonstratives as subjects, such as

(28) It was Cameron.

When such a sentence is asserted, its meaning obviously depends on some antecedent question. If the question is “Who was at the door?” then it specifies the person who was at the door; if it is “Who received the mysterious envelope?”, it specifies the person who received the mysterious envelope. But it makes no sense to ‘specify’ a person apart from some question being asked about who satisfies a certain condition. Similarly, a specificational sentence sounds odd when the question left open by its subject does not agree with one asked in prior discourse:

(29) Who was at the door?

(30) ??The person who received the mysterious envelope was Cameron.

Unless it is clear from the context how the question of who received the mysterious envelope is related to the question of who was at the door, it sounds terrible to assert (30) after (29). A speaker who does so is felt to be inexplicably changing the subject.

This sensitivity to prior questions is easily understood when we view specificational sentences in terms of their role in investigations. A specificational sentence has a scope; it only makes sense in the context of an investigation structured by a question. It makes no more sense to utter (28) when no question is salient than it does to offer $x = 2$ as a solution when no algebra problem has been given. Similarly, uttering (30) after

¹¹See the discussion in Section 2.2.2.

(29) is like using $x = 2$ as part of a derivation for the problem given by $y^2 - 9 = 0$: it just isn't clear how giving a value for x is relevant to finding the values of y . The sensitivity of specifications to prior questions reflects the fact that a specificational sentence is part of an activity structured by the goal of answering a certain question. When it is not clear what the goal is, or how it advances that goal, a specificational sentence becomes unacceptable.¹²

Another interesting feature of specificational sentences concerns the distinction between specification and pure equation.¹³ There is an intuitively felt asymmetry in specificational sentences like (31) which is absent from pure equatives like (32):

(31) The month I got married was May.

(32) The month I got married was my uncle's favorite month.

We saw in Section 3.2 that the main problem with the equative analysis of specification is that it does not have a plausible account of this distinction. In assimilating specificational sentences to algebraic solution statements, am I not inviting the same objection? It is true that algebraic solution statements are made using equations, and I have proposed that we can make a similar use of equations in first-order analyses of specifications. But the important idea behind the assimilation is that specifications play a certain role in investigations, and only a special *kind* of equation can play that role. By inquiring into that role, we have arrived at better criteria to distinguish specification from pure equation: a specification occurs in the context of a problem which introduces the variable, it gives a value in the variable's range, it completely satisfies the request of the problem, and it is in canonical form. Pure equations in general do not satisfy these criteria—after all, some equations give problems, rather than solutions—so they don't count as specifications.

¹²It is worth pointing out that there is a rich literature on dynamic semantics, beginning with Heim (1983/2002), which deals in part with the problem of how to represent inter-sentential anaphoric relationships like those I have pointed to in (28)–(30). I believe the techniques developed in this literature will prove useful for giving a formal representation of specificational sentences like (28), which will elude the strategy of first-order analysis I explained above because the scope needed to make sense of 'it' extends beyond the sentence itself. I do not pursue this issue here, though.

¹³For the distinction, see the discussion in Section 2.1.3. The distinction plays an important role in my criticism of the equative analysis in Section 3.2.

These criteria, especially those concerning range and canonical form, are just as useful for distinguishing specification and pure equation in natural language. The difference between (31) and (32) is that ‘May’ is a canonical expression for a month, but ‘my uncle’s favorite month’ is not. This is hardly an isolated example. There are general, stable conventions in natural language about which sorts of expressions are suitable for specifying values in the range determined by a certain noun. Consider the following examples:

- (33) The month I got married was ...
 a. {January, February, March, ...}
 b. {my uncle’s favorite month, the last month of the Alaskan summer, the month before my divorce, ...}
- (34) The color of the paint is ...
 a. {red, yellow, blue, ...}
 b. {the color symbolizing our Glorious Revolution, the color which attracts the bees, the cheapest color the store sold, ...}
- (35) The way she traveled was ...
 a. {by caravan, in an armored car, along the Rhine, ...}
 b. {the way her father had expressly forbidden, however she liked, the way everyone travels in Stockholm, ...}
- (36) The tournament winner was ...
 a. {Armando, Beatrice, Christine, Deep Blue, ...}
 b. {last year’s champion, the mysterious newcomer, the best European player, ...}

In each of these cases, the continuations in the (a)-group will yield a specificational sentence, while those in the (b)-group generally yield a pure equative. This shows that we distinguish in practice between expressions which can and cannot be used to specify months, colors, ways of traveling, or tournament winners. It is a stable feature of the (a)-group expressions that they are canonical expressions for values in these ranges, while the (b)-group expressions are not. That is, the (a)-group expressions are immune to further ‘But which do you mean?’ questions, while the (b)-group expressions by default are not. Moreover, it is clear that the (a)-group expressions do not have this property absolutely, but only

relative to a certain conception of a range: ‘January’ cannot be used as a canonical expression for a color or a way of traveling, for example; and ‘Deep Blue’ can be used to specify a tournament winner, but not a person. It would be easy to multiply such examples. The contrast between specificational sentences and pure equatives reflects our knowledge of which expressions are canonical for the values in a given range.¹⁴

There is a subtle point here concerning the role of the head noun in indicating a range of values. Suppose I ask:

(37) When did you get married?

You might respond in two different ways:

- (38) a. I got married in May.
b. I got married in 1982.

Just given the form of my ‘When?’-question, both responses are open to you, and to give a helpful answer you must think about my reason for asking. If you think I am interested in what a good time of year to get married would be, you will use the first sentence rather than the second. If you think I am interested in what sort of gift to buy for your anniversary, you will use the second rather than the first. As we observed already in Chapter 1, it is often communicatively useful to be more explicit about what is at issue in such exchanges. Thus, to express your understanding of my question, you might instead answer using a specificational sentence containing a noun that indicating the type of answer you thought I wanted:

- (39) a. The *month* I got married was May.
b. The *year* I got married was 1982.

¹⁴Similar ideas have been put forward by several other philosophers. Brandom (1994, Ch. 7), whose account of existential quantification in natural language has influenced the story I’ve told in this chapter, speaks of *canonical designators* as being organized into addressed *spaces*; which ‘space’ a designator belongs to would be made explicit by a noun such as ‘month’ or ‘color’. Similarly, Moltmann (2013, p. 533) gives a quasi-fictionalist account of reference to numbers, describing number terms as akin to names for fictional characters like ‘Hamlet’. She thinks of claims made with such terms as being true or false relative to different *strategies* of evaluation, and thinks of nouns like ‘number’ or ‘character’ as indicating such a strategy.

What's interesting about these answers is that, while they are largely equivalent to the ones in (38), the different nouns import different restrictions on the class of canonical expressions available for specifying when you got married. While an answer like

(40) The month I got married was September.

is merely false, an answer like

(41) #The month I got married was 1982.

seems ill-formed and uninterpretable, even though '1982' would be a perfectly fine (short) answer to the original question.

There is a natural representation of these facts on the strategy of analysis I have pursued in this chapter. In effect, the head nouns in (39) alter how the language is to be interpreted in this context. We can think of them, in the game semantics, as marking a difference in how the game can be correctly played. The restriction introduced by 'month' or 'year' is not merely an additional condition on *what* the Verifier must be able to find for the sentence to be true; it is a condition on the range *where* she may search—that is, on the domain of the model. Consequently, it restricts the expressions which the Verifier may use to give a value in that range without stipulatively altering the language when making a quantifier move. The difference between (39-a) and (40) is that they declare different strategies for the Verifier in the same model. The difference between (39-a) and (39-b) is that they require models with different domains. This is one way to explain why (41) is not merely false, but uninterpretable.¹⁵

These examples point toward a final interesting feature of specificational sentences: their logical relationship to 'reduced', non-copular paraphrases. Many specificational sentences seem to be logically equivalent to such a paraphrase, or at least entail it. For example, the specificational sentences in (39) seem equivalent to their counterparts in (38).

¹⁵Hintikka (1973, Ch. 3) makes a similar claim to explain the communicative difference between logically-equivalent statements like 'All swans are black' and 'No non-black things are swans'. His idea is that the head noun of the subject specifies the domain of the model, which forms the players' field of search in the game. We prefer the former sentence in communication because it is much more determinate what counts as seeking and finding a swan than a non-black 'thing'.

Similarly, (42-a) is equivalent to (42-b), and (43-a) is equivalent to (43-b) (supposing that only one person can win the tournament):

- (42) a. One good book to read is *Waverley*.
 b. *Waverley* is a good book to read.
- (43) a. The tournament winner was Armando.
 b. Armando won the tournament.

There is no great mystery about why this is so. To explain the equivalence, we can simply appeal to the logic of identity. The first-order form of these specificational sentences, according to the analysis outlined above, is

$$\exists x(\phi(x) \wedge x = c)$$

This is materially equivalent to

$$\phi(c)$$

which corresponds to the form of the ‘reduced’ paraphrase.

At the same time, the game semantics provides some resources to account for the felt differences between the two forms. It is easy to see that, despite their *material* equivalence, the forms are *practically* non-equivalent. The game for the reduced form of these sentences is trivial; the Verifier has won or lost after zero rounds. By contrast, the game for the specificational form begins with a choice by the Verifier and lasts two rounds, and so it will require the Verifier to have a non-trivial strategy. In this game, the Verifier faces the same strategic problem as in the game for the corresponding question-formula:

$$\exists x\phi(x)$$

Determining a winning strategy for this game is, in general, a non-trivial problem; it’s just that the specificational form already encodes its solution. In other words, the specificational form represents a practical situation intermediate between a difficult problem, and no problem at all: it is the situation of facing a problem that has already been solved. Insofar as we can see our own practical and informational situation as we learn about our world reflected in the Verifier’s situation as she investigates a model, the game semantics provides an illuminating description of the unique position of specificational sentences in our investigations.

These observations permit me to answer a final methodological question about specificational sentences in natural language. What makes a sentence count as ‘specificational’? Is the category of specificational sentences a syntactic, semantic, or pragmatic category? In the first instance, the category of specificational sentences is a pragmatic category. A sentence ultimately counts as specificational because it plays a certain role in practices of investigation: it links an investigatory problem to its solution. This role is related to, but distinct from, the pragmatic role of sentences which present the problem without presenting the solution, and of sentences which present the solution without explicitly linking it to the problem.

In the second instance, though, a sentence may be particularly well-suited to play this pragmatic role because it has certain syntactic properties. For example, a language might encode the distinction between the roles of canonical and non-canonical expressions for values in a certain range as a distinction between two syntactic categories. Since differences in pragmatic role may explain the syntactic differences between specificational and non-specificational sentences, specificational sentences may form a syntactic category in a derivative sense. They may also form a semantic category, though this depends on how we characterize semantics. On the analysis I have given, a specificational sentence is materially equivalent to a reduced, non-specificational sentence. If semantics is narrowly construed as concerned only with material truth conditions, then there is no semantic category of specificational sentences. But if semantics is construed more broadly, as concerned with general facts about how sentences are interpreted, then specificational sentences do form a semantic category, and the game-theoretical semantics gives us the resources to say what is distinctive of that category. Specificational sentences differ from their reduced forms by requiring the Verifier to make additional quantifier moves; they differ from the corresponding question sentences by displaying the terms she should use in those moves.

Our troubles are not over yet, however. The most famous example of an equivalence between a specificational sentence and its reduced form is the one with which I began this dissertation:

- (44) a. The number of moons of Jupiter is four.
 b. Jupiter has four moons.

I am not quite ready to say that we understand the relationship these sentences exhibit. In this chapter, I have focused on specificational sentences where the complement is a clear case of a name or singular term. These are the specificational sentences that have a straightforward first-order analysis. But as we saw in earlier chapters, and again in examples like (34) and (35), English contains many examples of specificational sentences where the complement contains a different kind of expression, such as an adjective or adverb, which is not obviously a semantically singular term. 'Four' is one such complement. Should we treat these expressions as names, or not? Should we say they denote objects, or something else? That is the puzzle with which I began, and it is still outstanding.

The problem presented by such sentences is that the specificational forms appear to involve first-order quantification over individuals, but in the reduced form, the specified individuals function as predicates or other modifiers. Yet the two forms seem equivalent. So can we give them a first-order analysis, in the manner I've suggested? Or do we need further logical resources? I shall take up these questions in the next chapter.

Chapter 5

Numbers, again

It is time to return to the puzzle about Frege's famous sentence:

- (1) The number of Jupiter's moons is four.

If the analysis of specificational sentences in Chapter 4 is right, this sentence has a form which can be represented like this:

$$\begin{aligned} \exists x((x \text{ is a number of Jupiter's moons} \\ \wedge \forall y(y \text{ is a number of Jupiter's moons} \rightarrow y = x)) \\ \wedge x = 4) \end{aligned}$$

The important feature of this analysis is that '4' is treated as a canonical expression for a number, the kind of expression you can use to say which number you mean without being liable to a request for clarification. This differentiates the specifier clause ' $x = 4$ ' from a predicational clause like ' x is even', or a non-specificational equation like ' $4x = x^2$ '.

This analysis treats the definite description as introducing a first-order variable. Accordingly, it treats ' \dots is a number of Jupiter's moons' as a first-level predicate, and '4' as a semantically singular term. So if the analysis can be applied to this case, it entails that numbers are objects. We have not yet decided, though, whether we can follow Frege that far. Should we say that numbers are objects, or not?

In Chapter 1, we saw some reasons for doubt. Frege himself noted that (1) is equivalent to

- (2) Jupiter has four moons.

Here, ‘four’ appears attributively, as an adjective or determiner, and the concept of number does not appear. If we follow the analysis above, we don’t seem to have a way to capture this equivalence. It is hard to see how we could read the formula above, or anything logically equivalent to it, back into English as (2). Precisely because we have taken ‘four’ as a singular term in our first-order analysis, it seems we cannot treat it as an adjective or determiner, which are expressions that require an argument.

Here is the trouble. Within the confines of first-order logic, no one expression can be interpreted as both a singular term and a predicate—in Frege’s terms, as an expression that stands both for an object and for a concept.¹ The syntactic categories of terms and predicates are *exclusive* in formal languages, as are the corresponding semantic categories of objects and concepts. That means that when giving an analysis of (1), we have to choose between representing ‘four’ as a term, or representing it as a predicate (of some level or other).² At the same time, because (1) and (2) are non-accidentally equivalent, it is natural to think that ‘four’ makes the same contribution in each, and its role in the two sentences should be analyzed in the same way. Yet in (1), ‘four’ appears to be a term, while in (2), it appears to be a predicate. Logical analysis forces us to choose between these options. If we represent it as a term, we can use the first-order analysis for (1), but we can’t represent the equivalence between (1) and (2). If we represent it as a predicate, there is still hope for representing the equivalence, but only if we give up the first-order analysis of specification developed in Chapter 4 for one that allows higher-order expressions in specifier clauses.

¹To avoid prolixity, and to remain as close as possible to Frege’s own terminology in a crucial passage in *Grundlagen* §27, I will use ‘concept’ in a slightly wider sense than Frege himself usually does. Frege distinguishes concepts (what monadic predicates stand for) from relations (what polyadic predicates stand for). I will essentially use ‘concept’ as he would have used ‘concept or relation’, that is, as a term for what predicates stand for in general. It is harmless to paper over this distinction, because I am interested in the way that *both* concepts and relations contrast with objects, and nothing in my argument depends on the contrast between the monadic and polyadic cases.

²It would be open to us to analyze ‘four’ as a quantifier, since in a Fregean logical system, quantifiers are second-level predicates—that is, predicates of first-level concepts. In the type theory standardly used by linguists, the analogous choice is between assigning ‘four’ to type e , on the one hand, or to a higher type such as $\langle e, t \rangle$ or $\langle \langle e, t \rangle, t \rangle$ on the other.

In *Grundlagen* §57, Frege responds to this puzzle by taking (1) as the more revealing form, opting to analyze it as an identity statement, in which ‘four’ occurs as a singular term:

I have already drawn attention above to the fact that we speak of “the number 1”, where the definite article serves to class it as an object... what we have is an identity, stating that the expression “the number of Jupiter’s moons” signifies the same object as the word “four”. (Frege, 1884/1980, §57)

He says we should not be bothered by the fact that ‘four’ occurs as an adjective in (2), since “that can always be got round” via its equivalence with (1).

Frege’s response looks increasingly unsatisfying, though, once we see that the puzzle is not limited to number words like ‘four’. As the previous chapters have emphasized, the relationship between (1) and (2) is just one example of a general pattern. A specificational sentence is generally equivalent to a corresponding reduced sentence. The equivalence holds even when the head noun of the subject disappears in the reduced sentence, and the complement appears as an adjective, adverb, clause, or other non-nominal constituent. Here are three more examples:

- (3) a. The color of Io’s surface is yellow.
b. Io’s surface is yellow.
- (4) a. The way I’m losing weight is by giving up sweets.
b. I’m losing weight by giving up sweets.
- (5) a. The reason Henry fled was that the cops had found his stash.
b. Henry fled because the cops had found his stash.

Adopting Frege’s attitude in these cases would apparently lead us to conclude that the adjective ‘yellow’, the adverbial phrase ‘by giving up sweets’, and the clause ‘the cops had found his stash’ each stand for an object. On any ordinary understanding of what it means to be an object, this result is unintuitive, and might make us hesitant to adopt Frege’s response to the puzzle.

The unintuitive consequences are not the most serious problem. More worrisome is that Frege’s response threatens to obliterate any distinction between the semantic roles of adjectives, adverbs, complete clauses, and

other non-nominal expressions. If the fact that an expression can appear opposite a definite description in a specificational sentence is reason enough to analyze it as a singular term, then it appears that just about any kind of expression can be a singular term. Frege himself would surely want to resist this conclusion, and regard at least adjectives and adverbs as predicates or concept-words.³ But once we recognize the general pattern, we have just as much reason to treat ‘four’ as standing for an object in (1) as ‘yellow’ in (3-a) and ‘by giving up sweets’ in (4-a).

What we’d like is a general analysis of specification, one that more easily accommodates the equivalence between a specificational sentence and its reduced form, even when the specificational complement appears to be a non-nominal expression. As things stand, (1) and (2) and the pairs of sentences in (3)–(5) pull us in two different directions. If we focus on the specificational sentences, we’ll be drawn (with Frege) toward treating their complements as terms. If we focus on the reduced forms, we’ll be drawn toward treating them as predicates. It would be nice to have an analysis that avoids this tension, or at least provides clear criteria for going in one direction rather than the other.

As I’ve suggested, the tension arises because our contemporary logical systems treat the distinction between objects and concepts as exclusive. These systems ultimately derive from Frege’s *Begriffsschrift* and are based on his function-argument analysis of propositions. The function-argument analysis is one of Frege’s major achievements. It is what allowed him to abandon the subject-predicate analysis of the earlier logical tradition, which in turn allowed him to represent sentences containing multiple levels of generality—that is, nested quantifiers. This was the key innovation that made Frege’s logical system adequate for representing mathematical reasoning in a way that no previous system was. In the most basic case, the function-argument analysis of propositions *just is* an analysis into objects and concepts. So it is not too much of an overstatement to say that Frege’s distinction between objects and concepts

³In the *Grundlagen*, for example, he argues that numbers are importantly different from colors in that they are not properties, either of external things or of concepts; he therefore supposes that color words are predicates while number words are not (Frege, 1884/1980, Cf. §§21–27, 45–57). On the other hand, Frege did later come to accept that a sentence or clause stands for an object. After he drew his distinction between *Sinn* and *Bedeutung*, he held that a sentence generally stands for a truth value, which he regards as an object.

made modern logic possible. But it is also the source of our present problem, because on Frege's understanding of analysis, the roles of function and argument are exclusive. Any logical system that shares the function-argument analysis of the *Begriffsschrift* will force us to choose between representing 'four' as a term or as a predicate, as standing for an object or for a concept.

A radical solution to this problem would involve developing a whole new logical system for the analysis of natural language, one that did not force the choice on us.⁴ I think it is possible that we might need such a system to fully resolve the problem, one that is not based on the function-argument analysis. But the radical solution is far outside the scope of what I hope to accomplish here. In this chapter, I will argue that even within a basically-Fregean framework, there is a way of thinking about the choice that makes it less puzzling. Although the roles of objects and concepts are exclusive *within* a given proposition or thought, we should recognize that one and the same thing can be both an object and a concept with respect to different propositions.

My strategy will be to interpret Frege's understanding of the concept-object distinction, and the reasons he took it to be exclusive. The distinction is often understood in ontological terms, as an exclusive distinction between two kinds of *entity*, or being. This reading faces serious interpretive challenges, however. It also makes the puzzle very urgent: if 'four' necessarily stands for two different entities in (1) and (2), it is unclear how or why these sentences are equivalent, since they seem to be talking about different things.

On my interpretation, the concept-object distinction should instead be understood as a distinction between two logical or epistemological *roles*. To call something an object or concept is to say something about the role it plays in a system of scientific thought. By recognizing the distinction as a distinction between two different epistemological roles, I leave room for the possibility that one and the same thing may play both roles at different times or places in our system of thought. There is then no tension between saying that 'four' stands for something that plays the

⁴I have some sympathy for this idea, because I think it not only could help us understand the semantics of specification, but also the semantics of predication and quantification. These devices are intimately connected, and there are deep puzzles associated with how we represent each of them in the logical systems we have today, because each of those representations depends on Frege's function-argument analysis.

role of an object in (1), but plays the role of a concept in (2). This simply reflects the fact that these sentences do different work in our language.

What are these two roles, and when does something play one role rather than the other? Although Frege himself does not provide clear answers to these questions, I will suggest that we can derive clearer criteria from the game semantics introduced in Chapter 4. Roughly, an object is something that can be sought, found, and specified by the players as they make quantifier moves in the semantic game. A concept is what guides a player as she makes such moves: it is how she distinguishes a strategic choice from a non-strategic one. It thus plays the role of constraining her search. We will see that this way of thinking about the distinction preserves the important features of Frege's understanding, and can even be motivated by his early views. In particular, it preserves the idea that the distinction is exclusive with respect to a given interpretation of the language, and it allows us to endorse the claim that numbers are objects. At the same time, it divests that claim of the ontological consequences it is often felt to have. Numbers are objects, as are colors, ways, and reasons—but platonism about these things does not follow.

5.1 FREGE'S UNDERSTANDING OF OBJECTS AND CONCEPTS

A standard reading of Frege's concept-object distinction construes it in ontological terms. According to this reading, Frege's distinction is on a par with the traditional distinctions between particulars and universals, or substances and attributes, although it is importantly different from them. This ontological distinction corresponds to a syntactic one. There are two different kinds of expressions in language: proper names, and concept words. Proper names denote, or refer to, or stand for, objects; concept words denote, or refer to, or stand for, concepts (of various levels). At least in a logically-ideal language, there will be a perfect mapping between the types of expressions in the language and the kinds of beings that they stand for in the world. I will refer to this reading as the *Ontic Picture*. On the *Ontic Picture*, to classify something as an object or a concept is to make a metaphysical claim about the kind of being it is, and it is appropriate to analyze an expression as a proper name or concept word

just in case it stands for the corresponding kind of entity.⁵

My main goal in this section is to show that Frege does not hold the Ontic Picture. This is because he does not understand the concept-object distinction primarily as a distinction between ontological categories. Instead, he understands it as a distinction between epistemological categories, between different roles things can play in thought. There is good textual evidence for this interpretation in the *Grundlagen*, where Frege makes his claim that numbers are objects. I will also argue, more tentatively, that Frege maintains this view throughout his career. We should not think of Frege's concept-object distinction in ontological terms, even after he draws his distinction between *Sinn* and *Bedeutung* in 1891.

What do I mean by 'ontological' and 'epistemological' in this context? I don't want to quibble too much over the terminology. Here is the issue, in terms that would have been familiar to Frege: does the concept-object distinction apply to things as they are in themselves, just as a matter of what they *are*, and independently of their relationship to us? Or does it apply to things in virtue of how we grasp them, or understand them? If it is the former, then Frege's distinction is an ontological distinction, and the Ontic Picture remains a viable interpretive option. If it is the latter, then it is an epistemological distinction, and we must find an alternative to the Ontic Picture to interpret Frege.

I should say that I am not trying to construct a comprehensive interpretation or a maximally-consistent reading of Frege's remarks about concepts, objects, or the distinction between them. I don't have the space for this task here, and it has already been carried out in more detail by others. In fact, much of what I have to say here is anticipated by other interpreters, such as Tugendhat (1970), Sluga (1977, 1-4), and Sluga (1980). Instead, I want to extract one line of thought from Frege's writing that I think is particularly relevant for understanding his claim in the *Grundlagen* that numbers are objects, and that sheds light on the logical form of specificational sentences. This is what's needed to show that there is a sense of 'object' such that both we and Frege can endorse the claim

⁵Versions of the Ontic Picture can be found in, for example, Dummett (1981), Wright (1983), Burge (1992), and Hale and Wright (2001). Dummett, for example, says: "Frege's use of the ontological term 'object' is strictly correlative to his use of the linguistic term 'proper name': whatever a proper name stands for is an object, and to speak of something as an object is to say that there is, or at least could be, a proper name which stands for it" (Dummett, 1981, p. 55).

that numbers are objects, without having to endorse any particular metaphysics.

5.1.1 *Concepts and objects in the Grundlagen*

The main evidence for my interpretation of the concept-object distinction comes from several passages in the early parts of the *Grundlagen*. I will begin with these passages, and then relate them to other claims Frege makes about the distinction in his early work.

In the introduction to the *Grundlagen*, Frege states three principles that have guided him throughout the work:

In the enquiry that follows, I have kept to three fundamental principles:

1. always to separate sharply the psychological from the logical, the subjective from the objective;
2. never to ask for the meaning of a word in isolation, but only in the context of a proposition;
3. never to lose sight of the distinction between concept and object. (Frege, 1884/1980, p. X)

The third principle is most immediately relevant, though as we will see, it is importantly related to the first two. Frege thinks the distinction between concepts and objects is fundamental, and crucial. He goes on to argue that numbers are neither concepts nor subjective, that we must take them to be objects if we are to understand arithmetic, and that various other writers have erred because they have failed to recognize these facts in one way or another.

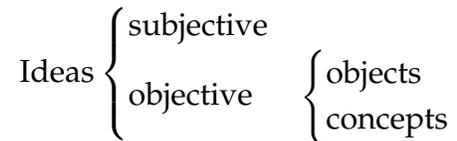
It's a bit puzzling, then, why Frege does not immediately tell us how he is thinking of these two categories. What *are* objects, as opposed to concepts? Are they species of a common genus? What is it for something to belong to one category or the other?

Frege answers these questions in a footnote to §27. There, he classifies concepts and objects as species of 'objective ideas', which he contrasts with subjective ideas:

An idea (*Vorstellung*) in the subjective sense is what is governed by the psychological laws of association; it is of a sensible, pictorial character. An idea in the objective sense belongs

to logic and is in principle non-sensible, although the word which means (*bedeutet*) an objective idea is often accompanied by a subjective idea, which nevertheless is not its meaning. Subjective ideas are often demonstrably different in different men, objective ideas are the same for all. *Objective ideas can be divided into objects and concepts.* I shall myself, to avoid confusion, use 'idea' only in the subjective sense. It is because Kant associated both meanings with the word that his doctrine assumed such a very subjective, idealist complexion, and his true view was made so difficult to discover. The distinction here drawn stands or falls with that between psychology and logic. If only these themselves were to be kept always rigidly distinct! (Frege, 1884/1980, §27 n. 1, my emphasis)

So concepts and objects are each a kind of idea, though they are to be distinguished from subjective ideas like sensations or imaginings. The taxonomy Frege is giving us in the passage therefore looks like this:



True to his word, Frege refrains from calling objects and concepts 'ideas' hereafter. But this passage is the only place in the *Grundlagen* where Frege explains what he means by 'object' and 'concept', and the point is that this explanation is very puzzling if he is thinking of objects and concepts as ontological categories, categories that things belong to independently of how we think about them and represent them. Why would he call them 'ideas', and take pains to distinguish them from subjective ideas, if he didn't think the distinction had anything to do with how we think and know about things? It seems that if Frege were thinking of objects and concepts as ontological categories, he could simply say: there's two kinds of stuff in the world, and this is what they're like. Our ideas about them, or representations of them, have nothing to do with their status.

One might feel that Frege is emphasizing that concepts and objects are *objective*, but not that they are *ideas*. That is certainly true. But he makes it clear that he does not understand this objectivity as being grounded in

ontology. In §26, he gives a detailed characterization of his understanding of objectivity. He begins by comparing the objectivity of numbers to that of physical objects, saying that “number is no whit more an object of psychology or a product of mental processes than, let us say, the North Sea is.” But the analogy is somewhat misleading, because he does not think of objectivity as confined to physical reality:

I distinguish what I call objective from what is handleable or spatial or actual (*Wirklichen*). The axis of the earth is objective, so is the centre of mass of the solar system, but I should not call them actual in the way the earth itself is so. We often speak of the equator as an *imaginary* line; but it would be wrong to call it a *fictitious* line; it is not a creature of thought, the product of a psychological process, but is only recognized or apprehended by thought. (Frege, 1884/1980, §26)

In this passage, Frege is trying to convince us that there is a third category, in addition to the objective realm of natural or physical or empirical things, and the subjective realm of things created purely by thought, such as impressions or imaginings. This third category consists of things that are objective, but not ‘actual’. This is the category to which numbers belong.

But what is it to be objective, if being objective is not limited to natural or physical or ‘actual’ things? Frege begins to give his positive characterization of objectivity in the following paragraph, using the example of geometry:

What is objective in [geometry] is what is subject to laws, what can be conceived and judged, what is expressible in words. (Frege, 1884/1980, §26)

He continues:

I understand objective to mean what is independent of our sensation, intuition and imagination, and of all construction of mental pictures out of memories of earlier sensation, *but not what is independent of reason*—for what are things independent

of reason?⁶ To answer that would be as much as to judge without judging, or to wash the fur without wetting it. (Frege, 1884/1980, §26, my emphasis)

His picture, then, is that objectivity does not derive from ‘actuality’, but from having a certain relationship to reasoning and judgment. We should distinguish what is objective from what is subjective or psychological, and different for different people. But for Frege, this does not banish what is objective from the realm of thought entirely. Instead, things are objective because they are grasped *in a certain way* in thought, the way we grasp things when we reason and make judgments, which is different from how we grasp things in our subjective sensations or imaginings.

Frege is here expressing a Kantian or neo-Kantian understanding of the objective. What is objective is not simply what’s ‘out there’, as it is in itself, apart from our capacity to recognize it and form judgments about it. He regards *that* kind of objectivity as out of reach for us, as impossible as ‘washing the fur without wetting it’. Instead, objectivity concerns how *we grasp* things, or the role they play in our cognition. That is why it is appropriate for Frege to speak of objective ‘ideas’ in his footnote on the following page. Since the distinction between objects and concepts is a distinction among objective ideas, it is a division of things insofar as they play a certain role in our cognition. This makes it an epistemological rather than ontological distinction.

It follows that Frege’s claim that numbers are objects should not be understood as a claim about the kind of entity they are. Frege makes it clear that his main concern is with grounding our *knowledge* of arithmetic, not with locating numbers in a particular ontological category. This point emerges clearly in §57, the crucial passage in which Frege finally decides that numbers are objects, as opposed to concepts, and considers the relationship between sentences (1) and (2). It’s worth quoting his reasoning there more fully than I did above:

It is time to get a clearer view of what we mean by our expression “the content of a statement of number is an assertion about a concept”. In the proposition “the number 0 belongs

⁶Austin’s translation uses ‘the reason’ here for *der Vernunft*; but it seems pretty clear that Frege intends something more like Reason with a capital ‘R’, not ‘the’ reason. He is speaking of our capacity for reasoning, not some particular reason.

to the concept F'' , 0 is only an element in the predicate (taking the concept F to be the real subject). For this reason I have avoided calling a number such as 0 or 1 a *property* of a concept. *Precisely because it forms only an element in what is asserted, the individual number shows itself for what it is, a self-subsistent object.* I have already drawn attention above to the fact that we speak of “the number 1”, where the definite article serves to class it as an object. In arithmetic this self-subsistence comes out at every turn, as for example in the identity $1 + 1 = 2$. Now our concern here is *to arrive at a concept of number usable for the purposes of science*; we should not, therefore, be deterred by the fact that in the language of everyday life number appears also in attributive constructions. That can always be got round. For example, the proposition “Jupiter has four moons” can be converted into “the number of Jupiter’s moons is four”. . . . what we have is an identity, stating that the expression “the number of Jupiter’s moons” signifies the same object as the word “four”. And *identities are, of all forms of proposition, the most typical of arithmetic.* (Frege, 1884/1980, §57, my emphasis)

Here is what is happening in this passage. Frege has previously argued, starting in §46, that a judgment as to how many things are F is a judgment about the concept F . In this passage, Frege is rejecting the view that numbers are the predicates of such judgments, that is, higher-order concepts (what he here calls a “property of a concept”), in favor of the view that they are objects.⁷ Notably, he gives no substantive metaphysical argument for this conclusion. Instead, he reasons that our use of the definite article with number words provides evidence that we treat numbers as objects. This evidence is not misleading, because in fact, numbers must be objects if they are to be “usable for the purposes of science”. Only the assumption that numbers are objects can ground our knowledge of arithmetical propositions, such as $1 + 1 = 2$. This is because

⁷In the earlier parts of the *Grundlagen*, Frege rejects the view that numbers are *first-level* concepts, or “properties of external things”, as well as the view that numbers are subjective. So by also rejecting the view that numbers are higher-level concepts in this passage, he leaves only one place for them in the taxonomy of ideas given in §27: they are objects, as opposed to any kind of concepts.

numbers appear as objects in such propositions, and numerals appear as proper names in our ways of expressing them.

To some interpreters, Frege's reasoning here is too simple. Dummett (1991, Ch. 9), for example, thinks Frege's choice to treat numbers as objects is completely unjustified. He shows that it is possible to develop arithmetic consistently using the 'adjectival strategy', treating number words instead as expressing higher-level concepts. Thus, Frege has failed to provide any strong support for the position he has chosen. For Dummett, a weightier argument is required, because he thinks Frege accepts the Ontic Picture. Dummett writes that "ontologically expressed, [Frege] is trying to establish that numbers must be regarded as objects" (Dummett, 1991, p. 101). Because he thinks an ontological conclusion hangs on the argument in §57, Dummett sees Frege's failure to rule out the alternative position as philosophically costly, and unsatisfactory.

On my reading, by contrast, Frege's reasoning here is perfectly adequate. Frege is choosing between two different reconstructions of the role of numbers in the science of arithmetic. One of those roles makes sense of the form of our actual arithmetical knowledge and our usual way of expressing it; the other doesn't. Dummett is correct that an arithmetic based on concepts can be made to work. But such an arithmetic would not be *our* arithmetic, in which we use number-words as proper names for objects. It would thus leave us puzzled about the ground of our actual arithmetical knowledge. In this sense, Dummett's concept-arithmetic is not 'usable for the purposes of science', even if it is more faithful to the way we use number words outside of arithmetic proper. Frege dismisses the alternative position so casually because he sees the issue in epistemological terms, and to him, it is clear that our actual arithmetical knowledge is possible only if we regard numbers as objects.

There is a natural objection to my reading, though, that can be raised at this point. In §57 and in many other passages, Frege speaks of the 'self-subsistence' or 'independence' (*selbständigkeit*) of numbers in connection with their status as objects. Doesn't that show that he regards numbers as something external to us, and their status as objects as independent of our way of thinking about them? No. Frege explicitly qualifies this language a few pages later:

The self-subsistence which I am claiming for number is not to be taken to mean that a number word signifies something

when removed from the context of a proposition, but *only to preclude the use of such words as predicates or attributes*, which appreciably alters their meaning. (Frege, 1884/1980, §60, my emphasis)

His appeal here to the Context Principle, the second of the fundamental principles set out in the introduction, provides further evidence for my reading. Frege is saying that his claim that numbers are ‘self-subsistent’ should not be understood in a way that is independent of their role in arithmetical thought. Instead, it is intended to mark the contrast between their actual role as objects from the role of concepts.⁸

A remark in the *Begriffsschrift* further supports this response. In §9, where Frege introduces his function-argument analysis of propositions, he says that ‘the number 20’, a proper name, “gives rise to an independent idea (*selbständige Vorstellung*)”, unlike quantifier expressions like ‘every positive whole number’ (Frege, 1879/1967, §9). This explicit combination of *selbständig* with *Vorstellung* shows that for the early Frege, something’s being ‘self-subsistent’ or ‘independent’ is entirely consistent with its being an idea in the objective sense. Frege is here, as in *Grundlagen* §57, using *selbständig* to mark the contrast between the roles of objects and higher-order concepts *within* thought, not to imply something about the ontological status of objects *apart from* their role in thought. Thus, the objection is misguided. Frege’s talk of the ‘independence’ of objects does not straightforwardly imply that when something is an object, it has this status independently of how we grasp it in thought.

5.1.2 Three theses about the distinction

So far, I have argued that the concept-object distinction should be understood as a distinction between two different roles things can play in thought, based on passages in the *Grundlagen*. We can bolster this argument if we look at three theses about the concept-object distinction which

⁸Why does Frege use terminology that has such an ontological ring to it? A native German speaker has suggested an explanation to me: perhaps Frege uses ‘self-subsistence’ (*selbständigkeit*) because it shares the common root *-stand* with ‘object’ (*Gegenstand*). Since it bears this etymological relationship to ‘object’ but not to ‘concept’ (*Begriff*), Frege’s remark in §60 might indicate that he is simply using ‘self-subsistence’ as a term for *whatever* it is that differentiates the role of an object from that of a concept.

are at work in the *Grundlagen*, but articulated more clearly elsewhere in Frege's writing. They are:

Exclusivity. The distinction between concepts and objects is exclusive, and the two roles are complementary.

Priority of the proposition. Concepts and objects are obtained by splitting up or analyzing complete thoughts.

Multiple analysis. Thoughts can be analyzed in multiple ways; there is no uniquely correct analysis of a thought into objects and concepts.

Taken together, these three theses yield a straightforward argument that Frege's concept-object distinction should be seen in epistemological rather than ontological terms. I will first explicate the three theses, then give the argument.

It is pretty clear that Frege thinks the roles of concepts and objects are exclusive. As we saw above, Frege takes it as a fundamental principle in the *Grundlagen* "never to lose sight of the distinction between concept and object", and he adds that "it is a mere illusion to suppose that a concept can be made into an object without altering it" (Frege, 1884/1980, p. X). Later, in 'On concept and object', he is more explicit. He criticizes Kerry for misunderstanding his view that nothing can be both an object and a concept, and says clearly that on his way of understanding these categories, an expression that stands for an object cannot stand for a concept:

the three words 'the concept "horse"' do designate an object, but *on that very account they do not designate a concept*, as I am using the word. (Frege, 1892/1997e, 184, my emphasis)

Why does Frege think the distinction is exclusive? The answer lies in the second thesis, his doctrine of the priority of the proposition. Frege thinks of concepts and objects as being obtained by analyzing, or splitting up, a complete thought. Complete thoughts are explanatorily prior to objects and concepts. Identifying an object or a concept within a thought is something we do *after*, and in virtue of, grasping the thought as a whole. A particularly explicit formulation of this doctrine appears in a letter Frege sent to Anton Marty, shortly before the *Grundlagen* was published:

Now, I do not believe that concept formation can precede judgment, because this would presuppose the independent existence of concepts, but I think of a concept as having arisen by decomposition from a judgeable content. I do not believe that for any judgeable content there is only one way it can be decomposed, or that one of these possible ways can always claim objective pre-eminence. (Frege, 1882/1997d, p. 81)

Notice that Frege immediately follows his statement of the priority of the proposition with a statement of the doctrine of multiple analysis. The two doctrines are closely related: it is in some sense because we grasp a complete thought prior to splitting it into parts that we are able to split it up in multiple ways.

These two theses can be traced all the way back to the *Begriffsschrift*. There, in §3, Frege identifies a criterion for the sameness of thoughts (or 'judgeable contents'), namely, that they have all the same possible consequences. Thus, for example, while "The Greeks defeated the Persians at Plataea" differs grammatically from "At Plataea the Persians were defeated by the Greeks", they express the same thought, because they have all the same consequences. This gives us a grip on what it is for two thoughts to be the same or different, prior to and independent of any way of splitting them up into objects and concepts. He then goes on to describe that process in §9, where he introduces the function-argument analysis of propositions:

Let us suppose that the circumstance that hydrogen is lighter than carbon dioxide is expressed in our formula language. Then in the place of the symbol for hydrogen we can insert the symbol for oxygen or that for nitrogen. This changes the sense in such a way that 'oxygen' or 'nitrogen' enters into the relations in which 'hydrogen' stood before. If an expression is thought of as variable in this way, it splits up into a constant component, which represents (*darstellt*) the totality of relations, and a symbol which can be thought of as replaceable by others and which denotes (*bedeutet*) the object that stands in these relations. The former I call a function, the latter its argument. *This distinction has nothing to do with the conceptual content, but only with our way of grasping it.* (Frege, 1879/1967, §9, my emphasis)

Frege is here operating at the level of syntax, assuming that a complete thought has been expressed in the notation of the *Begriffsschrift* and then considering how it can be parsed into a function and an argument by regarding one or another of its symbols as replaceable by others. But he already clearly has the idea that these symbols stand for something, speaking of the symbol we regard as replaceable as *denoting* an object, and the function symbol (what he will later call a ‘concept word’) as *representing* the relations it stands in. So we may pass from the level of syntax to the level of content, or what the symbols mean. Frege’s picture is then that we obtain a concept from a complete thought by thinking of some part of that thought as replaceable, or letting it vary. The concept is what’s ‘left over’ when we regard an object position as variable within a thought.

This explains why the roles of concept and object are exclusive: to vary a part of something, you have to isolate that part from something else that you hold fixed. The roles are exclusive because they represent complementary parts of a complete whole; they are not independent constituents, but instead have a kind of figure-ground relationship within a thought. Concepts and objects are exclusive, and complementary, because of how we recover them from thoughts which we grasp antecedently, by understanding their logical relationships to other thoughts.

Importantly, Frege thinks that analysis into objects and concepts is something *we* do to thoughts. As the last sentence in the passage shows, thoughts do not *come* with a structure; we impose a structure on them, through our choices about how to express them in language, and which parts of those expressions to regard as variable. Frege thinks we are free to do this in different ways. This is the doctrine of multiple analysis, which he asserts frequently throughout his career, from the *Begriffsschrift* onward. Another clear statement of it appears in ‘On Concept and Object’, which was published after the *Grundlagen*, in 1892:

In the sentence ‘There is at least one square root of 4’, we are saying something, not about (say) the definite number 2, nor about -2 , but about a concept, *square root of 4*; viz. that it is not empty. But if I express the same thought thus: ‘The concept *square root of 4* is realized’, the first six words form the proper name of an object, and it is about this object that something is being said. But notice carefully that what is being said here is

not the same thing as was being said about the concept. This will be surprising only to somebody who fails to see that a thought can be split up in many ways, so that now one thing, now another, appears as subject or predicate. The thought itself does not yet determine what is to be regarded as the subject. (Frege, 1892/1997e, pp. 187–188)

Furthermore, it is clear that Frege maintained the doctrine of multiple analysis in the *Grundlagen* itself. The doctrine is at work in §§64–66, where Frege asserts that we can ‘re carve’ a proposition about parallel lines into one about directions. Frege is also applying it in §57, to relate “Jupiter has four moons” and “The number of Jupiter’s moons is four”.

Why is the doctrine of multiple analysis important to Frege? Because we grasp the logical relations of a thought by analyzing it in different ways. Consider again the example from *Begriffsschrift* §9, that hydrogen is lighter than carbon dioxide. By splitting this thought up in one way, regarding hydrogen as an object and being lighter than carbon dioxide as the concept, we identify the common concept in this thought and certain other thoughts, such as the thought that something is lighter than carbon dioxide, or that everything lighter than carbon dioxide is lighter than silicon dioxide. Identifying this common element is what we need to grasp the logical relationship of the original thought to a certain family of other thoughts, and to represent those relationships in a formal language.

But of course, the relationships that this analysis allows us to grasp do not exhaust the logical relationships of the original thought. It also has relationships to the thought that hydrogen is lighter than nitrogen, or that hydrogen is lighter than everything heavier than helium. To grasp those relationships, we need to split the original thought up differently, regarding carbon dioxide as the object, and being heavier than hydrogen as the concept. The doctrine of multiple analysis ensures that it is possible and legitimate to do this, and therefore to represent the logical relationships of a thought to different families of other thoughts. In effect, the point of the doctrine of multiple analysis is to give us the flexibility needed to represent all the logical relationships of a thought in a system like the one of *Begriffsschrift*.

So much for the three theses. Now that we understand them, we can give a straightforward argument that Frege’s distinction between concepts and objects is not an ontological distinction. Call this the *argument*

from analysis:

1. Concepts and objects are obtained by splitting up thoughts.
2. The way in which a thought is split up is not a feature of the thought itself, but of how we grasp it.
3. So: the division of a thought into concepts and objects is a matter of how we grasp it.
4. So: something's being an object or a concept is a matter of how we grasp the thought in which it occurs, not a feature that it has independently of us.

The first premise is just the doctrine of the priority of the proposition. The second premise is part of Frege's understanding of the doctrine of multiple analysis. These two premises imply the intermediate conclusion (3), that analysis of a thought into objects and concepts depends on how we grasp that thought. But if that is right, it follows that (4) things count as objects or concepts in virtue of how we analyze the thoughts in which they occur. Note that this conclusion concerns *why* things count as objects or concepts, regardless of what such things are, or whether anything plays both roles. The point is that a thing's status as an object or a concept is an epistemological matter, a matter of how *we grasp* it in thought, not an ontological feature it has independently of us.

Here, then, is the picture that emerges. When Frege emphasizes that "Jupiter has four moons" can be converted into "The number of Jupiter's moons is four" in *Grundlagen* §57, he is thinking of these sentences as expressing the same thought, which is independent of its analysis, and which we are free to express in either way. How we analyze it will depend on our choice of expression, and the point of analyzing it in some particular way is to exhibit its logical relationships to certain other thoughts. This particular thought obviously has relationships to other thoughts in, say, the science of astronomy; but in the *Grundlagen*, Frege is concerned to bring out its relationships to arithmetical propositions. Arithmetical propositions have a standard means of expression, in which terms for numbers are always regarded as arguments, not as function symbols. So standard arithmetical practice analyzes arithmetical thoughts in a way that assigns numbers to the object role, rather than the concept

role, within them. For the purpose of exhibiting its logical connections to arithmetical thoughts, the best choice for an expression of the thought about Jupiter's moons is therefore one that also treats 'four' as a proper name. For Frege, this is the specificational sentence: "The number of Jupiter's moons is four."

So in §57, Frege is arguing that *we can and should* analyze thoughts involving numbers in such a way that assigns numbers to the object-role, as opposed to the concept-role, *given* that we want to make their logical relationships to arithmetical thoughts and our usual means of expressing them clear. Other analyses of thoughts involving numbers are possible, and appropriate, for other purposes. His claim that numbers are objects is not a claim about their ontological status, but about how arithmetical thoughts are standardly grasped and analyzed by us. *We could* develop arithmetic based on a different analysis, as Dummett shows. But as long as we haven't in fact done so, the best way to make it intelligible why arithmetical knowledge is relevant to claims with non-arithmetical content is to assign numbers the same epistemological role in the latter as in the former.

5.1.3 *Concepts and objects after the distinction between Sinn and Bedeutung*

So at least in Frege's early work, concepts and objects should be thought of as exclusive epistemological categories. The distinction between them is a functional distinction, between two complementary roles that result from the analysis of complete thoughts. Specifically, it is a distinction between two roles found in *objective* thoughts, the kind of thoughts that can be shared between different people, and that we grasp and use in science. To call something an object or a concept is to say it plays one kind of role in scientific thought, rather than the other.

I would now like to address an objection to the line of interpretation I have offered so far. The objection is based on Frege's later work, and runs as follows. In the *Begriffsschrift* and the *Grundlagen*, Frege had not yet made his distinction between *Sinn* and *Bedeutung*. Once he does make the distinction, he speaks of two different levels, or realms. Thoughts and their parts belong to the realm of *Sinn*; these are what we grasp and express in language. But the realm of *Sinn* should be sharply distinguished from the realm of *Bedeutung*. Frege thinks of the realm of *Bedeutung* as

a realm of entities 'out there', the things we study in science, things like planets or protons. It is the realm of things we think *about*, and which make thoughts true or false.

Having distinguished these two realms, Frege faces a choice about where to place objects and concepts, and it is very clear that he places them both in the realm of *Bedeutung*. This shows, according to the objector, that the Ontic Picture must be correct: concepts and objects are two different sorts of *entities* out there in the world. In light of this, Frege should be seen as jettisoning the epistemological remarks about concepts and objects from the earlier period, such as his classification of them as species of 'ideas' and his claim that they are obtained from our choices about how to analyze thoughts. These remarks translate most naturally to claims about the realm of *Sinn*. Since objects and concepts themselves belong to the realm of *Bedeutung* rather than *Sinn*, my thesis cannot be right, at least for the post-1891 Frege. In his later thinking, the concept-object distinction is an ontological, rather than epistemological, distinction.

In response, I first want to note that reading the distinction between *Sinn* and *Bedeutung* back into the *Grundlagen* is a bad idea. Perhaps it is true that the later Frege thinks objects form an ontological category, and takes the claim that numbers are objects as an ontological claim. But if so, that just means Frege needs a much stronger argument for this claim than the one he gives in *Grundlagen* §57 from considerations about grammar and the purposes of science. The interpretation I have offered so far is still preferable for understanding that passage, because it does not construe the argument as a weak argument for a strong ontological conclusion, but a sound argument for a weaker epistemological conclusion. Thus, thinking of concepts and objects in terms of epistemological roles still makes better sense of how Frege understands his claim in the *Grundlagen* that numbers are objects, and the equivalence between (1) and (2).

Still, I think a case can be made that Frege continues to think of the categories of concepts and objects as epistemological roles even after he draws the distinction between *Sinn* and *Bedeutung*. There is no denying that later on, Frege places both objects and concepts in the realm of *Bedeutung*. He says that proper names *bedeuten* objects, while concept words *bedeuten* concepts. But the crucial premise in the objection is that Frege thinks of the realm of *Bedeutung* in ontological terms, as the realm of stuff 'out there' that our thoughts are about, and that make thoughts true or

false. I think we have good reason to resist this premise as a characterization of Frege's official view of *Bedeutung*. My interpretation here is more tentative, but let me try to spell it out briefly.

Frege's habit of classifying ordinary physical objects as the *Bedeutungen* of proper names is what makes the crucial premise in the objection look plausible. Here is just one example of a passage where he does so, from a 1906 diary entry Frege titled 'Introduction to Logic':

If we say 'Jupiter is larger than Mars', what are we talking about? About the heavenly bodies themselves, the *Bedeutungen* of the proper names 'Jupiter' and 'Mars'. (Frege, 1906/1997c, p. 295)

Remarks like this motivate the translation of *Bedeutung* as 'reference', and lead us toward thinking of the relationship between an expression and its *Bedeutung* along the lines of the Ontic Picture, as a kind of correspondence between expressions and independent entities in the world. Frege is saying here that 'Mars' refers to Mars, the planet itself. Since 'Mars' is a proper name, it refers to an object; so Frege is thinking of planets, at least, as examples of objects. It is natural to read Frege as saying that objects, and perhaps all *Bedeutungen*, have the kind of ontological status that planets have: they exist, independently of us and how we think of them.

But such remarks are probably misleading when we look at them in the context of what Frege says about *Bedeutung* in general. When Frege argues that other kinds of expressions also have *Bedeutungen*, he does not picture the relationship as a correspondence between expressions and independent entities. He does not try to establish that concept words, for example, have *Bedeutung* by identifying any entities to which they correspond. Instead, there is an important continuity between the kind of reasoning he employs in these arguments and in *Grundlagen* §57: he argues that they *must* have *Bedeutungen*, on the basis of what's needed to understand scientific practice and thought.

The passage I just quoted occurs in the context of one such argument. Frege argues that since words like 'Mars' and 'Venus' have *Bedeutungen*, we must take concept words to have them as well. Here is the complete argument:

As far as the mere thought-content is concerned it is indeed a matter of indifference whether a proper name has a *Bedeutung*, but in any other regard it is of the greatest importance; *at least it is so if we are concerned with the acquisition of knowledge.* It is this which determines whether we are in the realm of fiction or truth. Now it is surely unlikely that a proper name should behave so differently from the rest of a singular sentence that it is only in its case that the existence of a *Bedeutung* should be of importance. *If the thought as a whole is to belong to the realm of truth, we must rather assume that something in the realm of Bedeutung must correspond to the rest of the sentence...* If we say 'Jupiter is larger than Mars', what are we talking about? About the heavenly bodies themselves, the *Bedeutungen* of the proper names 'Jupiter' and 'Mars'. We are saying that they stand in a certain relation to one another, and this we do by means of the words 'is larger than'. This relation holds between the *Bedeutungen* of proper names, and so must itself belong to the realm of *Bedeutungen*. (Frege, 1906/1997c, 295, my emphasis)

In short, concept words must have *Bedeutungen* because if they did not, it would be impossible to understand how 'Jupiter is larger than Mars' as a whole could express something that is true or false. The fact that the names of the planets and 'is larger than' each have a *Bedeutung* distinguishes this thought from mere fiction or poetry, such as 'Odysseus observed Mars', or 'Mars vorshooned upon the Jabberwock'.

Notice that this argument is quite weak if we understand *Bedeutungen* in ontological terms. Why would we insist that every scientific expression *bedeutet* an independent entity? Consider words like 'the' or 'all'. They are clearly important for science, and Frege would count them as concept words. But what is important about these words is that they are *meaningful* or *significant*, not that they correspond to some *entity*. That is, it's important that they make a clear contribution to a scientific thought—or, once we separate the content of the thought from its *Bedeutung*, that they play a clear role in determining its truth. That requirement can be satisfied by treating these terms syncategorematically, by giving a rule for computing the *Bedeutung* of the expressions in which they occur from the *Bedeutung* of the other constituents, without assigning any entity as

the *Bedeutung* of 'the' or 'all' itself. So if Frege is insisting in this passage on a categorematic treatment of words like 'the' or 'all', on which they are conceived as standing for certain entities *in addition to* merely making a determinate contribution toward the truth value of a complete thought, his conclusion is clearly unjustified.

Instead, we should read this argument as telling us something about what the concept of *Bedeutung* is for. Frege is clearly saying here that all the expressions in a statement must have *Bedeutungen* if that statement is to express something truth-evaluable, something that is a candidate for knowledge. It seems to me that this is the best way to understand *Bedeutung* in general: having a *Bedeutung* is what's required for an expression to contribute to the statement of a scientific thought. Frege says this succinctly in 'Comments on *Sinn* and *Bedeutung*':

Of course in fiction words have only a sense, but in science and wherever we are concerned about truth, we are not prepared to rest content with the sense, we also attach a *Bedeutung* to proper names and concept words; and if through some oversight, say, we fail to do this, then we are making a mistake that can easily vitiate our thinking. (Frege, 1891/1997a, p. 173)

Thus, even the later Frege should be read as thinking of *Bedeutung* in primarily epistemological terms. The point of claiming that an expression has a *Bedeutung* is to say it is suitable for making truth-evaluable claims, claims that are candidates for knowledge and scientific demonstration. Such claims are objective, but as we saw already in *Grundlagen* §26, there is no need to understand this objectivity as a correspondence between each part of the claim and an entity 'out there'. For an expression to have a *Bedeutung* is simply for it to have a determinate semantic role in scientific language. Accordingly, the categories of concepts and objects represent two different kinds of semantic role, two different kinds of contribution that an expression can make in the expression of a truth-evaluable thought.

Why does Frege so often insist that the *Bedeutung* of a proper name like 'Mars' is something like Mars, a physical object in the ordinary sense? After all, if all that is required for an expression to have *Bedeutung* is that it makes a clear contribution toward the truth value of statements in

which it occurs, we could assign it a very different *Bedeutung*. For example, we could say that 'Mars' *bedeutet* $\lambda P.[P(\text{Mars})]$, as Montague does. Does Frege's preference for Mars as the *Bedeutung* of 'Mars' reveal that there is more to his notion of *Bedeutung* than semantic role?

At this point, Frege's interpreters divide. Dummett (1981), for example, sees a tension in Frege's thinking about *Bedeutung*. He thinks Frege conceives of it *both* in terms of semantic role, in the way I have been suggesting, and in terms of what he calls the 'name/bearer' model. It is the name/bearer model that leads Frege to say that 'Mars' *bedeutet* Mars, a real entity in the world, rather than something like $\lambda P.[P(\text{Mars})]$. For Dummett, the two ways of thinking about *Bedeutung* are complementary, and both are essential to Frege's view of *Bedeutung* in the case of proper names, but they are in tension in the case of concept words, precisely because concept words do not seem to name any entity. Interpreters like Tugendhat (1970), on the other hand, think that Frege is only committed to thinking of *Bedeutung* in terms of semantic role. Either we can derive the fact that 'Mars' *bedeutet* Mars from this idea, or this is simply an extra assumption on Frege's part, unjustified by his official characterization of *Bedeutung*.

I will not try to settle this dispute. I simply want to note that to be faithful to Frege, we must at least *allow* that an object in the ordinary, physical sense could play the role of an object in Frege's official, epistemological sense. This is an appealing idea, because it helps us see why scientific claims give us knowledge about physical things in the world, and why their truth or falsity depends on those things. Still, I think it would be a mistake to read Frege's characterizations of *Bedeutung* as *requiring* that objects or concepts have the kind of ontological status that planets do. Although objects and concepts belong to the realm of *Bedeutung*, that is compatible with thinking of them in epistemological terms, as two different roles to be played in determining the truth of a scientific thought.

5.1.4 *What are the two roles?*

I have now argued at some length that Frege's distinction between concepts and objects should be thought of as a distinction between two epistemological categories. This interpretation follows naturally from remarks that he makes in the *Begriffsschrift* and the *Grundlagen*, and it ap-

pears that he maintains this view even in his later work, after he distinguishes *Sinn* from *Bedeutung* and classifies concepts and objects as belonging to the realm of *Bedeutung*.

I have not yet said much about how we should think about these two roles, though. What kind of contribution to thought is associated with each role? When does something play the role of an object, as opposed to a concept? This is the question we need to answer to really understand Frege's claim that numbers are objects. So far, we mostly seem to have learned that if numbers are objects, then they do not play the role of concepts in (arithmetical) thought. Without an independent grip on each of the two roles, this does not get us very far. I began by asking whether Frege is correct to classify numbers as objects. How can we decide which role numbers really play? In order to answer this question, we need some criteria that will tell us how to *apply* his notions of 'concept' and 'object' in the case of numbers.

Unfortunately, this is the point at which Frege's own explanations become unsatisfying, because his official position is that we cannot define what it is to be an object. He says as much in 'Function and Concept':

the question arises what it is that we are here calling an object. I regard a regular definition as impossible, since we have here something too simple to admit of logical analysis. (Frege, 1891/1997b, p. 140)

Frege does think it is possible to give a definition of concepts: in the same essay, Frege characterizes them as functions whose values are truth values. But because the truth values are merely a special sort of object for Frege, this definition cannot be understood apart from the idea that functions relate things to objects, so it presupposes that we already have an understanding of what objects are. Unless we know what it means to be an object, we don't know what it means to be a concept, either.⁹

Frege does say some things in various places to try to help his readers understand how he thinks of the two roles. Most of these remarks do

⁹Dummett (1981), Wright (1983) and Hale (2001a) criticize Frege on just this point. If we don't know what it is to be an object, then we don't know what it is to be a first-level concept, since a first-level concept is a function from objects to truth values. They conclude that unless we have some satisfactory criteria for what it means to be an object, Frege's whole hierarchy of levels is threatened, since each level consists of functions whose domain is the functions defined at the previous level.

little more than characterize concepts and objects in terms of each other, though. For example, Frege in several places characterizes objects and concepts in terms of singular predication. He says that “a concept is for me a possible predicate of a singular judgement-content, an object that which can be subject of the same”, and that “with a concept the question is always whether anything, and if so what, falls under it”, while the question of what falls under an object makes no sense—an object is, rather, what falls under a first-level concept. (Frege, 1884/1980, §66 n. 2 and §51). But in Frege’s view, predication is the fundamental logical relationship. He offers no independent account of it; the claim that singular predication is the relationship of an object falling under a concept *is* his account. Thus, knowing that concepts are essentially predicative while objects are not will not help us apply the distinction.

There are two ideas in Frege’s work, though, that may seem to characterize the roles of objects and concepts in more informative terms. The first is the idea that objects are the *Bedeutungen* of proper names, while concepts are the *Bedeutungen* of concept words or incomplete expressions. The second is the idea that objects, unlike concepts, have identity conditions. Let’s see if either is any help.

The first idea is to characterize the roles of concepts and objects by leveraging Frege’s claim that they are the *Bedeutungen* of concept words and proper names. Provided we can classify expressions into these two syntactic categories, we can then use our general understanding of *Bedeutung* to say what the roles of concepts and objects are more specifically. Call this the *syntactic priority* strategy for characterizing the distinction between the two roles.

This strategy immediately runs into a problem. According to the strategy, if we want to understand what an object is, we must first understand what a proper name is. But what is a proper name? Frege does not give a theoretical answer to this question. He mostly seems to rely on intuitions about whether or not an expression is ‘complete’, as well as a few grammatical heuristics like the use of the definite article, to identify proper names.

Dummett (1981, Ch. 4), Wright (1983, Ch. 2), and Hale (2001a, 2001b) all criticize Frege for failing to provide satisfactory theoretical criteria for an expression to be a proper name, and attempt to develop a better account. In their view, the main problem is to distinguish genuine proper names from other expressions that can occupy the same grammatical po-

sitions, such as quantifier expressions like ‘nobody’. Dummett supplies some tests to distinguish proper names from such expressions by their behavior in inferences. Wright and Hale both argue that Dummett’s criteria are not sufficient on their own, but can be made to work if appropriately supplemented. These efforts are not an unqualified success, even by their authors’ own lights.¹⁰

Even if we could find satisfactory criteria, though, there is a bigger problem with the syntactic priority strategy. Frege does think that the categories of objects and concepts are correlated with the categories of proper names and concept words, but that is not because he thinks the syntactic distinction is prior. Here’s an intuitive way to see the point. It is true that Mars is an object and ‘Mars’ is a proper name, and these facts correspond to each other; but it is hard to imagine Frege saying that Mars is an object *because* ‘Mars’ is a proper name. Instead, Frege takes the syntactic and epistemological distinctions to be correlated because he is a reformist about syntax: he wants to *design* a formal language whose syntax will *reflect* the distinction between concepts and objects. This is something we can only guarantee by artifice, after careful analysis of the logical relationships between complete thoughts in some scientific domain.

In a logically-ideal language, we use a term like ‘Mars’ as a proper name because that is what allows us to express thoughts about Mars in a way that makes their logical relationships to other thoughts clear and explicit. It can happen that natural languages are already ideal in this respect: perhaps the syntax of ‘Mars’ in English is already well-suited to helping us grasp the relationships between, say, thoughts about Mars and thoughts about Jupiter. But Frege is not committed to the idea that we can always discover the correct category for an expression’s *Bedeutung* in a logically-ideal language based on its syntax in an existing language, as his remarks about the ordinary attributive use of ‘four’ in *Grundlagen* §57 show.

Thus, the first idea is not going to work. We cannot characterize con-

¹⁰Wright, for example, remains unsettled by the possibility that we will not be able to state the criteria in a language-independent way, and that there will consequently be no language-independent characterization of objects. There will then be no such thing as “International Platonism” about numbers, no conclusion that numbers are objects *tout court* (Wright, 1983, p. 63). Unless the criteria can be made more universal, we would only be able to say that numbers are objects for speakers of this or that language.

cepts and objects by the fact that they are the *Bedeutungen* of concept words and proper names except with respect to a logically-ideal language, and we are not in a position to call any language logically ideal before we have done the work of analyzing the thoughts it can be used to express. To decide whether numbers play the role of objects or concepts, we have to ask which way of analyzing and expressing thoughts containing numbers makes the logical relationships between them clear. It is not sufficient to show that 'four' occurs in a particular syntactic category in an English sentence; we also need to know whether that sentence has a syntax suitable for representing those relationships.

What about the second idea, that objects have identity conditions while concepts do not? In the famous passage in *Grundlagen* §62, Frege writes:

How, then, are numbers to be given to us, if we cannot have any ideas or intuitions of them? Since it is only in the context of a proposition that words have any meaning, our problem becomes this: To define the sense of a proposition in which a number word occurs. But we have already settled that number words are to be understood as standing for self-subsistent objects. And that is enough to give us a class of propositions which must have a sense, namely those which *express our recognition of a number as the same again* ... In doing this, we shall be giving *a general criterion for the identity of numbers*. (Frege, 1884/1980, §62, my emphasis)

Frege is saying that it is a necessary condition on our using words as proper names of objects that they can figure in a special kind of statement. This is a statement that expresses our recognition of the object as 'the same again', the judgment that the object is identical with itself. Elsewhere, Frege claims that concepts cannot stand in the identity relation (though they can stand in the analogous relation of being co-extensive). So it appears to be necessary and sufficient for something to be an object that it makes sense to judge it to be the same as, or different from, another object. More succinctly, objects, but not concepts, have identity conditions. If we could establish that numbers have identity conditions, we would establish that they play the role of objects.

This thought is less helpful than it seems. The trouble here is parallel to the trouble with the first idea: we have no *prior* grip on recognition

judgments or the concept of identity, just as we have no prior grip on the syntactic categories of proper names and concept words. As Wright (1983, p. 55) emphasizes, the only plausible way to pick out the concept of identity is by its logical behavior, as a predicate R such that for any terms X and Y and any context Φ , $R(X, X)$ is true, and if $R(X, Y)$ is true, then $\Phi(X)$ is true if and only if $\Phi(Y)$ is true. The problem is that this criterion does not distinguish identity from the analogous relation of co-extensiveness for first-level concepts, unless we already know whether X and Y are proper names or concept words. And as I have just argued, we won't know that prior to logical analysis of the thoughts we can express in the language.

This should not be surprising in light of Frege's doctrines of the priority of the proposition and multiple analysis. Frege explicitly holds that some judgments can be analyzed *both* as identities and as non-identities. In *Grundlagen* §64, for example, Frege asserts that the thought expressed by "line a is parallel to line b " is the same as that expressed by "the direction of line a is identical with the direction of line b ". Thus, there is no fact of the matter about whether such judgments are 'really' identities or not; for Frege, a judgment is not a recognition judgment independently of a particular analysis into objects and concepts. This means that we will not be able to use the idea that objects have identity conditions to characterize the role of objects in a way that helps us decide whether numbers play the role of objects in the thoughts expressed by statements like "The number of Jupiter's moons is four". We cannot decide whether this statement expresses an identity prior to deciding whether four is an object, or 'the number of Jupiter's moons' and 'four' are proper names.

So both ideas fail to tell us much about how to apply Frege's distinction between concepts and objects in the troublesome cases, and I am not aware of anything Frege says that might be more helpful. The situation we are in is something like that of an unfortunate physiology student, who has learned that the liver produces bile and the gall bladder stores it, but when dissecting a specimen finds herself unable to say which organ is which. She knows they are different and functionally related; but without some additional information, she will not be able to apply what she knows to decide whether *this* bilious organ is a liver or a gall bladder. Likewise, we still need some additional information that will help us apply the concept-object distinction to the case of numbers. I turn now to deriving this additional information from game-theoretical semantics

and the concept of an investigation.

5.2 THE INVESTIGATORY CONCEPTION OF CONCEPTS AND OBJECTS

The discussion in the previous section has given us some desiderata for an explication of Frege's concept-object distinction. To align with Frege's understanding of the distinction, any interpretation of it should at least do the following:

1. It should interpret the distinction between objects and concepts as a distinction between two *epistemological roles*, or between two ways that things are apprehended in thought.
2. It should explain why the distinction is *exclusive*, and the two roles are complementary.
3. It should preserve the idea that objects and concepts are *objective*, because they represent two roles in objective (truth-evaluable, scientific) thoughts.
4. It should be *intuitive*, allowing us to say that things like planets or persons, which are 'objects' in the ordinary physical sense, can play the role of objects in the official epistemological sense.
5. It should be *applicable*, in the sense that it helps us see whether something is playing the role of an object or a concept, and why. In particular, it should be applicable even when we do not know whether our existing language for talking about that thing is logically ideal.
6. Finally, it should be *analytical*, respecting the doctrines of the priority of the proposition and multiple analysis. In particular, it should preserve the idea that pairs of sentences like (1) and (2) can express the same thought, independently of its analysis into objects and concepts.

I now want to argue that an interpretation of the distinction based on the concept of an investigation can satisfy all of these desiderata. I therefore call this interpretation the *investigatory* conception of concepts

and objects. Because it satisfies these desiderata, the investigatory conception agrees with Frege's own. It also improves on his explanation of the distinction, because it provides us with an account of what it means to be an object that we can apply to the case of numbers. According to this account, there is a clear sense in which numbers are objects. We may therefore agree with Frege, and accept the first-order analysis of (1).

I will first explain the investigatory conception of concepts and objects using game-theoretical semantics, which was introduced in Chapter 4. I will then argue that this interpretation satisfies the first four desiderata: it conceives of concepts and objects as epistemological roles, and makes the roles exclusive, objective, and intuitive. It requires more work to show how this interpretation applies to the case of numbers, and that it is analytical, so I will discuss those issues separately.

5.2.1 *Concepts and objects in game-theoretical semantics*

To see how we can derive an interpretation of the concept-object distinction from the game semantics, recall the rules for the quantifiers in the game for a sentence S played on model M :

G.∃. If S is of the form $(\exists x)S'$, the move is made by the Verifier. She goes to the domain of individuals in M , chooses an individual, and gives it a name ' a '. The name is substituted for ' x ' in S' . The resulting sentence is called ' $S'(a/x)$ '. The game continues in the next round with respect to $S'(a/x)$.

G.∀. If S is of the form $(\forall x)S'$, the same thing happens as in G.∃, except that the individual is chosen and named by the Falsifier.

These two rules connect players' moves in the game with the objects in the domain of M , and they are the only two rules that do so. Thus, they characterize the role of objects in the semantic game. Essentially, I propose that we take this role as *constituting* what it means to be an object. An object is what a player searches for and finds in the domain of a model, and specifies with the name ' a ', when making a move in accordance with one of these rules.

Coordinately, if an object is what a player searches for and specifies, a concept constrains such a search. It is what determines whether a particular choice of an object is strategic or unstrategic in the game. This is an

appropriate characterization of concepts because of the relationship between strategic play and truth. To see this, notice that a formula with one free variable will express a concept. In the game for $\exists x\phi(x)$, the choice of object that the Verifier specifies with ' a ' is strategic just in case ' $\phi(a)$ ' is true in the model.¹¹ Thus, ' ϕ ' expresses a constraint on her search for this object, so it signifies a concept in the sense I have identified. We may therefore say that an expression for a concept is true or false *of* an object, and that a concept is predicative.

We can also say that a player *grasps* a concept to the extent that she knows how to play strategically during quantifier moves. To play strategically, the Verifier must choose an object that forces the game to end with an atomic sentence which is true in the model, regardless of what other choices the Falsifier makes in the meantime. These considerations constrain her search in the sense that they determine what she *should* choose. But they also determine what she *will* choose—provided that she is playing to win, and can recognize her choices as complying or failing to comply with them. Insofar as she understands the constraints, she will be guided in her search for an object and will know how to play strategically.

Characterizing objects and concepts in terms of their roles in the semantic game suffices to characterize them more generally, in terms of their roles in investigations structured by which-questions. I developed the ideas which are crucial for this characterization of objects and concepts in Chapter 4. Recall that an investigation is an activity structured by a question, and that we can view an investigation structured by a which-question from two perspectives: it is an attempt to answer this question, but also an attempt to find and give the value of a variable (or some values of some variables). One answers a which-question *by* finding and specifying that value.

These two perspectives set up a correspondence between investigations and instances of the semantic game. On the one hand, quantifier moves in the semantic game can be seen as investigations. Making a strategic choice during a quantifier move requires the player to undertake an investigation. For example, to make a strategic move in accor-

¹¹Again, the Falsifier's case is symmetric: the choice of object that the Falsifier specifies with ' a ' in the game for $\forall x\phi(x)$ is strategic just in case ' $\phi(a)$ ' is false in the model. Because of the symmetry, I leave the Falsifier's role aside in what follows.

dance with $G.\exists$ for a formula of the form

$$\exists x\phi(x)$$

the Verifier should undertake an investigation structured by the question, which objects are ϕ ? On the other hand, for any investigation structured by a which-question, there is a corresponding instance of the semantic game. A which-question has the general form, which F s are ϕ ? This question is exactly the question that structures the investigation the Verifier should undertake to determine her initial moves in the game for

$$\exists \bar{x}\phi(\bar{x})$$

on a model whose domain consists of the F s.¹² So for every quantifier move in the semantic game, there is a corresponding investigation structured by a which-question, and for every investigation structured by a which-question, there is a corresponding instance of the semantic game. Thus, to characterize objects in terms of their role in the game just is to characterize them in terms of their role in investigations. Something plays the role of an object when it can be sought, found and specified in an investigation. Something plays the role of a concept when it constrains such a search.

So that is how to derive an investigatory conception of the concept-object distinction from game-theoretical semantics. Let's see now why this interpretation makes sense, and satisfies the desiderata laid out above.

First of all, notice that like Frege, the investigatory interpretation understands the concept-object distinction in functional terms. What it means for something to be a concept or an object is for it to play a certain role in investigations. We have characterized concepts and objects in terms of how they are used in a certain kind of activity, not in terms of what kind of entities they are.

¹²I have here used the notation \bar{x} as an abbreviation for x_1, \dots, x_n . This notation is convenient because it permits both collective and distributive readings for the predicate ϕ ; we may read the formula $\exists \bar{x}\phi(\bar{x})$ either as $\exists x_1, \dots, \exists x_n(\phi(x_1) \wedge \dots \wedge \phi(x_n))$, requiring the Verifier to select n objects, *each* of which is ϕ , or as $\exists x_1, \dots, \exists x_n\phi(x_1, \dots, x_n)$, requiring her to select n objects that *together* stand in the relation ϕ . The former case can be viewed as a series of individual 'nested' investigations; but it is also a special case of the latter, which is best viewed as a single investigation requiring multiple specifications of objects, like an algebra problem with multiple unknowns.

The distinction between the roles is an epistemological distinction, as I argued that it was for Frege. An investigation is an activity characterized by its epistemological structure. The distinction between concepts and objects corresponds to the epistemological distinction between questions and answers, between an investigator's different epistemic states at the beginning and end of an investigation. To begin an investigation, you must understand the question that structures it. That is, you must understand the constraints on its possible answers, which is to say that you must grasp a certain concept. To conclude an investigation, you must select among these possible answers, which is to say that you must specify which object you take to answer the question. Thus we may say that the concept-object distinction concerns how we apprehend things in thought. It is not a division between things in virtue of what they are in themselves, but in virtue of their possible epistemic relations to we who carry out investigations and can play the semantic game. This shows that my interpretation satisfies the first desideratum.

The epistemological structure of an investigation also implies that the two roles are exclusive, and complementary, within a given investigation. A concept is something known, which constrains and guides players as they search for an object; an object is something unknown, which the search aims to make known. Searching for an object, as opposed to merely guessing or casting about at random, is a deliberate, goal-directed process. An investigator cannot search without some idea of what she is looking for, that is, without knowing how to tell when the aim of her search is fulfilled. On the other hand, a player cannot search for something she already knows or grasps.¹³ So something playing the role of a concept cannot at the same time play the role of an object within a given investigation; the two roles are exclusive. But the structure of an investigation always involves both, so the two roles are complementary. This takes care of the second desideratum.

Concepts and objects are objective under the investigatory conception. Their objectivity derives from the fact that the game semantics, like any logical semantics in the Fregean tradition, is centered around the concept of truth. Claims are objective insofar as they admit of being true or false, that is, insofar as they have truth conditions. The game semantics describes the truth conditions of claims using the notion of a win-

¹³Readers of Plato will recognize the Meno problem here.

ning strategy in the game. We have characterized concepts and objects by their roles in that game; so we may consider them objective insofar as they form part of the conditions under which a claim is true or false. At any rate, it is no more (or less) doubtful that objects and concepts are objective as characterized by the game semantics than as characterized by a standard Tarski-style semantics, since the two styles of semantics are extensionally equivalent.

The game semantics also allows us to describe both objects and concepts as objective in the broadly Kantian terms underlying Frege's understanding of objectivity in the *Grundlagen*. Objects are objective in the sense of being intersubjective, both at the level of the language and at the level of a model or interpretation. The domain of the model is shared in the game; there are no objects that one player can choose but the other player cannot. Likewise, the canonical expressions used to specify objects are shared. A player does not count as making a choice of object unless she uses an expression which makes it clear to the other player which object she has chosen. This precludes private, subjective mental entities from being objects, and precludes using a private language to specify objects. Thoughts are objective because they are expressed in a shared language that speaks about objects in a shared world.

The objectivity of concepts might raise more worries. I have spoken of concepts as 'guiding' a player's search for an object, and this might make it seem as if concepts exist only in the minds of the players. I have been careful to qualify this language for this reason. According to the investigatory conception, a concept determines what is *correct* for a player to choose, in the sense that it determines what is strategic. Correctness, like truth, is an objective notion. A concept only 'guides' a player, in the sense of determining what she actually *will* choose, insofar as she grasps the concept and understands it as constraining what she should choose. The concept is independent of any particular player's grasp of it, or ability to be guided by it when searching. Thus, both objects and concepts are objective, satisfying the third desideratum.

Finally, it is clear that the role of objects is 'intuitive' in the sense I laid out above: ordinary physical objects can play the role of what the players seek, find, and specify in the game. We have already seen examples of this, in Chapter 4. The domain of a model may consist of physical objects in some part of the physical world, like pieces on a chessboard, or planets in the solar system, because it is possible to search for one of

these things, and to say which of them one is choosing. But so long as we can conceive of investigations that are not physical explorations, we do not have to think that all objects are physical (or spatial, or ‘actual’). Anything that can be sought, found and specified in an investigation—anything we might speak of when answering a which-question—plays the role of an object, according to the investigatory conception.

5.2.2 *Numbers as objects*

These observations at last put us in a position to answer my opening question: are numbers objects? That is, do they play the object role in scientific thought? According to the investigatory conception, the answer is a resounding *yes*.

Numbers are a clear example of something we seek, find and specify in investigations. Indeed, they are a paradigm case. In Chapter 4, I introduced the concept of an investigation through examples of elementary algebra problems. The practice of solving such problems makes a distinction between what does and does not count as finding and specifying a number. An equation like

$$x = \frac{1}{2}$$

can be used to specify a number and thereby provide the solution to a problem, because it is in canonical form. Equations like

$$4x^2 - 4x + 1 = 0$$

or

$$2x - 1 = 0$$

do not specify numbers; that is why they are used to give problems, not solutions. Elementary algebra can even be axiomatized into a deductively-complete first-order theory (Tarski, 1967), which can be interpreted via the game semantics. So it hardly seems possible to doubt that numbers play the role of objects according to the investigatory conception.

It follows that a first-order analysis of specificational sentences like

- (1) The number of Jupiter’s moons is four.

is correct and appropriate on the investigatory conception. Like an algebraic solution statement, this sentence specifies a number. A first-order

analysis requires that numbers are objects, but we are now taking objects to be *defined* as what can be specified at the conclusion of an investigation. So we may say that ‘four’ stands for an object in (1), just because it is doing the work of specifying a number. Under the investigatory conception of objects, the first-order analysis of this sentence simply makes explicit what we established by other means in Chapters 1–3: that ‘four’ answers the which-question left open by ‘the number of Jupiter’s moons’, a question we might otherwise express as “Which number is the number of Jupiter’s moons?” or, more colloquially, “How many moons does Jupiter have?”

Our puzzle, though, was that canonical expressions like ‘four’ do not appear to belong to a grammatical category of names or singular terms. They instead belong to the category of adjectives, or determiners, or some other syntactic category in natural language. The virtue of the investigatory conception of objects is that this does not matter. We do not decide whether something plays the role of an object by looking at the syntax of the expression used for it in natural language. Instead, we ask whether it is the kind of thing that can be sought in an investigation, and whether that expression is canonical. That is, we ask whether our expression for it is being used to indicate which thing it is, in a way that is not subject to a reply of ‘But which do you mean?’. For the purposes of logical analysis, we might design a language where all such expressions belong to the same syntactic category, but we have no reason to assume that the syntax of natural language is already structured that way. Fortunately, on the investigatory conception, we have no need to make that assumption in order to apply the concept-object distinction. To conclude that the number four plays the role of an object, we need only observe that ‘four’ is a canonical expression for it. ‘Four’ can be used to indicate a choice of the number four, and thereby complete an investigation structured by a ‘How many?’ question in a sentence like (1). This is so even if ‘four’ is an adjective or a determiner or something else as far as syntax is concerned.

It is worth stressing how different this approach is from the approach of neo-Fregean logicians like Wright (1983) and Hale and Wright (2001). Both approaches agree with Frege that numbers play the role of objects, but they take very different routes to that conclusion. Wright’s route requires a particular understanding of the syntax of expressions for numbers, while the route provided by the investigatory conception of objects does not. This means the investigatory interpretation of the concept-

object distinction is more readily applicable to the case of numbers. Thus, it satisfies the fifth desideratum better than Wright's interpretation, which is the main alternative available in the literature.

Let me explain this point in detail. On Wright's interpretation, Frege holds the thesis of the priority of syntax over ontology. According to this thesis, the category of singular terms is explanatorily prior to the category of objects; what it means to be an object should be explained in terms of what it means to be a singular term, not the other way around. He writes:

The really fundamental aspect of Frege's notions of object and concept is that they are notions whose proper explanation proceeds *through* linguistic notions. . . . For Frege it is the syntactic category [of singular terms] which is primary, the ontological one [of objects] derivative. It is because Frege holds this primacy of syntactic categories that he believes that he can legitimately argue that the syntactic behavior of numerical expressions immediately settles that numbers are if anything a kind of object. (Wright, 1983, p. 13)

On this way of conceiving of objects, to be an object is to be the referent of a singular term. Thus, we explain why numbers are objects by showing that our expressions for them are singular terms, and that they occur in certain true sentences. Because the sentences are true, these numerical expressions refer, and because they are singular terms, their referents are objects.

I explained above why I don't think we can read Frege as treating the category of singular terms as *prior* to the category of objects, though he does think the two categories are in correspondence in a logically-ideal language. My present point is that Wright's interpretation faces a problem in reaching the conclusion that numbers are objects which the investigatory interpretation does not. On Wright's approach, the syntax of the language we use to talk about numbers is fundamental. Under the syntactic priority thesis, we cannot apply the concept-object distinction in the case of numbers unless and until we have an account of the syntax of number words. Wright never really doubts that (1) has the syntax of an identity statement in which 'four' appears as a singular term, but that is exactly what comparing (1) with other specificational sentences

leads us to doubt. It appears that ‘four’ is an adjective or determiner in that sentence, just as much as in “Jupiter has four moons”. And as long as *some* of our uses of number words do not appear to be singular terms, then Wright’s strategy for arguing that numbers are objects is threatened. Someone who adopts the syntactic priority thesis either must explain that those uses really do count as singular terms after all, or explain why those uses can be ignored when applying the concept-object distinction.

The investigatory conception of objects does not take syntax as fundamental, so it does not face this worry. Instead, it takes investigations as fundamental. We recognize numbers as objects by recognizing them as the kind of thing we can search for and specify. It is obvious that we do search for and specify numbers; so it is obvious that they play the role of objects. A great advantage of the investigatory conception of objects is that this conclusion is not beholden to changes in the syntax of the language we use to talk about numbers, or changes in our understanding of that syntax. It is only beholden to our practice of asking and answering a certain kind of mathematical question, the kind of question asked in elementary problems in arithmetic and algebra.

5.2.3 *Specification without syntax*

I have argued that under the investigatory conception, numbers play the role of objects, and the roles of concepts and objects are exclusive. Does this mean that numbers do *not* play the role of concepts? Not in general.

We have reached a vantage point from which we can apply the distinction between concepts and objects independently of the syntax of natural language. To decide whether something is a concept or an object, we ask about its role in thought. Under the investigatory conception, this question becomes: is it something we search for and specify in investigations, or is it something that constrains such a search? But now that we have learned to think of concepts and objects in terms of their roles, it would be better to ask this question in a slightly different way. We should not presuppose that things always play the same roles across different thoughts. Rather than asking whether numbers or anything else play the role of concepts in *thought*, we should ask whether they ever play the concept role in *a* thought. I will argue that they do.

To reach that conclusion, we have to revisit Frege’s doctrines of the priority of the proposition and multiple analysis. We saw earlier that

Frege thinks of concepts and objects as having been obtained by splitting up or analyzing a thought, and that this is something we are free to do in different ways, so something does not play the role of a concept or an object independently of how we analyze the thought in which it occurs. I called this the argument from analysis. Given these doctrines, my argument is simple: in general, it is possible for one and the same thing to play the object role under one way of splitting up a thought, and the concept role under another way. This is no less true for numbers than for other kinds of things. Thus, by seeing how to accommodate the doctrine of multiple analysis under the investigatory conception, we will be led to see why numbers can play both roles. The only question that will remain is how we can best represent this insight in a language of analysis like that of the *Begriffsschrift*, a language which is designed to express the relationships between different ways of splitting up a thought.

First of all, to narrow our focus, let's make a distinction. Frege gives several different kinds of examples of how a thought can be analyzed in multiple ways. Here are a few:

- (6) Passivization (Frege, 1879/1967, §3)
 - a. The Greeks defeated the Persians at Plataea.
 - b. The Persians were defeated by the Greeks at Plataea.
- (7) Clause nominalization (Frege, 1879/1967, §3)
 - a. Archimedes perished at the capture of Syracuse.
 - b. The violent death of Archimedes at the capture of Syracuse is a fact.
- (8) Abstraction (Frege, 1884/1980, §64)
 - a. Line a is parallel to line b .
 - b. The direction of a is identical to the direction of b .
- (9) Concept-extension conversion (Frege, 1892/1997e, pp. 187–188)
 - a. There is at least one square root of 4.
 - b. The concept *square root of 4* is realized.
- (10) Specification (Frege, 1884/1980, §57)
 - a. Jupiter has four moons.
 - b. The number of Jupiter's moons is four.

Frege's normal way of explaining the relationships in pairs of sentences

like these is that the two sentences express the same thought. There is a common thread that runs through them: in each pair, there is something the same (the thought, the content, the truth-conditions) but also something different (the expression, the emphasis, the ‘coloring’). The sentences in each pair are equivalent for certain purposes, which we can broadly characterize as logical or scientific purposes. They are not equivalent for others, such as communicative or rhetorical purposes.

I will focus on the relationship in the last pair. It is clear from the text of the *Grundlagen* that Frege thinks of this relationship as comparable to the one in (8), which he characterizes by saying “we carve up the content in a different way”; so it is clear that Frege thinks these sentences represent or correspond to two different ways of carving up a single thought. Still, it is not obvious that the relationships in all five pairs are instances of the same idea, or the same kind of intellectual process.¹⁴ For this reason, I will be satisfied to explain why the sentences in the last pair count as two ways of expressing the same thought under the investigatory interpretation. I hope what I say might eventually shed light on the other examples, but I leave that task for another occasion.

So we have to explain why ‘specificational recarving’ counts as expressing the same thought in a different way. This is straightforward, given the strategy of analysis for specificational sentences I developed in Chapter 4, and the interpretation of the concept-object distinction I have offered here. The explanation centers on two ideas. The first comes from Frege’s discussion of analysis in *Begriffsschrift* §9. As we saw above, Frege’s model for splitting a thought into a concept and an object involves regarding a part of that thought as variable, while holding the remainder fixed; it involves seeing that part as one among several possibilities in a range. The part that we conceive as variable plays the role of an object. The part we hold fixed plays the role of a concept.

The second idea comes from the argument in Section 5.2.2. As we saw there, numbers play the role of objects because they are the kind of thing we can seek, find and specify. In fact, there was nothing special about the case of numbers. In general, when we express a thought using a specificational sentence, what is specified in the complement is playing the role of an object, precisely *because* it is being specified. This just follows from the role of specificational sentences in investigations. It is more or less

¹⁴See Dummett (1991, Ch. 14) for further discussion of this issue.

true by definition, under the investigatory conception of objects.

Both ideas are implemented in the strategy of analysis for specificational sentences I developed in Chapter 4. It's easiest to see all this in the context of a particular example, and it will be helpful to shift away from the example of numbers for a moment. Instead, I will work with the example of colors. Consider the thought we might express as:

(11) Io is yellow.

One way to carve up this thought is to regard Io as variable, thereby assigning it the role of an object. Suppose we are thinking of it as part of the range of Galilean moons (the others being Europa, Ganymede, and Callisto). Since Io is the only yellow moon in that range, the following specificational sentence expresses this way of analyzing the thought:

(12) The Galilean moon which is yellow is Io.

This sentence presents the thought as one in which Io is the object selected and specified in answer to the question, which Galilean moon is yellow? Under the strategy of analysis given in Chapter 4, this becomes:

(13) $\exists!x(x \text{ is yellow} \wedge x = \text{Io})$

where we think of the Galilean moons as forming the domain of quantification. Here we have used the syntax of a first-order language to represent the process Frege describes in *Begriffsschrift* §9, by replacing 'Io' with a variable with a certain range. By regarding Io's position as variable, we take the concept in the thought as that of being yellow, contrasting the possibility that Europa or Callisto is yellow with the fact that Io is.

The formula in (13) is the same as the one in (11) in the sense that they are logically equivalent. On the other hand, as I observed in Chapter 4, they are *practically* non-equivalent. Under the game semantics, (13) is true just in case the Verifier has a non-trivial winning strategy in a game of at least two rounds, while the game for (11) is trivial. Thus, a first-order language under the game semantics enables us to explicate the sense in which (12) is a recarving of (11). They express the same thought, but in different ways; they are logically equivalent, but not equivalent for other purposes.

This strategy of analysis allows for *multiple* analysis because exactly parallel reasoning applies to the color. Suppose, in the thought expressed by (11), we instead regard yellow as variable, thinking of it as part of a range of colors including red, silver, and so on. We thus regard yellow as an object; the concept is that of being Io's color. We can express this way of analyzing the thought with a different specificational sentence:

(14) The color of Io is yellow.

This sentence presents the thought as one in which the color yellow is the object selected and specified in answer to the question, which color is Io? Represented in a first-order syntax, this becomes:

(15) $\exists!x(\text{Io is } x \wedge x = \text{yellow})$

where the domain of quantification consists of colors. Again, under the game semantics, this formula is logically equivalent to (11), but practically non-equivalent. So the first-order analysis of (14) again vindicates the idea that it represents a recarving of the thought expressed by (11). And since (13) and (15) differ in both their syntax and their semantics—for example, a winning strategy for the Verifier will involve choosing a moon as the value of x in the game for (13), but a color in the game for (15)—they represent different analyses of the same thought.

Admittedly, the formula or quasi-formula in (15) will look strange to anyone trained in the Fregean hierarchy and the way it is standardly applied to natural language. Haven't I made a mistake by placing a first-order variable in the position of 'yellow', which is a grammatical predicate? This objection is the central challenge which I have struggled against in this chapter. My answer is no, the first-order analysis is correct and entirely appropriate. But it is only now, with significant interpretive work behind us, that we are in a position to see why. First-order variables represent the role of objects, as opposed to concepts. Now that we understand those roles in epistemological terms, we can see that one and the same color may play the object role in addition to the concept role, depending on how we carve up the thought in which it occurs.

This point emerges vividly from examples (12) and (14). Frege would regard these sentences as expressing two different ways of splitting up one thought, which is the same thought as (11) expresses. But it is obviously the very same color which is specified in (14) that is predicated

of the moon in (12). If it were not, one or the other sentence's equivalence with (11) would appear to rest on a fallacy of equivocation. The word 'yellow' does not pick out or correspond to different *things* in these sentences; rather, we use these different sentences to express different epistemological relationships to a single color.

Moreover, under the investigatory interpretation of the concept-object distinction, it is easy to see *why* a color can play both roles. It is possible to search for something on the basis of its color, but equally, it is possible to search for the color of something. That is exactly what will happen in the semantic games for (13) and (15). In the first game, the Verifier must search for an object among the Galilean moons, and her choice will be strategic just in case the moon she chooses is yellow. But in the second, the Verifier must search for an object among the colors, and her choice will be strategic just in case Io is the color she chooses.

The conclusion that yellow may play both the role of a concept and the role of an object may still make some readers uneasy. I will reply to some objections in a moment. But before we consider them, notice that the argument I have just given for the case of colors can be repeated, *mutatis mutandis*, for the case of numbers. Just like "Io is yellow", the thought expressed by "Jupiter has four moons" can be carved up in different ways, including both

(16) The planet with four moons is Jupiter.

and

(1) The number of Jupiter's moons is four.

And just like colors, it is possible both to search for a number, and to search for other things on the basis of that same number. When we ask "How many moons does Jupiter have?", and answer with (1), we are searching for and specifying a number, and thus treating it as an object. But it makes just as much sense to ask "Which planet has four moons?", and answer with (16). In that case, we are using the number as a constraint on a search for a planet, and thus treating it as a concept. Thus numbers, like colors, are things which can play both roles, and we ought to adopt a language of analysis that is capable of representing this idea. Frege was right that numbers most often play the object role in our usual ways of carving up mathematical thoughts. But I see no reason why that

excludes numbers from playing the role of concepts under other analyses, such as the one expressed by (16). If Frege believed that numbers must *always* play the object role, I do not see that he gave any good reason to follow him in that respect, in *Grundlagen* §57 or elsewhere.

The investigatory conception gives us room to expand this conclusion beyond numbers to other things, too. Reasons or ways, for example, can be said to be objects in the same investigatory sense that we have found for numbers and colors.¹⁵ Our practices of seeking, finding and specifying are just as determinate in the case of reasons or ways as in the case of planets, numbers, or colors. These things can all play a common epistemological role, the role of objects in the investigatory sense. But they may also play the role of concepts; and which role they play is not a matter of the syntax of our expressions for them in natural language, but of whether they are being sought or are constraining a search.

Now for the objections. First, how can it be appropriate to use a first-order variable to take the place of ‘yellow’ in (11)? Instead of (15), wouldn’t it be better to analyze the specificational sentence in (14) using a formula like this,

$$(17) \quad \exists!F(F(Io) \wedge \forall x(F(x) \leftrightarrow x \text{ is yellow}))$$

where a *second-order* variable takes the place of ‘yellow’? A second-order variable is a kind of hybrid. The values of second-order variables are like objects insofar as they are alternatives to one another; they are quantified over, and could be chosen and specified during a quantifier move in a semantic game. But they are like concepts insofar as they are predicated of other things. A second-order variable thus signifies both roles in a single form of expression. This may look like exactly what we need in the case of yellow: a kind of expression which stands for something that plays both roles. Perhaps that was part of Frege’s motivation for introducing this kind of variable into his language of analysis.

But it is not quite what we need. The point of using a formal language to represent a particular way of carving up a thought is to exhibit that thought’s logical relationships, which need not follow the syntax of any

¹⁵Consider for example that a way is specified in (4-a), while a reason is specified in (5-a), so they count as objects in the investigatory sense. At the same time, there are other specificational recarvings of (4-b) and (5-b) where they would play the role of concepts.

particular way of expressing the thought in natural language. When we are deciding how to represent the analysis that corresponds to (14) in a formal language, we should not consider the grammatical role of 'yellow' in English, but the logical relationships that this representation can help us see and understand. It is clear that we will want to use *some* sort of variable to capture the fact that in (14), yellow is being presented as one among several colors. This way of carving up the thought presents it as logically related to the thoughts that Io is *red*, or that Io is *silver*, to the general thought that Io is some color or other, and so on. The question is whether the hybrid role that a second-order variable signifies is the right device to make these relationships explicit. I am not convinced that it is.

Just considered as a piece of notation, the second-order form in (17) obscures the symmetry in the two ways of carving up the thought. If we represent the two analyses as (13) and (15), then the first-order notation makes clear that we use the same process to arrive at each analysis from the neutral expression (11): we let one part of that thought vary, and hold the rest fixed. Using a second-order variable, as in (17), introduces a notational distinction between the two cases, but we have no real logical difference for that distinction to express. Colors and moons are both things that we seek, find and specify in language. Our canonical expressions for them have different grammatical properties; but that is entirely beside the point, as Frege would be the first to point out. They also seem to belong to different ontological categories, but I have argued that Frege was not interested in representing this difference in his language of analysis, either.

Using a second-order variable would not just be a benign notational quirk, though. The problem is that a second-order variable seems to be a sign which indicates that its value plays both roles. Specifically, it plays them 'simultaneously', under a *single* way of carving up the thought, since that is what the formal representation depicts. That is not what we need in order to capture what is happening with the color yellow in (12) and (14). What we need is a way to express that something playing the object role in *this* way of carving up a thought can also play the concept role in *that* way of carving it up. We need a notation that is neutral between the two roles, not one that combines them, as second-order variables do.

This is already adequately achieved by the quasi-formal first-order notation in (13) and (15). The important feature of this notation is that,

unlike the second-order notation, the syntax representing predication is orthogonal to the syntax representing variation and specification. Because it uses a distinct sign (the copula 'is') to represent predication, a variable always represents something playing purely the role of an object—something to be specified, at the outcome of an investigation, by means of a canonical expression. Likewise, a canonical expression is neutral between uses which specify something and uses which predicate that thing of something else. For example, yellow is specified in ' $x = \text{yellow}$ ', but predicated in ' $x \text{ is yellow}$ '. In this way, the notation is able to represent the idea that one and the same thing can play both the role of an object and of a concept. That one and the same color can play both roles is captured by the common element 'yellow' in these signs; that it plays distinct roles is captured by the different elements '=' and 'is'.¹⁶

In (15), the expression 'is x ' captures the *potential* for something playing the object role to take on the concept role under a different analysis: placing the variable to the right of the copula expresses the idea that, *given* a specific value for that variable, we may go on to hold that value fixed and let another part of the thought vary, such as the part expressed by the grammatical subject. In other words, if there are things which play the role of objects, but never play the role of concepts, this will be reflected in the fact that variables ranging over them never appear to the right of the copula.¹⁷ The color yellow, because it is specified as the value of the post-copular variable in (15), is not such a thing, but that is no reason to think it does not play the object role.

¹⁶This is not to say that we have to recognize 'is' and '=' as significant parts of these expressions in the sense that they must be treated categorically, as symbols that stand for something. They are just pieces of syntax used to *depict* a certain distinction between predication and specification. They could be eliminated in favor of syntactic conventions that don't treat them as distinct symbols, such as putting a canonical expression in boldface when it is used specificationally, and italics when it is used predicatively. I am therefore not suggesting a departure from Frege's view that the copula is "a mere verbal sign of predication" (Frege, 1892/1997e).

¹⁷This is basically how Aristotle identifies primary substances in the *Categories*: primary substances are those things which are not predicated of anything else. This makes primary substances a special kind of individual, but not the only kind. The notation I am suggesting is therefore capable of making a distinction like Aristotle's distinction between primary substances and other kinds of individuals. An object, or individual, is the value of a variable; primary substances are those individuals in a model that can never be specified as the value of a post-copular variable without making the sentence false.

Still, the objector may press further: doesn't this proposal obliterate Frege's insights? Have I not simply eliminated the idea that a concept is 'essentially predicative' and 'unsaturated', or that "a concept cannot be made into an object without altering it"? In allowing that a color can play both roles, and seeking to accommodate that idea in logical notation, am I not making exactly the proposal that Benno Kerry made, which Frege assailed in 'On Concept and Object'?

Any such objection can be answered, I think, by reminding the reader that the concept-object distinction is not a distinction between ontological categories, but between epistemological roles. Being essentially predicative or 'unsaturated' is not a property of the color yellow insofar it is a certain kind of entity or thing; instead, yellow is predicative or unsaturated insofar as it plays the role of a concept. That is, yellow is essentially predicative *qua* concept, not *qua* color, and the argument from analysis implies that it only plays the role of a concept with respect to a given analysis of a thought. Its role as a concept is adequately represented in the notation I have proposed by the fact that 'is yellow' is an incomplete expression, a predicate suitable for forming a complete truth-evaluable sentence when completed by an expression like 'Io'. Similarly, the notation represents the alteration that occurs between taking yellow as a concept and taking it as an object as the transformation from 'is yellow' to '= yellow'. This is not to be thought of as a mysterious change in the ontological status of a color. It is an ordinary change in our epistemological relationship to that color, the kind of change that occurs between asking "Which of these fruits is yellow?" and "What color is the lemon?"

Here is another way to put the point. The function-argument notation that Frege developed in the *Begriffsschrift* and that has evolved into our modern notation was designed to express his epistemological distinction between concepts and objects, which it does very well. I am not recommending that we abandon the insights which led to it; I am only recommending that we recognize its limitations. Because it is so sharply focused on representing the distinction between two epistemological roles, Frege's notation is inadequate for representing the idea that one thing can play both roles. The reason for this is that Frege's notation represents the distinction between the roles using mutually exclusive syntactic categories. Frege's notation represents yellow in its object role using a term like 'yellow', and in its concept role using a predicate like 'Yellow(ζ)', and recognizes no common element in the semantics of these expressions.

I am proposing that, by adopting a syntax for canonical expressions which is neutral between specificational and predicational uses, we can escape this limitation, and so increase the expressive power of our language of analysis. Though it remains possible to distinguish the two roles, it also becomes possible to represent how a single thing plays one role under one way of carving up a thought, and another role under another carving. The example of colors shows that such a language is intuitively desirable, and well-suited to expressing the different ways we 'specificationally recarve' thoughts in natural language. So what is to stop us from adopting such a language for the purpose of analyzing what sentences like (12) and (14) say? Furthermore, once such a language is available, it will make just as much sense to apply it to the case of numbers, for the reasons given above: like colors, numbers play both roles, and we ought to represent both the distinction between those roles and the common semantic element between them in our language of analysis.

Much more needs to be said, of course, to work these insights into a complete formal system, and show that this system is consistent and adequate for the purposes of logical analysis. I am not worried that this can be done successfully.¹⁸ But I cannot carry out that project here, and must leave it for future research.

So here is where we have ended up. By taking Frege's doctrines of the priority of the proposition and of multiple analysis seriously, as well as his idea that a specificational sentence articulates one way of carving up a thought, we are led to the conclusion that things like numbers and colors can play the roles of both objects and concepts. We have thus avoided the forced choice that puzzled us at the beginning of the chapter: we do not have to assign numbers or colors to one role or the other once and for all; we can instead recognize that they play both, and devise a language of analysis which is capable of expressing this fact. Once

¹⁸Famous last words, perhaps. But I am optimistic because the notation I have used in (13) and (15) to represent the different ways of carving the thought in (11) is not so different from the language of first-order set theory, except that I have used 'is' instead of ' \in ', and made use of canonical expressions other than ' \emptyset '. There is every reason to think that a fully-specified formal language suitable for representing the logical relationship of specificational recarving will have a natural translation into the language of set theory. I expect that a consistency proof for a plausible theory in this language, relative to the consistency of the background set theory, would in most cases probably only require a straightforward set-theoretic construction.

we free ourselves from the idea that we can tell which role something plays in a thought just by looking at the syntax of the expression we use for it in natural language, this conclusion is not surprising or puzzling. Discovering a number's or color's role is a matter of grasping the logical relationships of the thought in which it occurs. The relationships which lead us to assign it to one role or the other may or may not correspond with distinctions between our natural language expressions.

We are able to abandon the tight link between syntax and the concept-object distinction because the investigatory conception of concepts and objects gives us a different way to apply the distinction. Thus, the investigatory conception is the linchpin of the solution to the puzzle. By reorienting our attention, away from syntax and toward epistemology, the investigatory conception gives us a substantive sense in which numbers are objects—but also a substantive sense in which they are concepts. Since it satisfies the other desiderata for a Fregean understanding of concepts and objects, it provides a desirable and correct interpretation of Frege's claim that numbers are objects, one that we can agree with. But since it also implies that numbers may play the role of concepts, few ontological conclusions follow directly from this claim. I will turn to that discussion in Section 5.3, after a brief detour.

5.2.4 *A further consideration: specifying what predicates denote*

In arguing that numbers and colors play both the object role and the concept role, I relied on the idea that when we use words in both ways, we are not using them equivocally. When we use the word 'yellow' as a predicate in (12) and as a canonical expression in (14), we are talking about one and the same *thing*, namely, the color yellow. Similarly, when we use the word 'four' attributively in (16) and as a canonical expression in (1), we are talking about one and the same thing, the number four. This seems to me the most natural and intuitive way of describing these cases.

Still, a Fregean who adopts the Ontic Picture might deny this intuition, and protest that, precisely because these words function as predicates in some cases and as terms for objects others, there *is* no one thing that each is about. That is, we *do* use these words equivocally, or at least polysemously. This 'Ontic Fregean' takes his general theoretical commitments to rule out the intuitive idea, and so he simply bites the bullet:

since concepts and objects are exclusive ontological categories, it simply cannot be the case that when we use, say, 'yellow' as a predicate and as a term, we are in both cases talking about the same thing, a certain color. Perhaps, following Frege (1892/1997e), we can say that we sometimes speak about objects which 'go proxy for' concepts. But strictly speaking, the two different uses never have a common denotation.

In response to this objection, I would like to point out that there is also a good theoretical reason for avoiding the Ontic Fregean's picture. The Ontic Fregean's theoretical commitments prevent him from *ever* being able to specify what predicates denote. This is a position we should avoid if we want to give a systematic semantics, for either a natural language or a formal language of analysis.

I will follow a discussion of this problem for the Ontic Fregean in Rieppel (2013a, 2016). To see the problem, consider the following simple predicational sentence:

(18) Oscar is happy.

What does 'happy' denote in this sentence? A natural answer is that it denotes the property of being happy. Likewise, when we existentially generalize into the position of 'happy', as in

(19) There is something Oscar is.

it is natural to say that the quantifier ranges over properties. Properties therefore seem like the perfect candidates for the denotation of predicative adjectives like 'happy' or 'yellow'.

On the other hand, it is also natural to assume that 'the property of being happy' denotes the property of being happy. But now we have a problem: even if (18) is true, it is obviously incorrect to say

(20) Oscar is the property of being happy.

Substituting 'the property of being happy' for 'happy' does not preserve truth. Thus it appears that 'happy' cannot denote the same thing as 'the property of being happy'. So perhaps 'happy' does not denote a property after all.

For our Ontic Fregean, there is a seemingly straightforward response to this substitution failure. He will explain the substitution failure by distinguishing the kinds of things that the post-copular phrases in these

two sentences denote: while ‘happy’ denotes a concept, ‘the property of being happy’ does not. Instead, it denotes an object, which is a fundamentally different kind of thing. The difference between the kinds of things denoted by the post-copular phrases ‘happy’ and ‘the property of being happy’ make for different kinds of truth conditions for (18) and (20). Since ‘happy’ denotes a concept, and concepts are essentially predicative, (18) is a predicational sentence. But since ‘the property of being happy’ denotes an object, and this object cannot be predicated of Oscar, the ‘is’ in (20) must express identity: this sentence is true just in case ‘Oscar’ and ‘the property of being happy’ denote the same object.

But as Rieppel points out, there is a serious problem with this strategy for explaining the substitution failure. The explanation apparently leaves the Fregean with no way to say what concept-words like ‘happy’ denote at all. For whenever he *tries* to say what they denote, by completing a clause like

(21) ‘Happy’ denotes ...

he will not be able to fill in the blank. If he fills it with a predicative expression, the attempt does not result in a well-formed sentence. In particular,

(22) ‘Happy’ denotes happy.

is ungrammatical and makes no sense.¹⁹ Syntactically, the only way to complete the clause is to fill in the blank with a nominal expression of some sort. But if he fills in the blank with a nominal expression, then by his own lights, the result will inevitably say the wrong thing. Since a nominal expression will denote an object, it will not denote a concept, so it cannot be used to specify the denotation of concept-words like ‘happy’. The Ontic Fregean cannot even say that ‘happy’ denotes the concept of being happy! For like ‘the property of being happy’, ‘the concept of be-

¹⁹Admittedly, with both ‘yellow’ and ‘four’, this problem is less compelling: “‘Yellow’ denotes yellow” and “‘Four’ denotes four” do not feel as ungrammatical as (22), perhaps because these words shift to the object role more readily than ‘happy’ does. This may explain Rieppel’s choice of example. On the other hand, the problem is a general one, and it is arguably even worse in some cases where disquoting the object language expression *does* result in a grammatical sentence, such as with adverbs like ‘occasionally’. For while “‘Occasionally’ denotes occasionally” is well-formed, it has the wrong meaning.

ing happy' is a nominal expression. It therefore denotes an object, not a concept, and gives rise to the same substitution problem as above.

Besides giving his view a whiff of paradox²⁰, this conclusion leaves the Ontic Fregean in a dilemma. Neither a nominal nor predicative expression can be used to fill the gap in (21), and since those are his only options, he has no way to say what 'happy' denotes. The problem clearly arises from the Fregean's adoption of the Ontic Picture: because he assumes there is a strict correspondence between the syntactic categories of expressions and the kinds of things they denote, and that objects and concepts are fundamentally different kinds of things, he is unable to specify the concept he intends.

Accordingly, Rieppel argues that we ought to abandon these assumptions. Instead of distinguishing predicative and nominal expressions at the level of *what* they denote, he proposes that we distinguish *how* they denote. When we distinguish predicative and nominal expressions at the level of how they denote things, we are then free to say that they denote the *same* things, in different ways. For example, we may say that 'happy' and 'the property of being happy' both denote the property of being happy; the difference is that 'happy' *ascribes* that property, while 'the property of being happy' *refers* to it. On Rieppel's proposal, ascribing and referring are two different semantic relations—two different ways that words denote things.

This proposal allows us to specify the semantics of 'happy' while avoiding the substitution problem. 'Happy' denotes the property of being happy, by ascribing it. It ascribes this property to Oscar in (18). But 'the property of being happy' refers to the property rather than ascribing it, and for that reason (20) has the form of an identity statement, rather than a predicational sentence. Since Oscar may have the property of being happy without being identical to it, the substitution isn't truth-preserving. And because this approach to the substitution problem does not bar us from specifying what 'happy' denotes, it leaves us in a better theoretical position than the Ontic Fregean's.

If Rieppel is correct that this is the best solution, there is independent theoretical motivation for a picture very similar to the one I have offered. Rieppel's idea that nominal and predicative expressions can bear differ-

²⁰The paradox here is the same as the one in Frege's original claim that "the three words 'the concept horse'... do not designate a concept" (Frege, 1892/1997e, p. 184).

ent semantic relations to the same things is akin to my claim that some things can play both the concept role and the object role. Where Rieppel would say that properties like being happy, or being yellow, or being four in number, can be both ascribed and referred to, I would say that they can play both roles. Regardless of how we put it, the point is that we need to give up the idea that nominal and predicative expressions denote fundamentally different kinds of things. If we are to get a systematic semantics for predicates off the ground, we should instead think of the distinction as reflecting two different *ways* our words can be related to things. Our words can sometimes be related to one and the same thing in different ways; one and the same thing can play different roles for us in our language and thought.

I don't wish to over-emphasize the parallel. Rieppel's proposal differs from the position I am taking here in several ways. First, Rieppel sees his proposal primarily as a methodological proposal about how to do semantics, and remains neutral about its extra-linguistic grounds. By contrast, I see the semantic proposal as grounded in an epistemological distinction, the distinction between concepts and objects in the investigatory sense. On the other hand, I wish to remain more neutral than Rieppel about whether adjectives like 'happy', 'yellow', or 'four' denote properties. If 'property' is just a piece of technical terminology for *whatever* such words denote, then I am happy to say that they denote properties, though I think we should first establish that their denotations really form a univocal semantic category. But if 'property' is a piece of *metaphysical* terminology, say for an ontological category of universals, then the thesis that properties play both the object role and the concept role represents a particular metaphysical position about our epistemic relationship to universals. This is one possible position about the metaphysics of properties—perhaps an attractive one. Still, I think the investigatory conception of the concept-object distinction is compatible with other metaphysical positions, and doesn't decide between them. I will conclude by spelling out this idea.

5.3 PLATONISM AND THE INDEPENDENCE OF OBJECTS

I concluded above that we can agree with Frege: numbers are objects—or at least, they play the role of objects, given our usual analyses of math-

ematical thoughts. When this claim is interpreted as I have argued it should be, it means that numbers are answers to which-questions. When we say that numbers are objects, we are not saying anything more surprising than that we give them in answer to the questions we ask with 'How many?'. They share this epistemological role with things we call 'objects' in the ordinary, physical sense: numbers are things we can seek, find and specify in investigations. That is the conclusion of this dissertation. Its main work, of articulating and justifying a conception of objects that applies equally to numbers, other abstracta, and physical things, is now completed.

At the same time, I concluded that we have no reason to think that numbers do not also play the role of concepts. A number can be something we seek, find, and specify; but it can also be something that constrains a search for something else. The temptation to deny this stems from the Ontic Picture, from the thought that concepts and objects are two exclusive ontological categories. But the Ontic Picture is not the right interpretation of Frege's theoretical commitments, and it also faces serious theoretical problems. So we should reject this picture, and once we do, nothing bars us from saying that numbers play both roles.

In this final section, as a kind of epilogue, I would like to explore whether there are any further consequences we can draw from these conclusions about the ontology of numbers. I have argued that such consequences are *not* part of Frege's concerns. But they are interesting to consider nonetheless, now that we have a clearer view of the content of his claim that numbers are objects. How far do my conclusions carry us toward an account of the ontology of numbers, and especially toward mathematical realism or platonism, positions that Frege at least sometimes appears to endorse?

I will argue that the investigatory conception permits us to endorse a moderate and commonsense mathematical realism. It does not, however, imply anything like a *platonist* position. As I shall understand it, the distinctive claim of platonism is that numbers are *independent* in an ontological sense. There is nothing in the investigatory conception of objects to encourage us to adopt platonism in contrast to a position on which numbers exist, but not independently, such as an aristotelian²¹ position,

²¹I am using 'platonism' and 'aristotelianism' as labels for these positions because I take them to be inspired by the views of Plato and Aristotle. But I am not making

or even a kind of idealism.

Let me start by saying what I think we clearly can say in the direction of realism. First, *are* there numbers? Yes, of course: the previous sentence contains four words; so *there is* a number of words belonging to the previous sentence, namely, the number 4. We make claims which existentially generalize over numbers, and by specifying the numbers which witness those existential quantifiers, we show that such claims are true. Nothing more is required to see that there are numbers, that numbers exist. The investigatory conception of objects vindicates this most basic realist claim about numbers because it takes such existential generalizations at face value, and assigns them truth conditions which obviously obtain.²²

Furthermore, it is clear that we make claims *about* numbers, and that such claims are *objective*. For example: the positive square root of 16 is even, and it divides 64. In making this claim, I am making a claim about certain numbers. (What else could it be about?) Specifically, I am making a claim about the number 4, because the positive square root of 16 is 4. This shows that the claim is about a number in the same sense in which any singular predicational sentence is about what its pre-copular phrase denotes: when I specify which object it denotes, I give a number. Furthermore, the claim is objective, just as claims about physical objects are objective. Indeed, one of the desiderata I gave for the investigatory conception was that concepts and objects in general are objective. Numbers follow suit, for the reasons I gave above. Numbers are objective in the sense that they are intersubjective, part of a shared language and shared world. And whether they play the object role or the concept role, the claims we make about them have truth as a standard of correctness: they have truth conditions and truth values; we offer evidence for their truth, and retract them if they are proven false.

In short, thinking of numbers according to the investigatory conception yields a moderate mathematical realism. Nothing I have said challenges the theses that numbers exist, that we talk about them, and that

the historical claims that Plato was a platonist or that Aristotle was an aristotelian; so I leave the labels uncapitalized.

²²Thus, as I am understanding it here, *moderate* mathematical realism takes an 'internal' reading of existential quantification over numbers, in the sense of Carnap (1956). The platonist, aristotelian, and idealist positions I discuss below go beyond moderate mathematical realism by saying more about the truth conditions of such quantifications from an 'external' point of view.

our claims about them are objectively true or false. I see no reason to deny those theses, and every reason to endorse them.

And yet. One might still feel there is a nagging question to be answered: do numbers *really* exist? Are they real in the same way that planets and paramecia and protons are? That is, do they exist independently, somewhere 'out there', independent of and uncreated by our thought and talk about them? Mathematical platonism, as I shall understand it here, is a position that gives affirmative answers to these questions. It goes beyond moderate realism in insisting that while numbers are abstract rather than concrete, they are nevertheless *beings* in exactly the same sense that concrete objects are. Although numbers are not in space and time, they exist, they are real, in just the same sense as planets and protons. We didn't create or imagine them; we only come upon them already existing, and discover facts about them which would still hold true if they remained undiscovered.

What's distinctive of platonism is a certain way of understanding what makes for objective truth. For the platonist, the objectivity of a true claim is tied to the *independent* existence of the objects that claim is about. When I say "Jupiter is a gas giant", what I say is true because Jupiter makes it true; but it is objective because Jupiter is an object with an independent existence. Such objective claims contrast with subjective claims (like, perhaps, "I am in pain") on the one hand, and fictional claims (like "Sherlock Holmes was a detective") on the other, in that they are about independently existing things. Because Jupiter is a thing 'out there', it would exist, and it would be a gas giant, whether or not anyone knew about that planet or knew it was a gas giant; not so for pains or fictional characters. Similarly, when I say that "Four is even", what I say is true because the number 4 makes it true, and it is objective because that number exists independently, and would be an even number whether or not anyone knew it.

The investigatory conception of numbers as objects is compatible with the platonist's understanding of objects as independently existing things. Certainly, any example that one could give of an independently existing thing would count as an object under the investigatory conception. The issue between the moderate realist and the platonist is just about the converse claim: if something counts as an object in the investigatory sense, does it exist independently? Here, I think, we can reasonably answer in the negative. There is room, under the investigatory conception, to hold

that not everything which plays the role of an object exists independently. In particular, there is room to hold that while numbers are objects in the investigatory sense, they do not exist independently. So while mathematical platonism is compatible with the investigatory conception of objects, it does not follow from it.

What does it mean for something to exist ‘independently’? I will consider two ways of understanding this idea, and two corresponding ways of rejecting mathematical platonism: a kind of aristotelianism, and a kind of idealism. Both positions, like platonism, are motivated by a desire to preserve realism about the truth of mathematical claims, while rejecting the platonist’s strong ontological stance about numbers. And both positions are compatible with thinking of numbers as objects in the investigatory sense.

First, aristotelianism. We have seen that the platonist spells out his position by drawing an analogy between numbers and concrete things. The platonist’s models for independent existence are things like stars, or icebergs, or lemons. These are things an aristotelian would consider independent because they are *primary substances*. Aristotle himself describes primary substances in terms of an asymmetry in the predication relations among things: “a *substance*—that which is called a substance most strictly, primarily, and most of all—is that which is neither said of a subject nor in a subject, e.g. the individual man or the individual horse” (Aristotle, 1963, 2a11 ff.). In more modern terms, primary substances are those things which are not properties of anything else; everything else, rather, is ultimately a property of them. For an aristotelian, this means that other things are dependent on primary substances: the existence of properties depends on the things they are properties of. Without the lemon, for example, its particular color would not exist. Primary substances are independent in the sense that they don’t depend on anything else in this way.

Thus, on the aristotelian way of understanding independence, the platonist’s claim that numbers are independently existing beings amounts to the claim that they are primary substances. From the aristotelian perspective, this claim is too strong, and misjudges what is required for mathematical realism. Like the platonist, the aristotelian wants to see numbers as real, objective, mind-independent features of the world. But that is compatible with the idea that numbers are not primary substances, because an aristotelian holds that there are things other than primary

substances which are real, objective, and mind-independent. Aristotelianism about numbers is the thesis that numbers are not primary substances, but properties of other things, and therefore *dependent* beings.

It is easy to see that this aristotelian position is compatible with thinking of numbers as objects in the investigatory sense. We have seen already that some things the aristotelian would count as properties can play the role of objects in the investigatory sense. Colors are one example: they are qualities of substances, not substances themselves; yet they are specifiable individuals, things which occur as one among many values in a range, things that can be sought, found, and given as the value of a variable. So if the argument I gave above is correct, something can be an object and yet still be a dependent being in the aristotelian ontology. We have only to add that the same is also true for numbers. An aristotelian may take numbers, like colors, to be properties of other things without denying that they are objects in the investigatory sense.²³ Any resistance to this idea, I think, would ultimately derive from the Ontic Picture, and a tendency to align the epistemological distinction between objects and concepts with an ontological distinction between primary substances and properties. By rejecting this picture above, we have already made room for the idea that properties may play the role of objects, and so for aristotelianism about numbers.

Of course, if the platonist is making the claim that numbers are primary substances, he is asserting a particularly strong ontological thesis. Thus, the aristotelian who denies it is not making a particularly strong claim. But there is also a weaker and more plausible way of understanding the platonist's thesis: numbers are *mind*-independent. Denying this weaker version of platonism requires making a stronger claim, one that is more difficult to reconcile with realism. Can a moderate realist hold that mathematical truths are objective, and that numbers are objects, yet

²³In spelling out her position, though, a modern-day aristotelian about numbers would have to take up Frege's arguments that numbers are not properties of 'external things', at least not in the same sense that colors are (Frege, 1884/1980, Cf. §§21–25). Unlike colors, numbers only belong to things with respect to the concept we regard those things under: the same external thing may be regarded as 1000 leaves or one pile, one book or 24 chapters. An aristotelian can answer these objections, I think, by holding that numbers are properties which belong to primary substances only in virtue of some *part* or *aspect* of their form. There is even an aristotelian idiom to express this idea: 1000 is a property of this thing *qua* leaves, but not *qua* pile; 24 is a property of this thing *qua* chapters, but not *qua* book.

hold that numbers are mind-dependent? I believe she can. The resulting position is a kind of idealism.

Against the idea that numbers are mind-dependent, the platonist emphasizes that mathematical truths are *discovered* rather than *created*. This is just part of the phenomenology of mathematics: in mathematics, we find truths which are new *to us*, but we feel they were 'already' true somehow. They are not like the claims in a work of fiction, which are products of human imagination, and were not in any sense true before they were authored. Instead, mathematical truths seem more like the truths of empirical science: we accumulate evidence for them, often hypothesizing that they are true before anyone actually proves them; and their truth can sometimes surprise us. The platonist takes this as evidence that mathematical claims are made true by independently existing things, things which are akin to the objects of the empirical sciences in that they are external to us and exist 'already', things which are not products of our imagination or will.

Frege himself often voices this sort of intuition. Thus we find him saying in the *Grundlagen*, for example, that

the mathematician cannot create things at will, any more than the geographer can; he too can only discover what is there and give it a name. (Frege, 1884/1980, §96)

Such remarks seem like grounds for putting Frege in the platonist camp.

But there is another line of thought in Frege that points to a subtler position. A few pages later, in §105, Frege writes:

In arithmetic we are *not concerned with objects which we come to know as something alien from without* through the medium of the senses, but with objects given directly to our reason and, as its nearest kin, utterly transparent to it. And yet, or rather *for that very reason*, these objects are not subjective fantasies. There is nothing more objective than the laws of arithmetic. (Frege, 1884/1980, §105, my emphasis)

This passage tells against attributing platonism to Frege. Frege is saying in this passage that the objectivity of numbers is in fact a *consequence* of their 'being given directly to our reason': it is *because* of their intimate

connection with reason that they are objective. In other words, it is precisely *not* their existence *independent* from minds which gives them objectivity (unlike things that we “come to know as something alien from without”), but something like their *dependence* on our cognitive capacities which affords them their status. Frege’s thought here is in keeping with the Kantian characterization of objectivity he gives in §26 and §27, which I discussed above. To be sure, Frege is not going so far here as to say that numbers are ontologically dependent on us. But the point is that for Frege, the kind of ontological independence that external things have is not necessary for something to be an object, or to be objective. Frege is not, therefore, committed to platonism, at least in the sense in which I am understanding it here.

Frege says that numbers are objective because they are “given directly” to reason, and therefore “transparent” to it. What does this mean? One sense in which numbers might be transparent to reason is that they are created by reason. They are known so intimately because they are, in Dedekind’s phrase, “free creations of the human mind” (Dedekind, 1963, p. 31): we know them not as a traveler knows a foreign landscape, but as a carpenter knows the house he builds. That is clearly not Frege’s position; as we have seen, he explicitly denies that we create the numbers. But perhaps there is room here for a distinction. When Frege rejects the idea that the numbers are created by us, he is rejecting the idea that they are fictional, or imagined—that they are products of human creative acts. But there is a looser sense in which they might be ‘created’ by reason, the same sense in which any law-governed phenomenon creates things. The movement of tectonic plates creates mountains; the orbits of the Earth and Moon create eclipses. These things are created in the sense that they are causal consequences of other things. Their existence and nature is *explained by* those things, and they couldn’t exist without those things. It seems to me that for Frege, the numbers could be like that: they are created by reason in the sense that they are products of the laws governing reason; their existence and nature are consequences of those laws. In that sense, they are created by reason, though they are not *authored* by it.

Whether or not this is in fact Frege’s position, it looks like the start of a moderate realism on which numbers are mind-dependent. According to this idealist position, numbers are ontologically dependent on our capacity for reason, and therefore on our minds, because their existence is a consequence of the laws governing reason. We have to be careful

here: to understand this claim as Frege would, we should not conceive of the laws that give rise to numbers as the descriptive laws of human psychology, but as normative, logical laws. The way that numbers are products of the laws governing reason is less like the way that mountains are products of the laws of plate tectonics, and more like the way that a duty to serve on a jury is a product of the laws of the state. This does not mean that the numbers are products of human convention or authorship. It means that their existence depends on a certain normative status: just as the existence of a duty to serve on a jury depends on the existence of citizens, people with the normative status of being subject to the laws of the state, the existence of numbers depends on the existence of reasoners, minds with the status of being subject to logical laws. So numbers do not exist independently; but given that we are reasoners subject to certain laws, they do nevertheless exist.

What is interesting about this position is that it can accommodate our intuitions about the phenomenology of mathematics. While the idealist holds that numbers are created by the laws governing reason, this is compatible with saying that numbers and the truths about them are *discovered* by reason in another sense. For one of the special features of reason is its ability to reflect on itself, to come to know its own laws as well as their consequences. Mathematical truths are perhaps discovered in something like the way that I discover that I have jury duty: even though this duty is the consequence of a law to which I am already subject, I may not know the details of the law, or its particular consequences for me, until an envelope arrives in the mail. In the same way, I may be ignorant of the laws governing reason, even though I am subject to them; and I can discover their particular consequences, even if an omniscient reasoner could see that those consequences held all along.

From this perspective, numbers and the truths about them are 'transparent' to reason in the sense that they are accessible to reason via reflection. But that does not mean that such truths are immediate, or that reflection requires no intellectual effort. Mathematical discovery is a process of using reason to understand itself, in all its fine detail, by drawing out and becoming aware of the consequences of its governing laws. Some of those consequences concern particular mathematical objects; the intellectual effort to discover which objects those are is ultimately to be understood as an investigation by reason into the products of its own laws. For example, a beginning number theorist may not know what the

greatest prime number less than 100 is, and she can discover through a process of investigation that it is 97. She has here come to know an objective truth about a mathematical object, and will experience this as a discovery, not an act of creation. But for all that, the number 97 and its unique properties ontologically depend on reasoners and minds like her own.

I have, of course, only sketched the idealist's ontological position. But if it is coherent, then it seems clear that it is compatible with the claim that numbers are objects in the investigatory sense. For articulating the position even as roughly as I have done makes it clear why the ontological dependence of numbers on minds does not undermine their objectivity, or the phenomena of mathematical discovery. Reasoners have only a limited capacity to survey the laws governing their own minds and the consequences of those laws. It is this limitation that makes it possible to conduct mathematical investigations and make mathematical discoveries; and it is the possibility of such mathematical investigations that implies that numbers are objects in the investigatory sense. I think we may take it, too, that the position *is* coherent. Although I can hardly hope to prove this, I take comfort in the fact that much greater minds have thought so. I turn again to Frege's own words:

On this view of numbers the charm of work on arithmetic and analysis is, it seems to me, easily accounted for. We might say...: *the reason's proper study is itself*. (Frege, 1884/1980, §105, my emphasis)

So the conclusion of this epilogue is that the investigatory conception of objects implies very little about the ontology of numbers. The old debate between the platonist, the aristotelian, and the idealist can carry on. Each has a position that is compatible with the claim that numbers are objects. I see this conclusion as a positive result, for it means that the status of mathematical objects *as* objects and as what makes mathematical claims objectively true is under no threat from metaphysics. Mathematicians, I am told, never really had any doubt about this. But I hope there is nevertheless philosophical value in having made a route to that conclusion explicit.

Bibliography

- Akmajian, A. (1970). "Aspects of the Grammar of Focus in English". Cambridge: Massachusetts Institute of Technology.
- AnderBois, Scott (2010). "Sluicing as Anaphora to Issues". In: *Proceedings of SALT 20*, pp. 451–470.
- Aristotle (1963). *Aristotle's Categories and De Interpretatione*. Trans. by J.L. Ackrill. Clarendon Aristotle Series. Oxford University Press.
- Barker, Chris (2016). "Why Relational Nominals Make Good Concealed Questions". In: *Lingua*. SI: Understanding Possession 182, pp. 12–29.
- Belnap Jr., Nuel D. and Thomas B. Steel Jr. (1976). *The Logic of Questions and Answers*. Yale University Press.
- Brandom, Robert (1994). *Making It Explicit*. Harvard University Press.
- Brogaard, Berit (2007). "Number Words and Ontological Commitment". In: *The Philosophical Quarterly* 57.226, pp. 1–20.
- Burge, Tyler (1992). "Frege on Knowing the Third Realm". In: *Mind* 101.404, pp. 633–650.
- Caponigro, Ivano (2003). "Free Not to Ask: On the Semantics of Free Relatives and Wh-Words Cross-Linguistically". University of California, Los Angeles.
- Caponigro, Ivano and Daphna Heller (2007). "The Nonconcealed Nature of Free Relatives: Implications for Connectivity Crosslinguistically". In: *Direct Compositionality*. Ed. by Chris Barker and Pauline Jacobson. Oxford University Press, pp. 237–263.
- Carnap, Rudolf (1956). "Empiricism, Semantics, and Ontology". In: *Second*. University of Chicago Press, pp. 205–221.
- Dedekind, Richard (1963). *Essays on the Theory of Numbers*. Trans. by Wooster Woodruff Beman. Dover Publications.
- Den Dikken, Marcel, André Meinunger, and Chris Wilder (2000). "Pseudoclefts and Ellipsis". In: *Studia Linguistica* 54, pp. 41–89.

- Dummett, Michael (1981). *Frege: Philosophy of Language*. Second. Harvard University Press.
- (1991). *Frege: Philosophy of Mathematics*. Harvard University Press.
- Felka, Katharina (2014). “Number Words and Reference to Numbers”. In: *Philosophical Studies* 168.1, pp. 261–282.
- (2016). *Talking about Numbers: Easy Arguments for Mathematical Realism*. Vol. 3. Studies in Theoretical Philosophy. Frankfurt: Vittorio Klostermann.
- Frege, Gottlob (1967). “Begriffsschrift”. In: *From Frege to Gödel*. Ed. by J. Van Heijenoort. Harvard University Press, pp. 1–82.
- (1980). *The Foundations of Arithmetic*. Trans. by J.L. Austin. Second. Evanston, Illinois: Northwestern University Press.
- (1997a). “Comments on *Sinn* and *Bedeutung*”. In: *The Frege Reader*. Ed. by Michael Beaney. Blackwell Publishing, pp. 172–180.
- (1997b). “Function and Concept”. In: *The Frege Reader*. Ed. by Michael Beaney. Blackwell Publishing, pp. 130–148.
- (1997c). “Introduction to Logic (Extract)”. In: *The Frege Reader*. Ed. by Michael Beaney. Blackwell Publishing, pp. 293–298.
- (1997d). “Letter to Marty”. In: *The Frege Reader*. Ed. by Michael Beaney. Blackwell Publishing, pp. 79–83.
- (1997e). “On Concept and Object”. In: *The Frege Reader*. Ed. by Michael Beaney. Blackwell Publishing, pp. 181–193.
- Groenendijk, Jeroen and Martin Stokhof (1997). “Questions”. In: *Handbook of Logic and Language*. Ed. by J. van Benthem and A. ter Meulen. MIT Press, pp. 1055–1124.
- Hale, Bob (2001a). “Singular Terms (1)”. In: *The Reason’s Proper Study*. Oxford University Press, pp. 33–47.
- (2001b). “Singular Terms (2)”. In: *The Reason’s Proper Study*. biblatex: Hale2001-ST2. Oxford University Press, pp. 48–71.
- Hale, Bob and Crispin Wright (2001). *The Reason’s Proper Study*. Oxford University Press.
- Heim, Irene (1979). “Concealed Questions”. In: *Semantics from Different Points of View*. Ed. by Rainer Bäuerle, Urs Egli, and Arnim von Stechow. Springer Berlin Heidelberg, pp. 51–60.
- (2002). “File Change Semantics and the Familiarity Theory of Definiteness”. In: *Formal Semantics: The Essential Readings*. Ed. by P. Portner and B. Partee. Blackwell, pp. 223–248.

- Heller, Daphna (2005). "Identity and Information: Semantic and Pragmatic Aspects of Specificational Sentences". New Brunswick, New Jersey: Rutgers, The State University of New Jersey.
- Heycock, Caroline and Anthony Kroch (1999). "Pseudocleft Connectedness: Implications for the LF Interface Level". In: *Linguistic Inquiry* 30.3, pp. 365–397.
- Higgins, Francis Roger (1979). *The Pseudo-Cleft Construction in English*. Outstanding dissertations in linguistics. New York: Garland Publishing.
- Hintikka, Jaakko (1973). *Logic, Language-Games, and Information*. Clarendon Press, Oxford.
- (1974). "Quantifiers vs. Quantification Theory". In: *Linguistic Inquiry* 5.2, pp. 153–177.
- Hintikka, Jaakko and Jack Kulas (1985). *Anaphora and Definite Descriptions: Two Applications of Game-Theoretical Semantics*. Dordrecht, Holland: D. Reidel.
- Hintikka, Jaakko and Gabriel Sandu (2011). "Game-Theoretical Semantics". In: *Handbook of Logic and Language*. Ed. by Johan van Benthem and Alice ter Meulen. Second. London: Elsevier, pp. 415–465.
- Hodges, Wilfrid (2013). "Logic and Games". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2013. Metaphysics Research Lab, Stanford University.
- Hofweber, Thomas (2005). "Number Determiners, Numbers and Arithmetic". In: *Philosophical Review* 114.2, pp. 179–225.
- (2007). "Innocent Statements and Their Metaphysically Loaded Counterparts". In: *Philosophers' Imprint* 7.1.
- (2014). "Extraction, Displacement, and Focus". In: *Linguistics and Philosophy* 37.3, pp. 263–267.
- Jackson, Brendan Balcerak (2013). "Defusing Easy Arguments for Numbers". In: *Linguistics and Philosophy* 36.6, pp. 447–461.
- Knowles, Robert (2015). "What 'the Number of Planets Is Eight' Means". In: *Philosophical Studies* 172.10, pp. 2757–2775.
- Link, Godehard (1983). "The Logical Analysis of Plural and Mass Terms: A Lattice-Theoretical Approach". In: *Meaning, Use and the Interpretation of Language*. Ed. by Rainer Baurle, Christoph Schwarze, and Arnim von Stechow. Berlin: de Gruyter, pp. 302–323.
- Lotze, Hermann (1887). *Microcosmos: An Essay Concerning Man and His Relation to the World*. Trans. by Elizabeth Hamilton and Emily Eliza-

- beth Constance Jones. Second. Vol. 1. 2 vols. New York: Scribner and Welford.
- Mikkelsen, Line (2005). *Copular Clauses: Specification, Predication, and Equation*. Vol. 85. *Linguistik Aktuell/Linguistics Today*. Amsterdam, Philadelphia: John Benjamins Publishing Company.
- (2011). “Copular Clauses”. In: ed. by Klaus von Heusinger Claudia Maienborn and Paul Portner. Vol. 2. Berlin: Mouton de Gruyter.
- Moltmann, Friederike (2013). “Reference to Numbers in Natural Language”. In: *Philosophical Studies* 162, pp. 499–536.
- Moro, A. (1997). *The Raising of Predicates: Predicative Noun Phrases and the Theory of Clause Structure*. Cambridge University Press.
- Partee, Barbara (1987). “Noun Phrase Interpretation and Type-Shifting Principles”. In: *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*. Ed. by J. Groenendijk, D. de Jongh, and M. Stokhof. Dordrecht, Netherlands: Foris Publications Holland, pp. 115–143.
- Rieppel, Michael (2013a). “Being Something: Prospects for a Property-Based Approach to Predicative Quantification”. Berkeley, CA: University of California, Berkeley.
- (2013b). “The Double Life of ‘The Mayor of Oakland’”. In: *Linguistics and Philosophy* 36.5, pp. 417–446.
- (2016). “Being Something: Properties and Predicative Quantification”. In: *Mind* 125.499, pp. 643–689.
- Romero, Maribel (2005). “Concealed Questions and Specificational Subjects”. In: *Linguistics and Philosophy* 28, pp. 687–737.
- Ross, John Robert (1972). “Act”. In: *Semantics of Natural Language*. Ed. by Donald Davidson and Gilbert Harman. Dordrecht: Reidel, pp. 70–126.
- Schlenker, Philippe (2003). “Clausal Equations (A Note on the Connectivity Problem)”. In: *Natural Language and Linguistic Theory* 21, pp. 157–214.
- Sellars, Wilfrid (1985). “Towards a Theory of Predication”. In: *How Things Are: Studies in Predication and the History of Philosophy of Science*. Ed. by J. Bogen and J.E. McGuire. Kluwer Academic Publishers.
- Sluga, Hans D. (1977). “Frege’s Alleged Realism”. In: *Inquiry* 20 (1-4), pp. 227–242.
- (1980). *Gottlob Frege*. Routledge.
- Strawson, P. F. (1950). “On Referring”. In: *Mind* 59.235, pp. 320–344.

BIBLIOGRAPHY

- Tarski, Alfred (1967). *The Completeness of Elementary Algebra and Geometry*. Reprinted monograph. Paris: Centre National de la Recherche Scientifique, Institut Blaise Pascal.
- Tugendhat, Ernst (1970). "The Meaning of 'Bedeutung' in Frege". In: *Analysis* 30.6, pp. 177–189.
- Williams, Edwin (1983). "Semantic vs. Syntactic Categories". In: *Linguistics and Philosophy* 6.3, pp. 423–446.
- Wright, Crispin (1983). *Frege's Conception of Numbers as Objects*. Scots Philosophical Monographs Two. Aberdeen University Press.