ONTOLOGY AND OBJECTIVITY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHILOSOPHY
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Thomas Hofweber
June 1999

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

John Perry
(Principal Co-Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Solomon Feferman
(Principal Co-Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Johan van Benthem

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

John Etchemendy

Approved for the University Committee on Graduate Studies:

Abstract

Ontology is the study of what there is, what kinds of things make up reality. Ontology seems to be a very difficult, rather speculative discipline. However, it is trivial to conclude that there are properties, propositions and numbers, starting from only necessarily true or analytic premises. This gives rise to a puzzle about how hard ontological questions are, and relates to a puzzle about how important they are. And it produces the ontology-objectivity dilemma: either (certain) ontological questions can be trivially answered using only uncontroversial premises, or the uncertainties of ontology are really a threat to the truth of basically everything we say or believe.

The main aim of this dissertation is to resolve these puzzles and to shed some light on the discipline of ontology. I defend a view inspired by Carnap's internal-external distinction about what there is, but one according to which both internal and external questions are fully factual and meaningful. In particular, I argue that the trivial arguments are valid, but they do not answer any ontological questions. Furthermore, I propose an account of the function of our talk about properties, propositions and natural numbers. According to this account our talk about them has no ontological presuppositions for its literal and objective truth. This avoids the ontology-objectivity dilemma, and solves the puzzles about ontology.

To do this I look at quantification and noun phrases in general, and at their relation to ontology. I argue that quantifiers are semantically underspecified in a certain respect, and play two different roles in communication. I discuss the relation between syntactic form and information structure, the function of certain non-referential, non-quantificational noun phrases, the uses of bare number determiners, and how arithmetic truths are learned and taught.

The more metaphysical issues discussed include: inexpressible properties, logicism about arithmetic, nominalism, Carnap's view about ontology, the problem of universals, the relationship between ontology and objectivity, different projects within ontology, non-existent

objects, and others.

A technical appendix deals with the relation that certain uses of quantifiers have to small fragments of infinitary logic.

Acknowledgments

I was very fortunate to be a graduate student at Stanford for the last 5 years, and to have had such dedicated, insightful and entertaining teachers. I have learned particularly much from Sol Feferman and John Perry during the many seminars I took with them, and from the numerous conversations we had. I'd like to thank them for their support when I considered working on this project for my dissertation, and for their supervising it together. During the time when I was working on the dissertation I profited tremendously from their help. They have given me constant feedback and have helped me to avoid many pitfalls. I couldn't be more grateful for their efforts.

Johan van Benthem has also helped me greatly in writing this dissertation. Through him I learned a number of things that play a central role in it, and attempts to copy his unparalleled energy kept me going at those late hours. My thanks for his efforts, as to John Etchemendy, for his support and helpful criticism.

I first learned about the problems this dissertation deals with when I was a visiting student at the City University of New York, Graduate Center, where I had the good fortune that Hartry Field and Stephen Schiffer were teaching there at the time. The conversations we had back then have always been of great value to me. Special thanks to Stephen for many e-mail exchanges and discussions throughout the years. That he had a great influence on me should be obvious from the view that I defend here.

It first occurred to me that the present topic might be my dissertation topic when I wrote a paper on (something like) it in a stimulating seminar taught by John Dupré and Peter Godfrey-Smith. I'd like to thank Peter for encouraging me to work this out and for his interest in the project.

Shortly after that a seminar by Robert Kraut on objectivity and many discussions with him and Philip Kremer helped me very much to get a lot clearer about what my view was, and what it shouldn't be. Later a seminar by Ed Zalta and lots of discussions with him at CSLI were most helpful in getting the project in the right direction.

I got many good comments on earlier drafts of this dissertation, had very helpful discussions after talks I gave on it, and had lots of very interesting conversations with many people about this set of problems. I'd especially like to thank the following, who in some way or other quite directly affected some part of this dissertation: John Bacon, Mark Balaguer, David Beaver, Mark Crimmins, Avrom Faderman, Kit Fine, Jeff King, Grisha Mints, Stanley Peters, Mark Richard, Stewart Shapiro, Peter Simons, Ken Taylor, Steve Yablo, and everybody mentioned above.

I'd like to gratefully acknowledge financial support from Stanford University, for a dissertation fellowship, from the Mellon Foundation, for a predoctoral fellowship, and from Johan van Benthem's Spinoza Grant, for funding a very helpful trip to ESSLLI 97. Thanks to CSLI for providing me with office space where this dissertation was written, and to the people who came up with LaTeX, which was used writing it.

Finally, special thanks to Rebecca Walker for lots and lots of all kinds of things.

Contents

Chapter 1

Ontology, Objectivity and Carnap

1.1 Ontology

Some philosophical disciplines attempt to answer fairly clear questions, others are driven by a general puzzlement that surrounds certain issues. Usually, disciplines of the first kind are easier to deal with. After all, you already know the question which you would like to answer. The philosophy of mind might be of this kind. It could be seen as trying to answer the question how the mind relates to the body. And even though that isn't easy to answer, it is at least a goal to work towards. Ontology seems to be of the second kind. That might seem surprising, since after all it is the philosophical discipline that tries to answer the question of what there is, what kinds of things make up reality. And that seems to be a fairly straightforward question, too. But the more one thinks about it, the less clear it is what is asked here. When we make a list of things that we think there are we soon get into an area where we feel very uncomfortable. For example, here are some things I think there are. I believe that there are:

- people
- cars
- something that we have in common
- infinitely many primes
- something that we both believe

- the common illusion that one is smarter than one's average colleague
- a way you smile
- a lack of compassion in the world
- the way the world is
- several ways the world might have been
- a faster way to get to Berkeley from Stanford than going through San Jose
- the hope that this dissertation will shed some light on ontology
- the chance that it might not
- a reason why it might not

and so on and so forth.

What is puzzling about this list is this: if you go down far enough on the list it just becomes quite hard to make sense of the claim that there really are such things as the ones listed. It seems that there isn't really such a thing the way way you smile, or the chance that Clinton might push the button. If we list the inventory of the world we will list cars and people, but chances, ways you smile and reasons won't be on the list. And all that even though there is a way you smile, and a chance that Clinton will push the button. And on top of this, it seems that just adding that there isn't *really* such a thing as the chance that Clinton pushes the button, even though there is a chance that he will push it, doesn't make any philosophically relevant difference.

To look at it from another point of view, just imagine someone writing a grant proposal asking to fund research into trying to find out more about what these things that I mentioned above really are. When you do this for people or cars, it seems all right. But when you do this for the way you smile, or the chance that you won't smile, it just seems really wrong headed. Trying to find out more about what these things are seems to be somehow confused. I find it quite legitimate that the NSF does not fund research that tries to find out what these things are, even though I think it should fund research into what kinds of things the world is made off. And all that even though I think there is a way you smile and a chance that you might not agree with me. Something has gone wrong somewhere, and it is not clear what or where. This dissertation tries to shed some light on what has gone wrong, and where.

To do this I think it is best to try to isolate what is puzzling about ontology, and how these puzzles arises. This, I think, can be done. In section 1.2 I will formulate two puzzles that arise about ontology. Understanding these puzzles seems to me to be the key to having a better understanding of ontology. After presenting the puzzles, and some sections trying to clarify them, I will outline how I think these puzzles should be solved. Spelling out this solution, and how it affects how we should understand ontology, will be the main task of this dissertation.

Before we get into this, let me make some brief remarks on what is commonly meant by "ontology". We will see much more about this in section 1.5.2, and in chapter 6. The word "ontology" is used in several different ways. First, ontology is the philosophical discipline that tries to find out what kinds of things the world is made of, or what there is, for short. This discipline is considered to be part of metaphysics. Secondly, ontology is used in the sense in which we talk of a person's, or a group of people's, ontology: what there is according to the beliefs of the person or persons. So, a priest's ontology contains God, and the 18th century chemist's ontology included phlogiston. Thirdly, we talk of a theory's ontology. This can mean either what there is according to the theory, or it can mean the entities that are somehow taken recourse to in the formulation of the theory. How these two senses of a theory's ontology relate to each other will be the subject of further investigation below. In the former sense we can say that electrons are in the ontology of certain physical theories, and in the latter sense that a certain type structure is in the ontology of Montague semantics. Finally, ontology is sometimes simply used for the collection or totality of things there are.

Today almost all of ontology, the discipline, is within the paradigm that started with Quine's classic essay "On what there is", (Quine 1980). According to Quine, there really isn't much of a special philosophical discipline called ontology. If we want to know what there is we should engage in science and we should ask what there has to be in order for our best scientific theories to be true. And this, Quine thought, would come down to what the variables of our best scientific theories have to range over for the theory to be true. Ontology isn't a special philosophical discipline, but a fallout of science. Still, Quine's view of ontology leaves room for philosophical discussions about ontology. And almost all of contemporary discussion in ontology falls within these options:

- Philosophers can disagree about whether or not our best overall theory requires certain entities in the domain over which the variables range. This is very explicit in the so-called indispensability arguments in the philosophy of mathematics. And similarly, they can disagree about whether or not our best theory will ultimately require its variables to range over properties or propositions, or whether speaking of them is convenient, but will ultimately have to be abandoned.
- Philosophers can construct formal theories of the entities that they think the best overall theory will require to exist. They can make mathematical theories about what events, facts and the like are, in particular, give precise conditions for their individuation (i.e. say which ones are the same and which ones are different).
- Philosophers can develop theories about which entities are special cases of which other entities. Certain philosophers might agree that there are events, states of affairs, properties, sets and possible worlds, but disagree about which ones are the most basic kinds of entities, and which ones are really just special kinds of the former. So, some philosophers believe that ultimately the world is made up of states of affairs, and properties and the rest are just special constructions of these, as in (Armstrong 1997). Or they might think that really the world is made up of possible worlds and sets of them and their inhabitants. Properties and the like are just special cases of these, (Lewis 1986).

Quine's view about ontology has a few problems of its own. One is to spell out the "must" in "must exist for a theory to be true". Another is the threat of circularity. In theory construction, don't we prefer theories that are ontologically plausible, just like we prefer theories that are simple? If so, can we really say that the resulting theory is our only guide to insights into ontology?[1]

Ontology, as it is done today, is puzzling to a degree that it seems dubious that we really know what we are doing when we are doing ontology. To make this point I would like to point to two puzzling features of ontology. If we can understand these puzzles better we

[1] See (Yablo 1998) and the references therein for more on this.

will understand better what is going on in ontology. And if we can resolve the puzzles we might put ontology on the right track.

A number of things that will be addressed below will have relevance for ontology quite generally. But the main focus of this dissertation will be somewhat more specific. We will mainly deal with ontological questions related to three kinds of things that have traditionally been of great interest to philosophers: numbers, properties, and propositions. Much of what will be said about them will be relevant to other kinds of things, too, like facts, events, chances, ways the world might have been, and even ways you smile. But I won't elaborate on these here. More work will has to be done to see how much of what is said below carries over to them, and as we will see, it all depends on the details.

1.2 Some puzzles about ontology

Here are two puzzles about ontology. One is a kind of a theoretical puzzle, or even an antinomy. The other is more of a puzzle about the attitude that philosophers have towards ontology. Both of these puzzles arise from the fact that there are apparently two equally correct, but contrary, answers to a simple question. The first puzzle comes from a question about the difficulty of achieving ontological results, the second comes from a question about the importance of such results.

1.2.1 Puzzle 1: How hard are ontological questions?

Answer 1: Very hard

Ontological questions are questions about what kinds of things make up reality. Some such questions might not be very hard to answer. For example, it's easy to answer the question whether or not there are thinking beings (meaning: at least one of them). Don't ask me how we manage to answer that so easily, but it seems clear that we can answer this one right away. It will be harder to answer the question whether or not there are material objects, but let's assume that this one we also can answer positively and rather easily. However, it will certainly be quite hard to answer the question whether or not there is a God. How can we decide this question? One might think that even if we have no religious experience that could count as direct evidence for the existence of God, we do have evidence that there are marks that God has left on the world. Or that we do have evidence for the absence of marks that God would have left, were there a God. Thus one might argue that we can answer the question whether or not there is a God by taking recourse to the fact that there is a world, or a world with order, or a world in which there is evil, or the like.

But how about numbers? They don't seem to leave any marks in the world as God might have, and it seems that the absence of such marks also doesn't speak against them. So, it seems that it's easier for us to figure out whether or not there is a God than whether or not there are numbers. The question whether or not there is a God is not an easy question to answer. So the question whether or not there are numbers isn't easy either.

Ontological questions are questions about the ultimate building blocks of reality. As such it seems that it might be more a subject of speculation than of serious results. After all, after centuries and centuries in which humanity has achieved a variety of results to numerous questions, it seems fair to say that neither the problem of universals nor basically any other ontological issue has received a definite answer. And the reason for that might just be that these are very hard questions, perhaps questions beyond our epistemic capabilities. At best, it seems, we can side with Quine. If our best overall theories of the world require the existence of certain entities for their truth then we will have to take that as our answer to the question whether or not their are such things. But then an answer to ontological questions will be as hard as coming up with our best overall theory. And that certainly isn't easy. Thousands of people are working on that theory all over the world, and hardly anyone knows more than a fraction of it.

Answer 2: Very easy

The question whether or not there are numbers, properties or propositions can be answered by simple arguments using only premises that everyone accepts and no more than three obvious steps of reasoning. Here is how this goes, for each case:

- **Numbers**

  1. Three men entered a bar.
  2. Thus: the number of men who entered a bar is three.
  3. Thus: there is a number which is the number of men who entered the bar, namely three.
  4. Thus: There are numbers, among them the number three.

- **Properties**

  1. Fido is a dog.
  2. Thus: Fido has the property of being a dog.
  3. Thus: there is a property that Fido has, namely being a dog.
  4. Thus: there are properties, among them the property of being a dog.

- **Propositions**

  1. Joe believes that Fido is a dog.
  2. Thus: there is something Joe believes, namely that Fido is a dog.
  3. Thus: there are propositions, among them the proposition that Fido is a dog.

  or alternatively:

  1. Fido is a dog.
  2. Thus: it's true that Fido is a dog.
  3. Thus: there is something which is true, namely that Fido is a dog.
  4. Thus: there are propositions, among them that Fido is a dog.

Some steps in these arguments might seem objectionable. It seems that the argument for properties and propositions use inferences where we suddenly start talking about properties and propositions, and it might be said that these are artificial philosopher's inventions, and therefore objectionable. That is a red herring, though. Sure enough, in ordinary English the word "property" is only rarely used to mean what philosophers mean by it. But there are plenty of other world in ordinary English that actually do mean just that, like "feature", "attribute" or "characteristic". We could just use those, if you prefer. The word "proposition" in philosophy stands for the kinds of things that that-clauses stand for, whatever they may be. Since the last arguments show that there is such a thing as that Fido is a dog, it shows that there are propositions.

Now, let's be honest. Are there any other philosophical problems that have a solution that is as simple and uses only as uncontroversial premises as the problem whether or not there are numbers, properties and propositions? I can't think of one example. And this is reflected in one attitude that philosophers and non-philosophers alike have towards the question whether or not there are numbers or properties. The answer is: of course there are, for example the number seven and the property of being red. It seems that the issues in ontology that interest philosophers the most, those dealing with properties, propositions and numbers, have completely trivial answers.

1.2.2 Puzzle 2: How important is ontology?

Answer 1: Not very important

It has to be admitted that not too many people lose any sleep over the question whether or not there really are numbers or properties. But lots of people do lose sleep over lots of other philosophical questions, philosophers and non-philosophers alike. Consider: Does the state have any authority? Is abortion legitimate in circumstances X? Is a life after the death of my body possible? Is there a God? And so on and so forth. The reason they don't lose any sleep over the former question seems to be that it isn't such an important question, perhaps even that it is a bit of a silly question.

Similarly, to raise ontological issues in, say, a seminar about the philosophy of mind is considered annoying and inappropriate. Even if the question under discussion in the seminar is whether or not every mental property supervenes on a physical property, not too many people will have much sympathy with you if you insist that before we go any further we have to solve the problem about whether or not there really are properties. And it seems that these annoyed people would have a point. The same holds for a math class. If you want it settled first whether or not there are any numbers, before you are willing to listen to any proofs in number theory, you will not meet much sympathy. These ontological worries are in practice treated like worries about the semantic paradoxes, or vagueness. Everybody knows that they are there, and everybody thinks that someone should work on them, but that whatever they come up with, we don't have to wait for that first in order to go on doing philosophy or mathematics as usual.

It seems you don't have to solve the semantic paradoxes first before you can do truth conditional semantics. At least Davidson didn't think so.[2] And that's just what we, in practice, think about ontology.

[2]Tarski thought that no Tarski-style truth theory could be given for natural languages because of the paradoxes. Davidson's answer to this: "...[this] point deserves a serious answer, and I wish I had one." He then goes on to develop his influential program in the philosophy of language using a Tarski-style truth theory for natural languages. See (Davidson 1984b, 28).

I'd like to point out that in practice the ontological worries are not considered annoying because the answers are so obvious. If you ask your local philosopher of mind what their view on properties is they will often not have one, or at least not have a firm one. Sure enough, many will readily admit to give properties whatever ontological status they think is required for them to have so that they, as philosophers of mind, won't have to deal with them any more. In general, it is a common opinion that if you are a platonist you don't have to do much theoretical work, but if you are a nominalist, you have to explain or translate away our property talk. Platonists seem to get a free ride, whereas nominalists have to do extra work.

Overall, it seems fair to say that overall the common opinion is that there are real issues about ontology, but that they are better left for the specialists.

Answer 2: Very important

In the introduction of *Psychosemantics* Fodor writes the following:

> "The main moral is supposed to be that we have, as things now stand, no decisive reason to doubt that very many common sense belief/desire explanations are - literally - true. Which is just as well, because if common sense intentional psychology were to collapse that would be, beyond comparison, the greatest intellectual catastrophe in the history of our species; if we're that wrong about the mind, then that's the wrongest that we've ever been about anything. The collapse of the supernatural, for example, didn't compare; theism never came close to being as intimately involved in our thought and our practice - especially our practice - as belief/desire explanation is. [...] We'll be in deep, deep trouble if we have to give it up." (Fodor 1987, xii)

And that seems to have something right. But note that belief/desire psychology seems to have ontological presuppositions. The literal truth of belief ascriptions seems to presuppose the existence of propositions. After all, and as we saw above, from the truth of "Joe believes that Fido is a dog" we can infer quite easily that there is a proposition that Fido is a dog. But what that means is that if there are no propositions then belief ascriptions aren't literally true. And that does not only apply to belief ascriptions, it applies to almost everything. That there are infinitely primes implies that there are numbers. But then again, if there aren't any numbers then it can't be literally true that there are infinitely many primes. And similarly for basically every branch of mathematics. And thus similarly for every branch of science that presupposes mathematics. And since basically everything seems to imply the existence of properties (as the "Fido is a dog" example seems to show) it seems that if there really were no properties then we have been very wrong about everything. Thus the greatest intellectual threat to our species is not eliminativism, or the like, in the philosophy of mind, but a potential discovery that the world does not contain such things as numbers, properties and propositions. If this were true, not only would belief/desire psychology have to go, but with it almost all of math and science, and of ordinary discourse.

Not all of the above can be true. It can't, let's hope, be true that ontological questions are both hard and easy, and both important and unimportant. But intuitively it seems that each of the above answers is right, or at least sort of right, or at least on to something. Ontological questions should be hard, but it seems that they have trivial answers. And ontological questions should be important, but we push them to the background and postpone them, and it seems that this is in fact the right thing to do. We will have to see how we can resolve this tension. But first, we should have a closer look at what threat ontology really poses to our talk about beliefs and numbers. This issue will be central for the rest of this dissertation, and is in much of the contemporary debates about ontology.

1.3 Ontology and objectivity

1.3.1 The ontology-objectivity dilemma

The threat of ontology

Let's assume that one day it will turn out, through the cooperation of philosophy, psychology and cognitive science, that people really do not have beliefs and desires, and that common sense psychological explanations aren't strictly speaking true. The true explanation of behavior has nothing to do with beliefs, but only with neural stuff. What would happen? In a sense, it would certainly be an intellectual catastrophe if this would turn out to be so. But on the other hand, it wouldn't change to much of our everyday practice. We would continue to ascribe beliefs and desires to people, and we would continue to explain their behavior the way we do. Sure enough, in more theoretical moments we would realize that this isn't strictly speaking correct, but in more practical moments we would take recourse to it as our only option. That would be a bit like our present statements about simultaneity. We know that strictly speaking such judgments are only true relative to a frame of reference. But in regular life we simply ignore this. Whatever the outcome of cognitive science will be, we will continue our practice of ascribing beliefs to each other. It just is so useful for some reason or other. We will continue to be able to predict what we will say, buy or do in the near future by taking recurse to that practice. And all that even if all belief ascriptions are strictly speaking false.

Even if all belief ascriptions are strictly speaking false, still, some will be correct and others won't be correct. Sure enough, none of them will be strictly speaking true, but some will be correct, in some sense of the word, and others won't be. It will be correct to say

(1) I believe that I am alive.

and (barring great surprises) incorrect to say

(2) The pope believes that God was made up by the ruling class.

Any practice that is in the ball park of communicating information and that is worth having has to have some notion of correctness, or some standards of correctness, that governs it. So, it seems that cognitive science isn't a threat to our practice of ascribing beliefs, nor to there being some standards for correctness governing that practice. And it seems that at least this much applies to our talk about numbers, properties and propositions. Ontology isn't a threat to either the practice of talking about them, nor to there being standards of correctness governing that practice. What then, is the threat of ontology to that practice, if there is one at all?

Let's compare this situation with another one, say talk about God, especially among religious people. It seems that there is a threat from ontology to the religious person's utterances of

(3) God created the world.

(4) God will reward the just and punish the guilty.

The threat is simply the following: if it turns out, for whatever reason, that God is not among the things that make up reality, then it seems that (3) can't be true, but rather has to be false. If there really is no God then it has to be false that he created the world.

So it seems that what is threatened by ontology about the religious persons talk about God is simply the truth or falsity of what they say, at least strictly speaking. Why strictly speaking? Well, because things aren't always that simple. Consider the case when we knowing adults say to a little child

(5) If you have been a good boy Santa will bring you a big present this year.

(6) Santa has a number of reindeers who work for him.

It seems that whatever truth there is to what we say here, it does not depend on ontology in the same way in which (3) does. If it turns out that there really is no Santa, it will not affect how we evaluate what I said. After all, we both know that there is no Santa. When I said (6) it was meant quite differently than when the believer says (3). (3) was meant to be true, strictly speaking, whereas (6) wasn't meant to be. But that is not to say that (6) wasn't meant to be true in some sense of the word. It certainly was correct, at least given the standards of correctness that govern our talk about Santa. It was correct, whereas

(7) Santa has a number of camels who work for him.

is incorrect. But the standard of correctness here is different than the standard that governs talk of the kind of (3). (3) was meant to be a purely descriptive stating of facts. (6) wasn't meant to be that. And the standards of correctness that govern these different sorts of talk are different. Now, it is controversial where the notion of truth falls in on all this. Perhaps only purely fact stating talk can be said to be true or false, or perhaps other talk can be, too. After all, if we accept (6), don't we have to accept

(8) It's true that Santa has a number of reindeers who work for him.

since we do accept

(9) It's true that p iff p.

We don't have to decide this here, but can rather focus on what is more important for our main topic here.[3] Whatever the appropriate use of "true" is when it comes to non-factual discourse, we can see that ontology can be more of a threat to some kinds of discourse

[3]For more on the relation between truth and non-factual discourse, see (Field 1994b).

than to others, and that different kinds of discourses are governed by different standards of correctness. If what we aim for in a certain domain of discourse is a certain standard of correctness, call it literal and objective truth, then the non-existence of certain things that we talk about will doom us to failure in our attempt to live up to these standards. If our aims are different then the non-existence of what we talk about might not matter in our achieving the standards of correctness we want to achieve. We will have to look at this more closely.

Objective correctness

Our talk about Santa was never meant to be literally and objectively true, since it wasn't meant to be purely fact stating talk. Its function was rather something else, like engaging in a practice of pretending certain things, continuing the tradition of a nice fiction that gets kids excited, and the like. There are standards of correctness governing that practice, but they are different than objective truth. Whereas the religious believer's utterance of (3) was intended to be objectively true and was meant as purely descriptive and fact stating talk. For it to objectively true, God has to exist. For (6) to be true or correct, in whatever sense it can be said to be true or correct, Santa doesn't have to exist. Thus it seems that different kinds of talks are different with respect to how ontology can be a threat for them to live up to the standards of correctness they are supposed to live up to. If a certain utterance is meant to be descriptive and fact stating then the standard of correctness for it is objective truth. If an utterance is meant differently, then the standard of correctness for it might be different. Results about ontology seem to be more of a threat to utterances of the first kind to live up to their standards of correctness than to utterances of the second kind. Sure enough, the non-existence of Santa is a threat to the objective truth of (6), but it is an empty threat. (6) was never meant to be objectively true, so no one cares if it isn't.

The standard of correctness we intended to apply to our talk about Santa were not meant to be standards of objective correctness. When I say this I mean that such talk was not meant to be fact stating, purely literal, world-mapping talk. But that doesn't mean that there isn't, in a sense, objectivity in our talk about Santa. Objectivity can be taken to mean nothing other than that there are intersubjective standards of correctness. And to be sure, we have intersubjective standards of correctness when it comes to making claims about Santa. To have such standards is quite independent from what the general features of a domain of discourse are. We might have such standards even if our talk is not factual, or not literal, or what have you. In the case of talk about Santa, such talk is constrained by the Santa story, and a certain tradition of Santa related practices. This rich Santa heritage gives rise to intersubjective standards of correctness of talk about Santa. Ones claims about Santa are constrained by more than just the individual speakers preferences or opinions. So, in a sense there is objectivity in ones talk about Santa, in the sense in which there isn't any objectivity in ones claims about how good pesto pizza tastes. It seems that when someone says that pesto pizza tastes good then there are no standards of correctness that govern this statement that could give rise to a procedure of resolving disagreement about this. There is no point in arguing about this. When someone things that pesto pizza tastes good then this is rock bottom. You can disagree with that, in the sense that you think that it doesn't taste good, but there is no point in arguing about it. This doesn't apply to the case of Santa. If someone thinks that Santa employs camels then there is a point in arguing about it, and then there are ways to resolve such disagreements. So, in a sense there is objectivity in our talk about Santa that isn't in our talk about how good pesto pizza tastes. But in another sense there isn't objectivity in our talk about Santa. Santa is just made up. It isn't like talk about who first made it to the top of Mount Everest. The former is dependent on a story, the later is dependent on real events in the real world. To be sure, it's correct to say that Santa employs reindeers, but is it objectively correct?

Yes and no. There is a confusing ambiguity hidden in the expression "objectively correct". "Objectively correct" could either be taken to mean that it is objective that it is correct, or it could mean that it is correct according to especially strict standards. Let me explain.

Suppose you believe, with Hartry Field, that for a mathematical statement to be true, or correct, is for it to follow from generally accepted mathematical beliefs.[4] Then you should believe that in a sense there is no objectivity in mathematics, and in a sense you should believe that there is. The sense in which there is no objectivity is simply that nothing in reality makes a mathematical statement true other then our shared, and perhaps arbitrarily chosen, or merely useful, basic mathematical beliefs. If Field is right then there will be an important difference between the correctness of, say,

(10) There are more than a million cars.

[4] See (Field 1989b).

and

(11) There are more than a million primes.

The second is only true because it follows from something we all accept, the first is true independently of what we accept. The standards for correctness between the two is different in important respects. How exactly is a bit tricky to spell out, but intuitively, in the first case the world independently of us makes it true, and in the second case we make it true by our picking certain basic mathematical beliefs and not others. There is a sense in which it is not objectively correct, because it is too closely related to our opinions and beliefs.

But in a different sense it is objectively correct that there are more than a million primes, even if Field is right. This sense is that it is objectively so that the standard of correctness, which we have for mathematical beliefs, applies in this case. If the standard of correctness for mathematical beliefs is merely to follow from what is commonly believed, then it applies to (11), and it is objectively so that it do. To deny this would be to deny that what follows from what is an objective fact. This could, of course, be done. One might claim that it isn't an objective fact what follows from what, but that rather our standards for logical consequence are related to our common opinions in a similar way. Field, however, doesn't do this, and rightly so, it seems. Thus if one believes that standards for correctness about what follows from what are strict in the sense spelled out above, and that the standards for correctness for mathematical beliefs are merely for them to follow from what we commonly believe, then it will be that it is objectively so that (11) is correct. But only in the sense that it is objectively so that it follows from what we believe. It is objectively correct that the standards of correctness for this domain of discourse apply to this utterance, or belief. This sense of objective correctness leaves it open what the standards for correctness for that domain of discourse are. They could merely be that everyone has to happen to agree about it, or that it says so in the Bible, or the like.

When I speak of objective correctness, or objective truth, I will always use it in the above strong sense (unless I explicitly says otherwise, of course). Objective correctness will be a strong form of correctness, a standard of correctness with a strong independence from our opinions. It will not be used in the sense in which something is objectively correct if it is objectively so that it is correct, however weak the standards for correctness may be. I will use this notion of correctness to exclude certain expressivist and other non-factualist or non-literal approaches to the dilemma we will encounter below.

Thus in the weak sense of objective correctness it is objectively correct that Santa employs reindeers. But it isn't objectively correct in the stronger sense. Objective correctness in this sense would require it to survive a reality check.

The dilemma

Let's get back to talk about numbers, properties and propositions. What threat is ontology to it? We can now say that it seems to depend on what the standards of correctness that apply to it are supposed to be. Is talk about numbers meant to be purely descriptive and fact stating? Or is it meant to be otherwise? If the first, then it seems that the standard of correctness governing such talk is objective truth. If the second it will be something different. And it if it is something different, if our talk about numbers, properties and propositions was never meant to be objectively true, then we can expect that ontology is not much of a threat to it living up to the standards it is meant to live up to. However, it seems that if it is meant to be objectively true then ontology is such a threat. After all, the trivial arguments seem to show that if our talk about numbers, say, is meant to be objectively true then it has to be objectively true that there are numbers for this to succeed. And it certainly seems to be meant to be purely fact stating and descriptive. In the case of talk about Santa I know that it isn't meant to be purely factual. In the case of belief ascriptions I have no such knowledge. So, what are the options? Roughly the following two:

1. Talk about numbers, properties and propositions might be quite analogous to talk about Santa. It might be talk relative to a certain background fiction, or within a game of pretense. Such a view about mathematics has been developed in (Field 1980) and (Field 1989b). It is played with more generally and more in the spirit of pretense theory in (Yablo 1999).

2. Talk about numbers, properties and propositions is not purely factual, but rather analogous to what expressivists claim about moral discourse, as (Stevenson 1963). According to them the function of such talk is not to communicate facts, but to express attitudes, emotions or the like. It seems less plausible to claim this about, say, math, but it might not be impossible, and certainly is a option to be listed.[5]

[5] Robert Kraut seems to have a view like this about ontology, which, I think, he will present in future publication.

In either of these cases, the standards of correctness would be different than objective truth. It would rather be some notion of correctness within the game of pretense in the first case, or would derive from our talk about numbers and the like to fulfill the expressivist's purposes it has. In either case, the non-existence of numbers, properties and propositions would (apparently) be no threat for such talk to live up to the standards of correctness it is supposed to live up to. Again, if there are no such things then our talk would not be objectively true. But if it wasn't ever meant to be then this is an empty threat.

So, it seems that what is threatened by ontological findings is the objective truth of our talk about numbers, properties and propositions. If it was never meant to be objectively true then this is no bad thing. The threat is empty. If, however, it is meant to be objectively true, if that is the notion of correctness that we mean to govern this kind of discourse, then the threat is real. Our talk about numbers, properties and propositions, it seems, can only be objectively true if certain things are part of reality, namely numbers, properties and propositions. If it is objectively true then its ontological presuppositions are satisfied. And if they aren't satisfied then it can't be objectively true.

This gives rise to the ontology-objectivity dilemma. This dilemma is, in fact, a dilemma that occurs in philosophy on several occasions. A classic example of it arises if we believe the following, which was a not unpopular belief not too long ago:

(12) The objectivity of morality depends on the existence of God.

If we believe (12) then we believe that morality is objective only if there is a God. This can be seen in two ways. One way of looking at it is to insist that the objectivity of morality is beyond reasonable doubt and that therefore (12) should be understood as being part of an argument for the existence of God. But another way of looking at it is to say that we have no more reason to believe in the objectivity of morality than we have to believe in the existence of God. For the latter, however, we have very little reason. Therefor, the objectivity of morality is rather uncertain. perhaps it has to be taken on the basis of faith without reason, just like the existence of God. Thus if we have a dependence of the objectivity of a certain domain of discourse on a certain ontological position then this can always be taken in two ways: either as an argument for the ontological position, or as casting doubt on the objectivity of the domain of discourse. In other words: one person's modus ponens is anther person's modus tollens.

And similarly for mathematics, or belief ascriptions. The above trivial arguments seem to show that they can only be objectively true if numbers and propositions exist. And this can, again, either be taken as an argument for the existence of such entities, or as an argument casting doubt on the objectivity of this domain of discourse. Since the existence of numbers is presupposed for the objectivity of number theory we really have no more reason to believe in the latter than in the former. But we have very little initial reason to believe in the existence of numbers, so we have very little reason to believe in the objectivity of number theory.

It is this dilemma which defines much of the debate in contemporary ontology, in particular in the philosophy of mathematics. The debate can be classified by looking at two conditionals and their constituents. Consider the following statements:

(i) Our ordinary mathematical, folk psychological, etc. talk is literally and objectively true.

(ii) Numbers, propositions, etc. exist.

(iii) Numbers, propositions, etc. are abstract objects.[6]

From these we can build the following conditions:

(A) If our ordinary mathematical, folk psychological, etc. talk is literally and objectively true then numbers, propositions, etc. exist.

(B) If numbers, propositions, etc. exist then they are abstract objects

Thus (A) is "if (i) then (ii)", and (B) is "if (ii) then (iii)". With these the different options that are usually taken can be classified as follows:

- **Platonists** Accept: (i), (A), (B). Thus: (iii)
- **Weak Nominalists** Accept: (i), (A). Reject: (B), and rather believe that numbers, propositions, etc. are concrete objects.
- **Strong Nominalists** Accept: (A), (B). Reject: (iii). Therefore: reject (i).

[6]This is read as presupposing that numbers, properties, etc. exist.

To put it simply: the standard platonist accepts our ordinary discourse as objectively and literally true, and believes that it implies that there are abstract objects, with numbers, propositions etc. among them. A weak nominalist accepts our ordinary discourse as objectively true, and accepts that it implies the existence of numbers, propositions, etc.. However, they deny that these objects are abstract. A strong nominalist accepts that if our discourse were objectively true then it would imply that there are abstract objects, with numbers, propositions, etc. among them. However, they deny the existence of such objects and thus deny the objective truth of the relevant discourse. Weak nominalists will spend most of their pages saying what things numbers etc. really are, strong nominalists will spend them on saying what function the relevant discourse has other than being an objective description of reality.

It is interesting to note that all three positions accept assumption (A). In fact, it is (A) that sets the issue for most of the debate in contemporary ontology, and a good part of contemporary philosophy of mathematics. It will be the main aim of this dissertation, and a central step in resolving the above puzzles, to show that (A) is in fact false. I will argue that our ordinary talk about numbers, properties, and propositions, is intended to be objectively true, but that there are no ontological presupposition that might stand in the way of achieving this aim. When we talk about numbers, properties and propositions we have *objectivity without objects*. Or so I will argue.

It might seem impossible to argue this without cheap tricks. How could it be, for example, that both the following are true:

1. Numbers don't exist.

2. It is objectively true that there are infinitely many primes.

It seems that if 2. is true then 1. has to be false. In a moment we will encounter a position that might make understandable how this could be so. But first, let's look at another relation that ontology is said to have to objectivity.

1.3.2 Ontological presuppositions for objectivity

So far we have seen one way in which ontology is a presuppositions for the objective truth of what one says. But in philosophy there is often also another way in which ontology is taken to be presupposed for objectivity. This other way is especially about certain philosophically distinguished and important entities: properties, facts, propositions, and the like. This line of thinking might not be independent of what we have seen so far, and goes as follows: For there to be objectivity about a certain domain of discourse the properties that are ascribed in that talk have to exist. Or there have to exist facts that this talk is about. For example, for there to be objectivity in morality the properties we ascribe in moral discourse, being good and being bad, have to exist. Similarly, there has to exist a fact that what someone did was good, or bad.

Believers in this way of looking at the objectivity of a domain of discourse will think that it is one thing to say that

(13) John's donating the money was good.

and quite another to say that

(14) John's donating the money had the property of being good.

The former is taken to be neutral about the objectivity of morality. It can be so whether or not morality is objective. But the later says that morality is a real feature of the world, since it explicitly claims that there are moral properties. If, and only if, there are such properties can morality be said to be an objective feature of the world. And a similar line of reasoning will apply to talk about facts, and possibly also to talk about propositions.

According to this line of reasoning the objectivity of morality is tied to the existence of moral properties and facts. And similarly for other domains of discourse, like mathematics. If certain properties or facts exist then we have objectivity in this domain of discourse. Objectivity comes from, so to speak, what facts, properties, and the like, exist. This is made plausible by considering that the existence of properties and fact should be understood as being quite independent of us. Thus moral talk can have its objectivity from something that is independent of us, the existence of certain properties and facts.[7]

This general way of tying the objectivity of a domain of discourse to ontology can be seen as a special case of what has been said above if we make on additional qualification. In the above trivial arguments we always allowed the inference from

(15) a is F

[7]This way of tying ontology to the objectivity of a certain domain of discourse is quite common. See, for example, (Boghossian 1990).

to

(16) a has the property of being F.

But one might argue that such an inference is only valid if there is objectivity in ascribing F-ness. One might argue that talk about properties, facts and the like can be used to distinguish talk that is governed by low standards of correctness, from talk that is governed by high standards of correctness: objective truth. (Similarly, it has been claimed that talk about truth can be used to distinguish the mere correctness of a statement from its objective correctness.) Once that is in place properties and facts will be ontological presuppositions of objectively true statements, but not of merely correct statements. And thus the existence of certain properties distinguishes domains of discourse where we do have objectivity from those were we don't.

After we have seen more about the function of talk about properties in chapter 4 we will return to this way of looking at the role of talk about properties. I can only announce now that I will argue that it is quite misguided.

1.3.3 Preliminary conclusion

While looking at the puzzles we have seen that there seem to be good reason to assume that ontological questions are hard, and that they are easy, that they are important, and that they aren't important. Something must have one wrong somewhere in these considerations. We have seen that a central point for the importance of ontology was the apparent dependence of the objective truth of certain domains of discourse on certain positions in ontology, and that this dependence can be taken in two ways. Either one takes it as an argument for the existence of these entities, or as an argument against there being standards of objective correctness governing this domain of discourse. We have also seen that the response to this situation in philosophy can fall into one of two camps. Either it is an outright endorsement of the ontology that apparently is presupposed, or it is a chipping away on the apparent objectivity of the domain of discourse. Below we will see another way of dealing with this situation.

It might seem easiest to simply to to find out which one of the answers to the questions that gave rise to the puzzles is false. To do this one might either say that ontological questions are in fact easy and it is a certain illusion that they seem hard (Frege, Wright, Schiffer), or that ontological questions are in fact hard, and that there are certain subtle mistakes in the apparently trivial arguments (Field). I don't think any one of these routes can be right. Rather, I will endorse both sides, that ontological questions are hard, and that the trivial arguments are correct. However, I will deny that there is any contradiction between them. I will argue that

(I) Even though the trivial arguments are perfectly valid, this has no implications for ontology

This will be a rather direct consequence of my main claim, namely that

(II) The objective and literal truth of our talk about numbers, properties and propositions has no ontological presuppositions.

This, by the way, will have nothing to do with general views of what objectivity is. I will not argue that a special analysis of the notion of objectivity will allow us to get rid of the puzzles. The solution will be much more down to earth.

I think the solution to these puzzles is hidden in the original program in ontology to which Quine's program was supposed to be an alternative to. The creator of this program, Rudolf Carnap, avoided the word "ontology" and would not want to be associated with the part of philosophy it belongs to. Still, he had a very important insight about ontology.

1.4 Carnap's views on ontology

Carnap is not a very popular philosopher these days. There are two main reasons for this. First of all, Carnap was a leading figure in the Vienna Circle, and it is generally believed that the Vienna circle consisted of a bunch of positivists who had really naive views about science and really radical views about language. Secondly, it is commonly believed that Carnap's views essentially involved notions like analyticity and truth by convention, notions that Quine successfully criticized, which made Carnap's views to be of merely historical interest. This attitude toward Carnap is partly unfortunate, and partly justified. It is unfortunate because the criticism of Carnap's views by Quine isn't as devastating as it is commonly believed to be. On the one hand, Quine didn't really show what he claims to show about, say, analyticity, on the other hand Carnap's views don't really depend too much on taking recourse to these notions in any bad way. These are themselves big issues that we should not get into now.

It is also unfortunate because the criticism by Quine lead to an almost complete elimination of Carnap from the cannon of present day philosophy. In fact, it is quite hard to get one's hands onto texts by Carnap. Presently only two books by Carnap are in print, his *Meaning and Necessity* and his *Introduction to Symbolic Logic*. There is no collection of essays by Carnap in print. In addition, very few essays by Carnap made it into presently available anthologies. This is unfortunate, I think, because Carnap not only had very insightful ideas in general, he in particular had some that are most relevant to the present debate. I will elaborate on this momentarily.

On the other hand, ignoring Carnap seems to have some justification. The verifiability criterion of meaningfulness, that motivates his whole negative attitude towards metaphysics, seems to be a mistake, and is hardly believed to be correct by anyone today. And his views on science seem to be somewhat dated, too.

I think Carnap had an important insight that will help us understand ontology. However, I don't want to defend Carnap's views about ontology. They are based, to a good extent, on his verifiability criterion of meaningfulness, which I reject. But Carnap had an insight which seems to help us answer the above puzzles. So, let me say a bit more about what Carnap's views were about (what we would call) ontology, and how his insight will help us to solve our two puzzles.

Carnap started his philosophical career with a good deal of contempt for the early 20th century metaphysics that he was exposed to in Germany. A step against that tradition was his first book (Carnap 1928), which deals with how to explicitly define certain concepts in terms of simpler concepts. It was assumed that the simplest concepts were ones that were purely phenomenal, or about sense data. Even though he tried to remain neutral about the status of metaphysics in general in that book, through the influence of Wittgenstein he later came to believe that any statements that were in principle not related to observation, and thus in principle not verifiable or falsifiable, were meaningless. By this, Carnap explains, he meant that these statements were without cognitive, or factual, content.[8] That doesn't mean that had no effect on the hearer, but their effect is merely to stir up emotions or the like. Classic examples of such meaningless statements were statements about the reality of

[8] See, for example, (Schilpp 1963, 45).

the external world, both for and against.

According to Carnap, the only real advance of knowledge comes through science. The main contribution that philosophy could make to the overall scientific project was to develop and describe various languages that might be useful to adopt while formulating scientific hypothesis or theories. Adopting a language he also called adopting a framework. What these languages were didn't matter, as long as their syntactic features where well defined, and as long as it was well defined how endorsement or rejection of a sentence in this language related to experience. Which language we accept is matter of which one turns out to be most useful. In case new experiences force us to make a change in our overall theoretical system this could be done in two ways. Either we change what we endorsed within the old language, or framework, or we adopt a new language or framework and accommodate the new experiences in it.

In his paper "Empiricism, Semantics and Ontology" (Carnap 1956), Carnap addresses some questions that are most relevant to the present debate about ontology. In this paper he wants to dispel worries that scientifically minded empiricists, like Carnap, might have about using a so-called platonistic language, a language that talks about numbers, properties or the like. Carnap argues that the use of such languages for science is perfectly compatible with empiricism, and is, as we might say, ontologically innocent. As long as these languages are well defined we should feel free to accept them if they turn out to be useful. And to accept them is nothing but "to accept rules for forming statements and for testing, accepting or rejecting them." (Carnap 1956, 208). Which language should be accepted is a mere practical issue. And once a certain language is accepted, it will be a trivial fallout of this that we will also accept, say, that there are numbers. This will be a fallout merely of accepting the number language. But this, Carnap argues, can't be what the philosophers are after when they raise issues of ontology. That there are numbers in this sense is trivial. What the philosophers are after has to be something else, since they can't be after such a triviality. What they want to know when they raise ontological issues is whether or not numbers are part of reality, or whether or not the linguistic framework corresponds to reality. These questions, however, are meaningless according to Carnap, and are thus like the rest of metaphysics. The acceptance of the number language, however, has nothing to do with that.

According to Carnap, there are two different kinds of questions people can ask when they ask whether or not there are numbers. One is a question that has a trivial answer which comes with the acceptance of the number language, or number framework. These questions Carnap calls **internal questions**, because they are asked within a framework. The answer to internal questions about whether or not there are things of a certain general kind X is always trivially “yes”. It would be pointless to adopt a framework that talks about Xs, when according to the framework there are no such Xs. The other question, which he calls **external questions**, are questions that aim to go beyond this, and that try to, so to speak, look at the world from outside of the framework. These questions are meaningless, and these questions are what philosophers are really after when they do ontology.

Carnap’s views about ontology have been criticized by Quine, for example in his paper (Quine 1966). The debate between Carnap and Quine was a major philosophical debate, and it is widely believed these days that Quine showed that Carnap’s views were based on some basic mistakes. I won’t get into this debate here. In particular, the way in which I will defend Carnap will be independent of the apparent mistakes Carnap has made. I will only defend Carnap in spirit, not in letter.

What is important for us now is that Carnap would have a nice answer to the above puzzles and the dilemma, at least as far as he would consider them meaningful. In particular, Carnap would have a nice answer to the first puzzle: it is no wonder that on the one hand “ontological” questions have at the same time a trivial answer, and are almost impossible to answer. That is because there are really two different things going on when we ask whether or not there are numbers. One is to ask an internal question, which has a trivial answer, the other is to ask an external question, which has no trivial answer. This seems to be quite insightful. It would give sense to the feeling that it is one thing to accept that there are infinitely many primes, and to realize that this implies that there are infinitely many numbers, but still wonder about the ontological status of numbers, i.e. whether or not numbers are part of reality. If we might be asking two different things when we ask whether or not there are numbers this would not only explain our attitude that certain questions about the reality of numbers have remain unanswered even if we accept that there are infinitely many primes, and realize what this implies, and it would not be surprising that we end up with certain puzzles about ontology.

But how could it be that we are asking two different things when we are asking whether or not there are numbers? In particular, how could it be if we reject Carnap’s theory? And if we are asking two different things, what are they?

1.5 Outline of a somewhat neo-Carnapian view

1.5.1 A strategy for solving the puzzles

We will have to try to solve the above puzzles. The solution I will defend won't give up either one of the two answers of the first puzzle that appear to be in conflict with each other. I will endorse that the correct answer to the question "Are there numbers?" is hard to figure out, and I will endorse that the correct answer to the question "Are there numbers?" is trivial to give. But to endorse this is not to endorse a contradiction. I will argue that with the two answers to the question of the first puzzle we are really answering two different questions. I will give an account of what these two questions are, and how they arise. Furthermore, I will offer an account why we so easily confuse these two questions. It will require some persuading to pull this off.

With respect to the second puzzle, I will defend the natural philosophical attitude towards (certain) ontological questions. Those, who push ontological questions into the background, those who don't want to listen to the ontological worries in the number theory course, are in fact doing the right thing. However, to see why that is so we will have to see what is wrong with the answer that says that ontological questions are very important. Here my line will be as follows: yes, ontological questions would be very important, and the natural philosophical attitude towards them would be a mistake, if the ontology-objectivity dilemma held. That is, if the objective truth of our talk about properties, propositions and numbers required the existence of certain entities. I will argue, however, that this assumption, an assumption that, as we saw above, defines much of the contemporary debate in ontology and the philosophy of mathematics, is in fact false. I will argue that the literal and objective truth of our talk about numbers, properties, and propositions, has no ontological presuppositions. To defend this claim I will argue that the function of our talk about numbers, properties, and propositions, never was in the first place meant to be about a domain of objects. Rather, I will propose an alternative account of what the function of such talk is. The account I offer is not a global large scale account about the unimportance of ontology in general. It is explicitly limited to these three kinds of things: natural numbers, properties and propositions. I will not deal with hopes, ways the world might have been, novels, sets, events, facts, etc.. However, the arguments I use and explanations I offer seem to me to have the potential to deal with much larger domains of discourse.

My answer to the first puzzle will be Carnapian in a sense. I will argue, with Carnap, that there are two quite different questions we can ask, and do ask, about what there is. One of them is trivial to answer, the other one isn't trivial. Carnap thought that this other kind of question was meaningless, but I won't follow him there. These other kinds of questions are meaningful, but very hard.

My answer to the second puzzle will be closely related to this. It will be that when we usually and ordinarily talk about numbers, properties and propositions we do so in a way that is closely related (in a sense to be spelled out below) to the trivial way of saying what there is. Based on this and a number of other considerations I will argue that our ordinary talk about numbers, properties and propositions has no ontological presuppositions for its objective truth.

1.5.2 Four ontologies

Through most of this dissertation we will focus on our ordinary talk about properties, propositions and numbers: how we talk about them in ordinary, everyday life. After this has been addressed we will go on to look at how we talk about them in theoretical enterprises, like mathematics, philosophy and semantics. This order of proceeding is motivated by the structure of the overall argument, and is not supposed to be a value judgment about which one of these ways of talking about them is more important. To simplify the overall argument it is very helpful to make the following distinction: let's call the collection of things that is implied to exist by what we all commonly believe our **natural ontology**. The word "natural ontology" is thus used in the ballpark of the use of the word "ontology" in which it doesn't denote a discipline, but a collection of things. Our natural ontology is simply all the things that exists according to our shared beliefs. Thus cars, people and many more things are in our natural ontology.[9] God isn't in our natural ontology, since there is a division throughout us whether or not there is a God. Similarly for ghosts, aliens, and the like. Thus, if our shared beliefs were all true then our natural ontology would be part of reality, it would be a collection of some, but most likely not all, the things that make up reality. The main question for the first part of this dissertation will be the following:

(17) Are numbers, properties, and propositions, part of our natural ontology?

To answer this question will be closely related to a better understanding of what is going on in the trivial arguments.

[9] I ignore the occasional idealist who doesn't believe they really exist.

One could, of course, also speak of an individual's natural ontology. But to establish whether or not properties and the like are in an individual's natural ontology might depend on the specifics of that person. In particular, ontological beliefs themselves might become relevant. When we start with our shared beliefs we will not end up with anything question begging. But that should also be enough. No more than commonly shared beliefs seem to be required to argue for the existence of properties. After all, the trivial arguments seem to settle the issue from completely uncontroversial premises.

Besides all the things that are such that it follows from our shared beliefs that they exist there are, of course, all the things that do, in fact, exist. Let's call the collection of all things that do exist **true ontology**. If all our shared beliefs are true then our natural ontology will be part of true ontology. But there will be things that are part of true ontology that aren't part of natural ontology, simply because nobody believes that there are such things. And not everything that is in our natural ontology will be part of true ontology, we can assume, simply because some of our shared beliefs about what there is will be false.

Natural and true ontology are not all there is to ontology, of course. Besides natural ontology there is **theoretical ontology**. Theoretical ontology is the collection of all the things that we take recourse to in theoretical enterprises, including the natural sciences, semantics, philosophy and mathematics. Thus Montague's theoretical ontology contained intensional types, and von Neumann's contained Hilbert spaces, and the present philosopher of mind's contains propositions, and mental and physical properties. Theoretical ontology is usually assumed to contain things that our natural ontology doesn't contain.

Besides these two we should distinguish another one: **philosophical ontology**. Philosophical ontology is not a collection of things like natural and theoretical ontology, but a theory. It is a philosophical theory about the things that are in our natural and (partly) our theoretical ontology. Giving a philosophical ontology involves giving an account of what the things in our natural ontology are, how they relate to each other, how our ontological beliefs can be nicely systematized and the like. Virtues of a philosophical ontology are usually conceived as being simplicity, mathematical precision, clear criteria of identity, sparseness in the basic entities and the like. Much of the philosophical work in ontology is, not surprisingly, in philosophical ontology.

How natural, theoretical, philosophical, and true ontology relate to each other is a bit problematic and will be the subject of an extended discussion in section 6.3.

1.5.3 Outline of the dissertation

A central part of the present view is based on what seems to me to be a core mistake that both Carnap and Quine make, and that shows up in particular in their view on ontology. This mistake comes from a too simplistic picture of language. I have to note, though, that to accuse Carnap, in particular, of a too simplistic picture of language is quite ironic. Carnap's investigations into the syntax and semantics of natural and artificial languages were the most elaborate and sophisticated of his time. But a lot has happened since. When Carnap wrote about language there was none of the work by Chomsky, Montague, or Grice, and especially none of what happened after those three. In Carnap's time formal logic was still in its beginning, and probably partly because of the huge success it had in understanding natural language, it also led to some quite mistaken views, in particular about the relation between the formal languages that were so successful and natural languages. One key point, I think, where both Carnap and Quine were wrong is what contribution the context of an utterance makes to what is said with this utterance. For them the contribution was pretty much restricted to filling in the values of the indexicals and demonstratives. This, however, is not true. In the next chapter we will see that the contribution that context makes to what is said is much more subtle, and much more difficult to get one's hands on. This will not be an original argument, but merely a review of a number of well known cases that received a lot of attention in the linguistic literature. I will then argue, still in chapter 2, that such a subtle contextual contribution to what is said also occurs when we speak about what there is. This will be motivated independently of any issues about ontology. It is, I think, a fact about language that goes through and through and only shows up in a confusing way in philosophical discussions.

In chapter 3 we will have a close look at expressions like "being a dog", "that snow is white" and "the number of men". Independently of the issues from chapter 2 we will have to have a look whether or not these expressions stand for objects, or whether or not their semantic function is something else. Here, again, I will argue that the standard referentialist view of the semantics of such expressions is based on a picture of language that is too closely tied to first order logic. I will argue that the semantic function of such expressions, in particular in the examples that are most relevant for the philosophical debate, is rather quite different and is in the ballpark of issues that are usually treated in the interface between semantics and phonology. Once we look at in what situations one would use these expressions we will see what their semantic function is.

After we have achieved some preliminary results about quantification and nominalizations we can then look at talk about properties, propositions and numbers in particular. In chapter 4 I argue that the function of ordinary talk about properties and propositions is not to talk about a domain of objects but rather something else. Some apparent problems of the view that I propose will be discussed. At this stage we will be able to give a partial answer to our main question, whether or not properties, propositions and numbers are in our natural ontology. The answer will be: no. After having offered an account of the role of talk about property and proposition in everyday life we will have a look at what role such talk has in philosophy.

In chapter 5 I address the same issue for numbers. However, no simple extension of the results of chapter 4 to numbers is possible. We will have to address a whole different set of issues, but one that will be very helpful for our overall picture. We will have to talk about the uses of words like "two" in ordinary discourse and in mathematics, and in particular how they relate to each other. A central part of the account of their relation will be an analysis of how basic arithmetic is learned. The answer to the main question will be extended to natural numbers.

In chapter 6 I will give a detailed account of how the view developed so far answers the puzzles we started with, how it accounts for the occurrence of the puzzles, and how it sheds light on what is going on in contemporary ontology. In particular, the relation between the four ontologies will be discussed. We will also have a look at some classic philosophical problems that closely relate to the main themes of this dissertation, in particular the problem of universals. After that some related views will be critically evaluated.

Finally, I will deal with a number of more technical issues that arise as we go along in an extended appendix.

While working on this dissertation I have become aware of three other approaches to revive ideas that Carnap had about ontology. All three of these are different, and all are different from the present approach. Robert Kraut defends a view according to which ontological claims are not factual claims, but have their function to articulate commitments to certain frameworks.[10] Huw Price endorses neo-Carnapianism in the sense that

[10] None of this has been published so far.

many different frameworks can all give rise to a true description of reality, which he calls "metaphysical pluralism".[11] Finally, Stephen Yablo has defend Carnap's internal-external distinction by associating it with the literal-metaphorical distinction. Yablo's version is close to the present one in overall motivation and in an endorsement that there is a certain natural puzzlement that one should have about ontology as done today. However, the way in which the neo-Carnapianism is carried out is quite different. For instance, Yablo endorses the above statement (A), which seems to me to be a central mistake that gives rise to the puzzles. Yablo's work will be discussed in section (6.9). There I will also discuss Stephen Schiffer's work, in particular his book *Remnants of Meaning*, which is in certain respects closely related to the present approach.

[11] See (Price 1992).

Chapter 2

Quantification

2.1 Quantification and ontology

One of the key pieces in understanding the above puzzles, and ontology in general, is to get clear about the relation that quantification has to ontology. In the above trivial arguments it was a key step to make the transition from a sentence without a quantifier to one with one, as in

(18) Fido has the property of being a dog.

(19) Thus: There is a property Fido has, namely being a dog.

And the latter seems to have direct ontological relevance since it immediately implies that there are properties, which we, again, express using a quantifier.

In particular, what quantified statements follow from what we commonly believe seems to have direct relevance to what is in our natural ontology. After all, if what we commonly believe implies that something is F, then it seems clear that our natural ontology contains Fs. Similarly, if it implies that there is a G then our natural ontology contains Gs. It seems reasonable to assume that there is a very close connection between quantification and ontology in this sense. What quantified statements are implied by a set of beliefs makes explicit what the ontological presuppositions of these beliefs are for their literal and objective truth.

But on the other hand, and as we saw at the very beginning of chapter 1, there do also seem to be a number of cases where we use quantified statements, but we apparently do not mean them to have any ontological implications. It seems that when I say

(20) There are three cars in the parking lot.

I do mean to make a claim that is only literally true if reality contains certain entities. But that is not at all clear when I say

(21) There are three ways to get to Berkeley from here.

This might, of course be mistaken. Maybe (21) is exactly like (20). Maybe both of them do require that three entities of a certain kind are part of reality. There are, however, a number of cases where it prima facie doesn't seem to be right to say that. Sometimes it seems we use quantifiers differently. And these uses relate to a number of classic philosophical puzzles, for example when we talk about what doesn't exist, as in

(22) There are three things we don't agree on whether or not they exist: Santa, Sherlock and the Easter Bunny.

To point out to the speaker that he is thereby quantifying over these exact things and thereby settling the dispute always has a ring of philosophical sophistry to it.

In this chapter we will have a closer look at quantification. In particular, we will have a look at for what purposes we take recourse to quantified statements in regular communication. And once we have seen this we will understand better what relation quantification has to ontology.

I take it beyond doubt, and in no need of subtle philosophical investigation, that quantification does have a close connection to ontology in the following sense. There are a number of uses of quantifiers where the utterances in which they occur are only literally and objectively true if reality contains certain entities. These are uses as, for example, in

(23) Something fell on my head.

A standard utterance of (23) is only literally true if reality contains a thing that fell on my head. Thus it has ontological presuppositions for its literal and objective truth. However, and this is the more controversial part of what I'll argue for in this chapter, there is also a different use we have for quantifiers. In this use they do not have a direct connection to ontology. So, I will argue that quantifiers in some uses do have a direct connection to ontology, and that in some uses they don't. In some uses of quantifiers we use them for a purpose quite similar to the one we use them for in (23), but in other uses we use them for a quite different purpose. But how is that supposed to go? And am I just saying this to avoid ontological commitments that I don't have the guts to face? No. I think that if we look at what we do when we use quantifiers, for what purpose we take recourse to them, and what function they have in communication, we see that there are at least two, and that only one of them has a direct connection to ontology. When I make this claim I don't take myself to be making a philosophically motivated claim about language. I will spell out how this particular claim about quantification is rather an instance of a much more general feature that natural languages have, one that unfortunately hasn't been widely recognized within philosophy. In particular, I will argue for this claim that there are more than one function that ordinary quantifiers have not on philosophical ground, (as I would if I would say that for reasons of parsimony we have to avoid commitments to certain entities), but rather on empirical grounds. In the following I will try to persuade you on empirical and linguistic grounds that there are two uses we have for quantifiers, and that only one of them is directly connected to ontology. I will do this by saying what these uses are, and in particular that one of them only works the way it is supposed to if it isn't directly connected to ontology. To do this I will first talk about a somewhat general picture of language, and in particular the role that context has in determining content. Then we will get back to quantification, and finally we will get back to ontology.

2.2 Some mistaken views in semantics

In this section I will point to two mistakes in semantics that are very relevant to our present discussion. When I call them mistakes in semantics I don't mean to suggest that present day semanticists would make these mistakes. Not at all. In fact, present day semantics has uncovered that these are indeed mistakes that have been made in the recent and not so recent past. But the lessons from these mistakes have not sunk into much of the philosophical debate, in particular not into the debate about quantification and ontology. Maybe because it didn't seem relevant. To show that and how it is will be the main work in this chapter.

2.2.1 The role of context

It was a common belief not all that long ago that the contribution context makes to the content of an utterance is basically to fill in the values of the demonstratives and indexicals that might occur in that utterance. However, there are a number of very plausible example that show that this is false. Contextual contributions to content are much more common, and much more subtle. What this means is that what is contributed to content from the shared language, meaning, is a lot thinner than one might think. In this section I will review three cases of this phenomenon: genitive, plural and reciprocals. These are well known cases in the linguistic literature. That these cases and the general phenomenon is relevant for our debate about ontology will be argued for later.

There are a number of obvious ways in which context contributes to what is said, and a number of more subtle ways. One obvious way is that it has to contribute the value of demonstratives and indexicals. The meaning of the word "I" or "this" alone is not enough to make the contribution to content that an utterance of this word makes. It is the context of the utterance together with the meaning of the word that make the contribution to content that it makes.

When I speak of a contextual contribution to content I use this very widely. With this I mean anything that is part of the content of an utterance, what is said with it, that is not already determined by the meaning of the words uttered. The meaning of the words uttered is whatever is fixed at the level of shared language. So, one might think that our shared language has words in them that stand for individuals or properties, and that an utterance of a grammatical sequence of such word has the content that certain objects have the properties that the words stand for. If some of the words for individuals are indexicals and demonstratives then we get the picture that the contribution of content of the utterance to content is to help fix what individuals the indexicals and demonstratives stand for. Thus we get following simple picture of what the context contributes to what is said:

- **The simple picture** An utterance of '$P(c_1, \ldots, c_n, i_1, \ldots, i_n)$' is true in context C iff $P(c_1, \ldots, c_n, F_C(i_1), \ldots, F_C(i_n))$, whereby c_k is a constant, or name, and i_k is an indexical or demonstrative, and $F_C(i_k)$ is the value of i_k in context C.

To also include quantifiers, we have to allow for contextual restrictions of the quantifiers, which is another commonly acknowledged way in which context can contribute to content.

Thus we have to extend the simple picture to allow for the following case:

- **The extended simple picture** An utterance of a quantified statement, say, "Every A is B" in context C is true iff every A which is F is B, whereby C determines what F is. Similarly for "Some A is B", and others.

In the following I will point to three rather well known cases where there are much more subtle contributions that context makes.

Genitive

Consider a standard use of genitive, like

(24) Joe's car has a flat.

It seems that what the genitive "'s" does in (24) is that it contributes to what is said with an utterance of (24) that Joe *owns* a certain car. It is the car that Joe owns that is said to have a flat. Thus it seems that what the genitive does is contribute the relation of ownership that holds between Joe and a certain car to the content of what is said.

But as it turns out, that is not always so. The genitive can be used in many different ways, and it can contribute many different relations to what is said. For example, in

(25) Joe's book is full of mistakes.

on would usually understand this as saying that the book that Joe wrote is full of mistakes, not the on he owns. But it can also be used to say that the one he owns is full of mistakes. And it can be used to say many different things. Consider, for example, the following situation:

(26) *At the beginning of the academic year the department requires that all graduate students meet in a room and bring the library copy of a book they read in the library last year and liked a lot. In the room the grad students sit around the table with library copies of books in front of them. Noticing the book grad student Joe brought to the meeting, one of the professors says to another one:*

Joe's book is full of mistakes.

In this situation what was said is that the book that Joe picked, or the one that is in front of him, is full of mistakes. It is clear in the context that Joe is neither the author nor the owner of that book. Furthermore, one can construct situations like this for almost any relation you please. A certain occasion of an utterance of a sentence with a genitive in it will contribute that relation to what is said. Take any relation that can reasonably be said to hold between a person and a book. We can find a situation where it will be clear that an utterance of "Joe's book is F" has the content that a book that has that relation to Joe is F.

Thus the contribution that the meaning of the genitive makes to what is said is not a particular relation. It is rather that there is some relation or other that is said to hold between, say, Joe and a book. What relation this is will be a contribution that the larger context of the utterance makes. This will partly be our knowledge about what kinds of relations usually hold between people and books, or what kinds of relations between people and books have been talked about just a minute ago, or the like.

Plural

Another, very widely discussed, example of a not immediately obvious kind of contextual contribution to content is plural. Consider the following example:

(27) Four philosophers carried three pianos.

The plural phrases "four philosophers" and "three pianos" can each make at least two different contributions to the truth conditions of the utterance. First, they could be read as being about individual philosophers, or individual pianos. Secondly, they could be read as being about a group or collections of philosophers or pianos. Thus an utterance of (27) can have at least the following truth conditions:

(28) a. Four philosophers together carried three pianos each (one after the other).

b. Four philosophers each carried three pianos each (one after the other).

c. Four philosophers together carried three pianos together (all of them at once).

d. Four philosophers each carried three pianos together (all at once).

Now, given what we all know about pianos and the strength of philosophers, the most natural way to understand an utterance of (27) is of course (28a). But an utterance of (27) can have the truth conditions as spelled out in (28d). That it usually doesn't comes from the fact that we talk about pianos. If we talk about other things then this will be the default reading, as in

(29) Four philosophers carried three books.

An utterance of this sentence will usually have the truth conditions that four philosophers each carried three books together.

This is a general phenomenon about plurals. They at least have a **collective reading**, being about a collection or group of things, and a **distributive reading**, being about the individuals in that group. At the level of language it is not determined which one of these two different contributions to the truth conditions a plural phrase will make in a particular utterance. This is determined by the context of the utterance. We can thus say that plurals are semantically underspecified in a certain respect. The semantics of such phrases, what is determined at the level of language, doesn't specify whether the phrase is about a collection of things, or about the things in that collection. This has to be determined by other features of the utterance. How this determination will work is, of course, a very tricky business, and we won't get into the details of this here. What matters here is the general phenomenon of semantic underspecification.

If a sentence contains a semantically underspecified item in it then an utterance of that sentence will have more than one reading. It will be possible to utter it with at least two different truth conditions. This is like ambiguity, but different in certain important respects from (standard cases of) ambiguity. In the case of ambiguity, too, one can utter a phonetically identical sentence with different truth conditions. But in the case of semantic underspecification we are dealing neither with lexical ambiguity, nor with structural ambiguity. It isn't the case that any one of the words in (27) has two different meanings, nor that the sentence can have two different structures. Rather, some of the items in the sentence are not semantically completely specified . The context of the utterance will have to fill in the details that the language left out.

Reciprocals

Another case of this phenomenon is the case of so called reciprocal expressions, expressions like "each other" or "one another". A sentence involving these expressions will specify some collection of things, and some relation that can hold among the things of that collection, and the reciprocal expression will specify how the things in that collection stand to each other with respect to that relation. In a simple case this goes as in

(30) The Smiths like each other.

The collection here consists of certain people, the relation is liking, and what is claimed is that each one of the Smiths likes each other one (except maybe themselves). However, "each other" does not always contribute the same to the content. For example, consider the following pair of standard utterances of the sentences:

(31) The people in the room tried to be as far away from the wall as possible. In the end they were no further than one yard from each other.

(32) The exits on the Santa Monica freeway are no further than one mile from each other.

In the case of (31) it is required for this utterance to be true that everyone in the room could touch everyone else by extending their arm. But in the case of (32) it would still be true even if the first and the last exit are 30 miles apart. All that is required here is that there is another exist every mile.[1]

Semantic underspecification

So, we have seen that the naive view about the role that context plays in determining the content of what is said is too naive. Context plays a much more subtle role than fixing the values of indexicals and demonstratives, and contextually restricting quantification. On the contrary, there is a general phenomenon of semantic underspecification that occurs in several different classes of examples. In these case the contribution that an expression will make to the content of an utterance is not completely determined at he level of language. The language only specifies certain options, the context has to do the rest. This should not be very surprising. It is most likely a very effective and useful way to communicate with

[1] See (Kim and Peters 1995) for a discussion of such cases.

such a language. Just think about how we would clutter our lexicon if we had a different word for each of the different contributions that "each other" can make to the content. Or how much more complicated communication would be if we would have to make things explicit in words that we in fact leave implicit and that are salient in the context of the utterance.

The above examples were examples where some of the words we used were semantically underspecified, but the utterances of the sentences had (in the situation we imagined) fully specified truth conditions. But this doesn't always have to be so. Sometimes we utter sentences and the utterances do not have fully specified content or truth conditions. When we ask: what precisely is the correct truth condition of this utterance in this context, TC_1 or TC_2? it might very well be that there is no correct answer. The truth conditions aren't specified in all details. And this, again, should not be seen as unreasonable. In communication not all the details matter, and in fact only rarely will details beyond a certain point have any importance. To illustrate this, when I say

(33) Germans don't like each other.

am I saying that no German likes any other German, including themselves, that some Germans don't like some other Germans, that many Germans don't like many other Germans (and if yes, how many), etc.? We can spell out many different precise truth conditions that an utterance of this sentence might be taken to have, but after a certain level of specificity there won't be a correct answer to the question any more which one is the correct truth condition of this utterance. And, again, that is only reasonable. These details won't matter in regular communication, and they will be left undetermined.

So, regular English expressions are often semantically underspecified, and subtle features of context fill in the details, or leave certain details undetermined. This will be quite useful in understanding our overall concerns about ontology. But first one more general point about quantification and semantics.

2.2.2 Two mistaken views about quantifiers

The following two views about quantifiers arise out of the first order logic based approach to quantification, and, it seems fair to say, are the ones that dominate the philosophical literature on quantification from Quine onwards. They are as follows:

- There are basically only two quantifiers, the universal one and the existential one. They are each expressed with different words, like "all" and "every", but basically they are the same.
- Quantifiers are semantically fully specified. They are thus not the kind of expressions that show any context sensitivity. The contribution to content that a quantifier makes is always the same (even though how it is restricted isn't the same).

It seems fair to say, though, that both of these points are in fact false. That the first point is false has given rise to a major trend in semantics in the 80's, and gave rise to what is called generalized quantifier theory (GQT). It is based on the observation that expressions like "some" and "all" really belong to a much larger category of expressions, and that this category can be given a unified syntactic and semantic treatment. They belong to the same semantic category as expressions like

- many
- few
- one, two, three etc.
- some, but no more than five,
- finitely many, etc.

These expressions, called determiners, together with nouns form quantified noun phrases, like

- many men
- few men
- one man, two men, etc.
- some, but no more than five, men
- finitely many men, etc.

Expressions like “something” and “everything” also belong to this category of quantified noun phrases. Thus the category is quite a bit larger than just the standard first order quantifiers, and it in particular contains plural and singular expressions.

Given this last point, and what has been said before, it should be clear that the second of the above two points is false, too. Not all quantifiers are fully semantically specified. Plural quantifiers are semantically underspecified at least with respect to having a collective and a distributive reading. This, of course, isn’t to say that all quantifiers are semantically underspecified. But some certainly are.

With this background information on quantifiers, semantic underspecification and subtle contextual contributions to content in mind, let’s now return to our discussion about quantification and ontology.

2.3 Quantification without ontology

2.3.1 Outline of the two readings

I said I would argue that there are uses of quantifiers that have no direct connection to ontology. With the things that have been said in the above sections in mind, I will formulate this as follows: I will argue that ordinary quantifiers, like “something”, are semantically underspecified in a certain respect. They play two different, but related, roles in communication. In particular, such quantifiers do not always make the same contribution to the truth conditions of an utterance in which they occur. At present, let’s focus just on “something” and “someone”. We will look at other quantifiers later. The claim I will argue for now is thus that this quantifier is semantically underspecified in a particular way. To argue for this I will point to a general situation in communication where we have a need to communicate certain information that we have, and where we can communicate the information we have using quantifiers. In this situation the quantifiers make a different contribution to the truth conditions than they do in other situations, like the one mentioned above where quantifiers have an obvious relevance to ontology. Let me first, before I will argue that there indeed are (at least) two such contributions to the truth conditions that quantifiers can make, say a bit about the one that they quite obviously make.

Quantifiers quite obviously have what I will call a **domain conditions reading**. According to it they contribute in a certain way to impose a condition on the domain of discourse that the domain has to satisfy for the utterance to be true. To put it simple, an utterance of “Something is F” is true just in case the domain of discourse satisfies the following condition: it contains an F. And the domain of discourse is whatever is out there in reality, all the things that make up the world, or a contextual restriction thereof. In these uses of quantifiers they have obvious ontological relevance, since the truth of an utterance in which they occur require reality to satisfy a certain condition with respect to what kinds of things it contains. And this seems to be just what we are doing when we utter

(23) Something fell on my head.

For reason that will become obvious in a few pages, I will also call the domain conditions reading the **external reading** of a quantifier. The truth conditions of quantifiers in their external reading are quite generally captured in generalized quantifier theory, and I do not want to take issue with this. In their external reading quantifiers are well understood and sufficiently described. The question we have to address is whether or not quantifiers are always used in their external, domain conditions, reading.

In the following section I will motivate that quantifiers also have a different reading. For reasons to become obvious soon, I will call it the **inferential role reading** or also the **internal reading**. This reading makes a different contribution to the truth conditions than the external reading. And there is a need we have for it in communication that is not served by the external reading.

2.3.2 Motivation: communicating in ignorance

To argue that quantifiers can make different contributions to the truth conditions is in effect to argue that we get more expressive power through quantification than one might think. And to motivate that this is so it is helpful to look at a situation where we need increased expressive power, and to show that in this situation the ordinary expressive power wouldn’t be enough to communicate what we want to communicate. In particular, to argue that quantifiers can make different contributions to the truth conditions it is a good strategy to try to find a situation where they would get different truth values. In particular, to try to find an example of an utterance of a quantified statement that has a certain truth value, and show that if the quantifier would make a certain contribution to the truth conditions then the truth value would be different. This is what we will do next.

Since I will argue that quantification gives us more expressive power than one might think, it is a good idea to look at situations in communication where we need increased expressive power to find ones where we can see that we use quantifiers in a different than the "standard" way. One such situation is communicating in ignorance. In the debate about the function that a truth predicate has it was often pointed out that one of the cases where we need to take recourse to a truth predicate in communication is when we try to communicate certain incomplete information. For example, if I do not know what the Pope said, but believe that whatever he said is true then I can only communicate this by saying

(34) Everything the Pope said is true.

If on the other hand I also happen to know what the Pope said then I can just repeat whatever he said, and say that this is all he said. So, if I have more information I can communicate the information that I communicate with (34) without using a truth predicate. However, if I only have certain incomplete information then I need a truth predicate. So, without the expressive power that a truth predicate gives us I can't communicate certain incomplete information I have. One of the needs we have for a truth predicate in communication is to communicate in ignorance.

I will also use this strategy here. I will look at communication in ignorance to point to certain expressive needs in communication. In particular, I will argue that these needs will motivate that quantifiers are semantically underspecified. The case of ignorance that I will look at is a special case of ignorance, the case where one has known partial information. This is the case where one does have only partial information, but knows that there is a certain aspect of it that one is missing, and one recognizes which aspect this is. In particular, I will consider the situation of forgetting part of the information that you had earlier, but at the same time realizing that one forgot part of it and which part it is.

Consider the following example. You are working on a psychological profile of Fred and one day you find out that

(35) Fred admires Clinton very much.

This is most useful information for you since it allows you to infer a lot about what kind of person Fred is, what character traits he values, and the like. The next day, however, when you actually sit down to write up the profile you simply can't bring yourself to remember who that was that Fred admires very much. You still remember that whoever it is that Fred admires is also admired by many democrats. But you just can't remember who that is, even though you realize very well that you knew this yesterday. What you know today is less than what you knew yesterday, but it is still a lot more than nothing. What you know today still gives you lots of useful information about Fred. It lets you conclude that the character traits that Fred admires in a person are much like the ones that many democrats value, and so on and so forth. And the incomplete information that you still have and that is useful to you can be communicated. You can do so by uttering

(36) There is someone that Fred admires very much, and who is also admired by many democrats. I just can't remember who that is...

So, if you lose the information that you did represent with, say, a name then you will have to use the quantifier "someone" or "something" to communicate the incomplete information that you end up with. And you do have a need in certain situations in communication to do so.

Now, this situation is completely general. It also applies if what you learn about Fred in the first place is

(37) Fred admires Sherlock Holmes very much.

Again, you might forget who it was that Fred admires very much, but you might still remember that whoever it was, that person is also admired by many detectives. Again, this is most useful information for you since it allows you to infer a variety of things that might be most useful. And, again, you can communicate the incomplete information you have by uttering

(38) There is someone that Fred admires very much and who is also admired by many detectives. I just can't remember who that is...

Whether or not Fred admires someone real doesn't matter for our situation here. What does matter is that we want to, and have the ability to, communicate certain useful incomplete information. In our case it is the information that whoever Fred admires is also admired by certain other people. Whether or not whoever is admired is someone real or not is of no interest here. We don't want the truth of what we say to depend on whether or not the person admired is real. That is just not relevant right now. Note now that if the only uses of quantifiers we had where the ones that are like the domain conditions, or external, reading then whether or not whoever is admired is real would be most relevant for the truth of (38). A quantified statements with a quantifier in the external reading in it is only true if the domain of discourse, all the things that make up reality, satisfy a certain condition. In the case of (38) the condition is to contain a thing which is admired by Fred and certain other people. Thus in the external reading (38) would only be true if Fred and these other people admire someone real. But this is not what we want here. In particular, we seem to be able to communicate the information we have just fine whether or not who is admired is real.

Consider a quite similar situation that also makes this point. You forget who it is that Fred admires. But this was very important to you, and you are angry at yourself for forgetting. However, you still have a number of clues about who that was. And you think that by reminding yourself what they were you might bring yourself to remember who it was that Fred admires. You start thinking about what you still remember about who it is that Fred admires, and you say to yourself (or somebody else who is supposed to help you):

(39) So, let's think hard. There is someone Fred admires very much, and whoever it is is also admired by many detectives. In addition, that person supposedly smokes a pipe. Who is that, again?

Consider also what would happen once you remember. Suddenly it appears to you:

(40) Of course, it's Sherlock!

Would you then say that what you said first is false? Would you now realize that your above reasoning or reminding yourself was really the uttering of a number of falsehoods? I don't think so. What you said was perfectly fine, and correct. Whether or not who was admired is real didn't matter, and shouldn't have.

To see more clearly what is going on in these examples, let's look at several aspects of them separately. Let's have a look at what we are trying to do when we use quantifiers to communicate incomplete information of the above kind (the goal), how these quantifiers allow us to do this (the means), and what assumptions I am making to use this situation as an argument for the thesis that ordinary quantifiers have more than one reading (the assumptions). We will address each one of these three issues separately.

2.3.3 Analysis: the internal reading

The goal

What are we trying to do when we use a quantifier as in (38) or in (36)? Thinking about the above situation, we can see the following. We learn a certain information that we represent linguistically using, amongst others, a proper name, or a term. Even when we forget the part of the in formation that corresponds to the contribution that the name or term makes to the overall information, we still have useful information, but now it has a certain gap. Whatever was contributed by the name or term is now missing. But all the other parts are still there. When we try to represent the information we now have linguistically we can't just articulate a sentence and simply leave out the term or name. We can't just say

(41) Fred admires . . . and . . . is also admired by many detectives.

This is nonsensical. We have to put something else in the gap that the forgetting has left. But since we have lost all the information that the name or term contributed we have to put something in that is neutral with respect to contributing any information. This is exactly what the quantifier in its above, internal, use does. It fills the gap in the representation of incomplete information, a gap that has been left by the forgetting of a certain part of the information that one once had, the part that was represented using a name or term. But in doing this the quantifier doesn't contribute any special information. It is neutral. It simply is a placeholder for the forgotten part that we put in to make the sentence we utter grammatical, and to indicate that this is the place where the information is only partial. How can a quantifier do all that?

The means

What we want the quantifier to do is at least the following. When we forget we lose information. Before we forgot who was admired we know than we know now. Thus the linguistic representation of the complete information has to be truth conditionally stronger than the linguistic representation of the incomplete information. That is to say, a statement that represents the complete information has to imply a statement that represents the incomplete information. Simply because the latter is less specific information, and thus whenever (a representation of) the former is true (a representation of) the latter has to be true, too. Note also that once you forgot a certain part of the information you had you end up in the same state of incomplete information, now matter how you started out. So, whether or not it was Sherlock, Clinton, the first man on the moon, etc., that is admired by Fred, once you forgot who was admired you end up with the same incomplete information, namely that there is someone Fred admires very much. So, a linguistic representation of any one of these complete information states will imply one and the same linguistic representation of the incomplete information. After all, you end up in the same incomplete information.

In a word, what we want from the quantifier in the above case is that it has a certain inferential role. What we want from it is that the inference from

(42) Fred admires X.

to

(43) There is someone Fred admires. (but who?)

goes through independently of what or who X is. Thus the quantifier has to contribute in a certain way to the inferential role of the sentence in which it occurs. Simply put, the quantifier "something/someone" has to have the inferential role such that

(44) F(n) implies something/someone is F.

In particular, this inference has to be valid independently of what n is. Whether or not n is real doesn't matter for the purposes for which we want the internal quantifier. If the quantifier would be read in its external reading then this would, of course, matter. Whether or not n is real is precisely what matters for whether or not "F(n)" implies "Something is F", used externally. Thus the quantifier in its external reading doesn't give use the above inferential role. But in the situation of communicating in ignorance as spelled out above we want and need this inferential role to do what we do. The external quantifier can't play this role.

The assumptions

I claim that the external or domain conditions reading doesn't give us for what we want the internal, or inferential role reading of the quantifier. In particular their contributions to the truth conditions have to be different. I will elaborate on how they differ in truth conditions in a minute, but first I'd like to say more about the assumptions that I make in arguing for this point. Some of them will be discussed at some length later in this chapter.

What I was assuming for the above example (38) to work can be labeled as **partiality**: I was assuming that not every name in our language stands for some thing that is part of reality. Some names stand for nothing, they are empty names. In (38) I took "Sherlock Holmes" to be such a name, but, of course, nothing hangs on this name being empty. Any other one would do just as well. Whether or not there are any empty names is controversial. Some people believe that names like "Sherlock Holmes" really do succeed in referring. Most people who believe this will also believe that what these names refer to are some kind of abstract object. I doubt that this is correct. I think we prima facie have lots of reason to believe that there are empty names in our language. I'd like to point out two especially:

- Whatever the mechanisms of reference are, however we manage that our names stand for things in the world, these mechanisms can break down. We are fallible creatures that make mistakes, and such mistakes can bring it about that a certain name fails to stand for an objects, even though it might have if we wouldn't have made that mistake.
- We sometimes intentionally make up names that were never meant to stand for anything. We make up stories to entertain ourselves, and when we do so we never in the first place try to talk about anything real. It would be quite surprising if we would after all talk about something real even though we never tried.

These points are, as I said, controversial. I'd only like to point to prima facie reasons why we should believe them. More importantly, however, these points are not needed for our purposes here. It is not necessary to take recourse to empty names to motivate that there are two uses we have for quantifiers. We can just as well take recourse to non-denoting descriptions. Consider the following modification to our above example. What you learn about Fred is that

(45) Fred admires the inventor of the wheel very much.

The next day you forget who that was that Fred admired very much but you do remember that whoever it was is also admired by many bicyclists. You can communicate this with uttering:

(46) There is someone Fred admires very much and who is also admired by many bicyclists. I just can't remember who that is...

I assume here that no one really invented the wheel. Either there were many people who invented it at the same time, or, more likely, no one invented it at all, but rather people adopted using round things gradually for certain purposes. The above example (46) has for our purposes all the same features than (38) but it doesn't take recourse to empty names. Taking recourse to them is an easy way to give an example of the above kind, but nothing essential about the example.

Another assumption that I made in the above example is that one can truly say that Fred admires Sherlock. This might seem controversial to people who want to take it very seriously that admiring is a relation that people have to other things, and therefore can only hold between two things that are real. I don't think this is right. More on this will be said in section 2.4.

Finally, it might strike you as quite problematic to think that one should articulate the incomplete information one has with

(47) There is someone Fred admires very much.

and not rather with

(48) Fred admires someone very much.

I will address the issue of their difference at some length in section 2.4. There I will argue that the difference between these two isn't a difference in scope, as one might think. I will argue that in fact both of them exhibit the same scope phenomena. There difference is in a different ballpark, and it will have nothing to do with our discussion about ontology. But to address this fully now would be too much of a distraction. We will get back to this.

2.3.4 Some remarks on the example

Let me make the following remarks on the example given above to make some things about it more explicit:

- The example does not rely on intuitions about the truth value, or truth conditions, of certain statements. The strategy was not to exhibit a certain sentence and then pump your intuition that it is true in a certain situation. This would be a poor strategy since such intuition might not be a good guide towards the truth. On the contrary, the example is an example of a general situation in communication, and it is supposed to illustrate a general need we have for expressive resources. I wanted to point out that we often have a need to communicate incomplete information and that the expressive resources we get from external quantifiers won't be good enough to do what we want to do. However, in such situations we still take recourse to quantifiers to successfully communicate the incomplete information we have. This is supposed to motivate that quantifiers are not always used in the same way. No intuitions about isolated examples are required.

- Taking recourse to fictional names is only a tool to motivate that there are two uses of quantifiers. The internal use is in no way only connected to quantification in the context of fiction, or only to quantification over fictional objects. The two uses that quantifiers have apply generally. Even a sentence like

  (49) There is someone Fred fell over.

  can be uttered to have two different truth conditions. However, in the case of (49) a competent speaker will utter it in such a way that the quantifier is used externally. Furthermore, whether or not it is used externally or internally will have no effect on the truth value in this case. Why will be spelled out more full below, but assuming that Fred fell over Jim then both the inferential role reading and the external reading will be true. The first because it is implied by the true "Fred fell over Jim", and the second because there is a thing out there in reality over which Fred fell.

- The above proposal that quantifiers are semantically underspecified and have two different readings, is not a radical proposal about language that I adopt for philosophical reasons. It is a proposal that fits nicely into a much more general framework about the role of context in determining content, a framework that has been widely accepted independent of all philosophical considerations.

- Finally, the scope of the proposal is, of course, not restricted to the above example, or examples like it, involving intentional verbs. This is just taken recourse to to get a difference in truth value with the different readings, just like taking recourse to fictional names. The above proposal applies very broadly. Any sentence involving "something" has two readings, since the two readings come from the quantifier. Whether or not the rest of the sentence is about fictional objects or involves intentional names is inessential for that. But it is in these cases where it can be made apparent in the simplest way that there indeed are two readings .

2.3.5 Inferential role and domain conditions

We have a need in communication for a phrase that plays a certain inferential role. And we have a need for a phrase that imposes certain domain conditions.[2] Both of them play an important role in communication. And as I said above, imposing certain domain conditions doesn't give a certain inferential role. The domain condition that "something" imposes on some of its uses doesn't give it the inferential role that I claim it has in other ones of its uses. However, this is only so given certain features of the language in which "something" occurs, features that our natural language has. In simple languages these two, inferential role and domain conditions, can coincide. In fact, in they do coincide, or almost coincide, in many of the languages that are studied in formal logic, and in many of the formal languages that are used for natural language semantics. Consider, for example, standard first order logic. In it every name (or constant) stands for some thing in the domain of every model of the language. Under these conditions, the inference from "F(t)" to "$\exists x$ F" will always be valid. After all, every term stands for some thing in the domain of a model. Under these circumstances one and the same contribution to the truth conditions can give rise to both, a certain inferential role and certain domain conditions. But as soon as these simple circumstances don't obtain any more these two will go apart. If the language contains empty names, or intentional verbs, then no one contribution to the truth conditions can

[2] When I say that we need a phrase that does certain things I mean that we need a phrase such that certain occurrences of that phrase in an utterance do that thing in that utterance. When I say that we need a phrase to do A and one to do B it is thus still open whether or not the same phrase sometimes does A and sometimes does B. A phrase is only a string of words or symbols in a language. I will not always make all this explicit in what is to come.

give rise to both, the inferential role and the domain condition. For example, if we have a free logic, and we allow for terms that do not stand for anything in the domain of discourse then this inference doesn't go through any more. For it to go through we need the further premise that "t" does stand for something in the domain of discourse. But this need for a further premise is to say that the quantifier does not have the inferential role any more that it has in non-free logic. Since natural language do contain empty names, and terms that do not stand for anything in the domain of discourse we have that in the situation of natural languages inferential role and domain conditions come apart.

The fact that they do coincide in simple cases is still of importance even for understanding the case of natural languages. Even if I am right that we have a need for a certain inferential role and certain domain conditions, and that these two do not coincide with respect to truth conditions, we still have to ask ourselves why one and the same phrase sometimes plays one of these roles and sometimes the other. Why don't we have two phrases that play these two different roles? Well, there is of course no a priori guarantee that we don't, but there is a nice explanation for why we in fact don't. The explanation is quite similar to the one why we in fact don't have two different expressions for plural, one for plural in its collective reading, and a different one for plural in its distributive reading. The explanation for this can be given along the lines that there is a very close connection between the collective and the distributive reading of a plural phrase. The collective reading is about a certain collection of things, and the distributive reading is about the members of that collection. Now, there is a very close connection between a collection and its members. Simply put, whenever you have one you have the other. They are not completely independent things. So, it seems only economical that a language only has one phrase for the plural, and that the context of the utterance determines whether or not this particular use of the phrase is about the collection or about the members of the collection. It seems reasonable that this is, overall, a much more economical way to communicate than to have different words in the language that each are context insensitive. To be sure, this is only a sketch of the story that has to be told here, but it makes plausible why we have semantic underspecification in the case of plural but not of a similar kind in the case of singular. And similarly we will be able to explain why one and the same phrase is semantically underspecified with respect to whether it plays a certain inferential role or imposes certain domain conditions. These two are not unrelated. In fact, they do coincide in simple circumstances. So, it may be most economical for a language to have one phrase that is underspecified with respect to which one of these two related roles it is supposed to play on a particular occasion of an utterance. And this is exactly what is the case.

2.3.6 The basis of the inferential role

So far we have only said that we want a phrase that is a placeholder for the forgotten part of the information, and that in order for some phrase to be able to do this it has to have a certain inferential role, an inferential role that allows for inferences independently if one talks about something real or not in the premises. But we haven't really seen how the quantifier in its internal reading achieves this. We have seen that it can't do this by imposing domain conditions. But how else does it do it? I will elaborate on this in this section.

There are basically two well known options from the philosophy of logic when we try to give an account of how a phrase comes to have the inferential role it has. These are:

- Either the phrase has this inferential role not in virtue of some other feature it has, but has it primitively
- or, it has it in virtue of the contribution it makes to the truth conditions of the statements it occurs in.

I will not get into this debate here, but merely state that the second option seems to me to be the right one. My making the choice of taking sides with the second option isn't too important for the overall project here, but it allows us to say more on how this inferential role arises. In addition, it will allow us to say more about how it increases the overall expressive power of our language. I think inferential role is had in virtue of the contribution that the phrase makes to the truth conditions. Thus expressions like "and" make a certain contribution to the truth conditions, and because the contribution is what it is, it turns out that the phrase will have a certain inferential role.

But then we have to ask ourselves: what is the contribution to the truth conditions that the quantifier makes in its internal use? We have seen that we have a need for an expression that has a certain inferential role. We have to take recourse to it to communicate incomplete information in a language with the complexities that our natural languages have. And if you agree with me that inferential role is not primitive, but had in virtue of making a certain contribution to the truth conditions then we should attempt to say more about what the contribution to the truth conditions is that gives the expression the inferential role for which we need it. What contribution to the truth conditions would "something" have to make in order for it to have the inferential role that "F(t)" implies "Something is F" independently of whether or not t is something real?

If we want to know what the truth conditions of a statement are we have at least two ways to approach this. One is to give a correct but uninformative account of the truth condition by simply saying

(50) An utterance of ". . . ." is true just in case . . . .

Even though this is undoubtedly correct it doesn't help us for our purposes. We want an informative account of the truth conditions. We already know (50). What we would like to know is when precisely are the truth conditions fulfilled? And we would like to know this in an informative way, a way that spells this out and makes it more explicit than by simply giving the disquotational truth conditions. This can often be done very informatively when the sentence uttered involves complicated words, like

(51) Supervenience doesn't help to understand mental causation.

One can spell out the truth conditions of such statements by elaborating on the sentence, by basically saying something with the same truth conditions, but without the fancy words. In this case one would simply say more explicitly what supervenience is, and what mental causation is. But how can we do this for

(52) There is someone Fred admires.

There aren't really any words in this sentence that would need spelling out. How can we be more informative in giving the truth conditions than by just saying that it is true just in case there is someone Fred admires?

What we can do here, and maybe the only thing we can do here, is to make the contributions to the truth conditions that the quantifier makes in these particular uses explicit by taking recourse to a formal model. If we give a formal model of a fragment of English that contains "something" in its internal reading then we can learn more about what this expression does by seeing how we have to accommodate it in the formal model. And looking at the truth conditions of a formal language sentence that is the model of a natural language sentence can give us insights into the natural language sentence. In particular, it can give us insights into what contributions to the truth conditions the individual parts of the natural language sentence make. But even though this is a very promising strategy it immediately faces us with a problem. It is a substantial and difficult task to give a formal model of any non-trivial fragment of English. In particular, to give formal models of fragments of English that deal with such complicated things as intentional verbs and empty names. But these latter ones were exactly what was relevant to motivate the distinction between the internal and the external readings of quantifiers. How can we learn from formal models of fragments of English that contain quantifiers in their internal reading without first going into the substantial task of giving a formal model of a large fragment of English?

Even if we can't right here and now give a formal model of a fragment of English large enough to contain all the aspects that are relevant for the present discussion, we might still be able to do the following. Suppose we have a formal model of a fragment of English that doesn't contain quantifiers in their internal reading, how could we extend it to one that does? If we could see how we have to extend a formal model to include quantifiers in their internal reading then we might be able to learn from this what contributions to the truth conditions these quantifiers make and we might learn this in a very informative way. And this is exactly the strategy I would like to pursue in this section. To do this we have to say a little bit more about what a formal model of a fragment of a natural language is, what it means to extend such a formal model, and, in particular, what the conditions are for such an extension to correctly cover internal uses of quantifiers.

What a formal model of a natural language has to do for us is to be an assignment of formulas in a formal language to natural language expressions. But not every assignment will do. First of all, it will have to have to preserve something like the syntactic compositionality of syntactic expressions. If a natural language expression is built up from certain parts then the formal language expression that is assigned to it should be built up from the formulas that are assigned to these parts. This aspect won't be too important here, and I will not elaborate much on it. Secondly, not just any compositional assignment will do. We will need some further conditions on what a correct assignment is supposed to be. Intuitively, the assignment should be such that the truth conditions are correctly modeled. This isn't a very precise condition, but fortunately we don't have to make it completely precise. We can simply state the following necessary condition that a formal model has to satisfy. There might be other necessary conditions, too, but we don't have to worry about these now. The condition is that the formal model has to model inferences correctly. If certain sentences $S_1, \ldots, S_n$ imply a sentence S then the formal sentences assigned to $S_1, \ldots, S_n$ in the formal model also have to imply the sentence assigned to S, and the other way round. Note, though, that when we speak of implication among natural language sentences we use the pre-theoretical notion of consequence. When we speak of implication among the formal sentences we use the semantic notion of consequence as spelled out by the semantics of the formal language. To make this a bit clearer, let's introduce some of these notions more precisely.

> **Definition** Let **F** be a fragment of a natural language. A formal model **M** of **F** is a pair $\mathbf{M} = < \mathbf{L}, \varphi >$, whereby **L** is a formal language (with a semantics) and φ is a recursively defined function that maps expressions of **F** onto formulas of **L** such that sentences of **F** are mapped onto sentences of **L**.
>
> A formal model **M** of a fragment **F** of a natural language is called an **Inferentially Faithful Formal Model (IFFM)** iff
>
> $\sigma_1, \ldots \sigma_n$ implies σ_{n+1} iff $\varphi(\sigma_1), \ldots \varphi(\sigma_n)$ implies $\varphi(\sigma_{n+1})$.
>
> Implication among sentences of **F** is the informal notion of implication. Implication among sentences of **L** is the notion of implication that is made precise in the semantics of **L**.

With this definition we can formulate our task anew. Given an IFFM of a fragment of English that does not contain "something/someone" in its internal use, how can we extend it to an IFFM for this fragment together with this phrase in this use?

To be clear about what the conditions for the correctness of such an extension are we have to say something about what the inferential behavior of the new phrase "something/someone" is supposed to be. We have seen that in its internal use the quantifier has the following inferential role:

(53) F(n) implies something/someone is F.

We now have to assign to this quantifier some formal expression that exhibits this inferential behavior in the formal language. How this can be done should give us some insight into the natural language quantifier in its internal reading.

To do all this we have to do the following. Assuming that $\mathbf{F}$ is modeled by an IFFM $\mathbf{M} =< \mathbf{L}, \varphi >$ we have to find an extension of $\mathbf{M}$ to a model $\mathbf{M}'$ such that $\mathbf{M}'$ is an IFFM of $\mathbf{F}'$, which is $\mathbf{F}$ plus "something/someone" in its internal reading. To do this we have to extend $\mathbf{L}$ to a language $\mathbf{L}'$, and φ to a function φ' such that the resulting model $\mathbf{M}' =< \mathbf{L}', \varphi' >$ is an IFFM of $\mathbf{F}'$.

This can be done with only a few assumptions about the original model $\mathbf{M}$ that we have to extend. The following assumptions are sufficient to uniformly extend $\mathbf{M}$ to $\mathbf{M}'$:

1. Being a formula of L is defined inductively such that we have a set of terms $\mathbf{T}$ of $\mathbf{L}$, and on the basis of this the atomic formulas $\mathbf{A_L}$ of $\mathbf{L}$, and being a formula is defined as the inductive closure of $\mathbf{A_L}$ under operations $O_1, \ldots, O_n$.

2. The semantics of $\mathbf{L}$ is based on the notion of satisfaction of open formulas by, say, assignment functions.

We will assume furthermore that there is a subset $\mathbf{N}$ of $\mathbf{T}$ which are the names in $\mathbf{L}$. Assuming that $\mathbf{L}$ satisfies these conditions we now have to find $\mathbf{L}'$ and φ'. As $\mathbf{L}'$ we take the fragment of $L_{\omega_1,\omega}$ built up on the basis of $\mathbf{L}$, as determined by the following inductive definition.[3] Define $\mathbf{L}'$ as:

- $\mathbf{A_L} \subset \mathbf{L}'$

- if $\Psi(x_1, \ldots, x_m)$ is in $\mathbf{L}'$ then for each $1 \leq i \leq m$
  $\bigvee_{n \in N} \Psi(x_1, \ldots, x_m)[n/x_i]$ is in $\mathbf{L}'$

- $\mathbf{L}'$ is closed under $O_1, \ldots, O_n$

[3] $L_{\omega_1,\omega}$ is an infinitary language that allows conjunctions and disjunctions to be formed over countable sets of formulas. See (Keisler 1971) and (Barwise 1975) for more on infinitary logics. The infinitary language we encounter here is smaller than the smallest ones studied in either of these books. Usually $L_{\omega_1,\omega}$ is just first order logic with these countably infinite conjunctions and disjunctions. But of course, one can take (basically) any language as a base language and then build a larger language by allowing countably infinite disjunctions and conjunctions as formulas, starting with the formulas of the base language. We will see more on all this in appendix A.

Here $\bigvee_{n \in N} \Psi(x_1, \ldots, x_m)[n/x_i]$ is the infinite disjunction over all results of substituting a name of **L** for the variable x_i in Ψ. Since **L** and **L**$'$ have the same atomic formulas the names of **L** and of **L**$'$ are the same. **L**$'$ is thus the fragment of $L_{\omega_1,\omega}$ that contains all and only the infinitary formulas that are required according to the above inductive definition.

We can extend φ, which was assumed to be recursively defined, to φ' as follows:

- φ' is defined just like φ except for the additional clause:
- φ'("Ψ(something/someone)") = Con('$\bigvee_{n \in N}$', Con(φ'("Ψ''), '$[n/x_i]$')).

Here Con is the concatenation function and double quotes are quasi-quotes.

The semantics of **L**$'$ is now a standard extension of the semantics of **L**, exploiting that the above assumption about the semantics of **L**, namely that is based on the notion of satisfaction. We simply extend the definition of when an assignment function satisfies a formula with the following clause:

- A disjunction is satisfied by an assignment function iff one of its disjuncts is.

It now only remains to be seen whether or not the extended model $\mathbf{M}' =< \mathbf{L}', \varphi' >$ is a IFFM of **F**$'$. This, however, is straightforward. We just have to see whether or not the inferences come out valid in this model that we observed sentences that contain "something/someone" in its internal reading figure in. These were the inferences of the form

(53) F(n) implies something/someone is F.

These inferences are supposed to hold whether or not n succeeds in referring. And this is true in our model. Since "something/someone is F" is modeled as a disjunction that contains "F(n)" as one of its disjuncts such an inference will always be valid, whether or not n exists. We can freely assume that the language used to model a fragment of English with empty names is a free logic. This is fully compatible with the way we defined our extension above.

"something/someone" in its internal reading can be nicely modeled by taking recourse to small infinitary extensions of the language that models the rest of a fragment of English. The inferential behavior of sentences containing this phrase corresponds exactly to the inferential behavior of certain simple infinitary sentences in an infinitary extension of the formal language that models the rest of a fragment of English.

Taking recourse to a small fragment of infinitary logic in modeling the truth conditions and inferential behavior of a quantifier in its internal reading is not the only way to do this. However, it is in a certain sense a distinguished way. It is, in a sense, the optimal solution. The infinitary truth conditions spelled out above do have the required inferential role, and anything else that has this role is implied by it. If an expression Ψ has the inferential role in (53) then it is either equivalent to the disjunction over all F(n), or implied by this disjunction.

So we have seen that by modeling the truth condition of the internal quantifier by certain infinitary expressions we can make explicit how it gets the inferential role it has. Taking recourse to the infinitary expressions is most useful in spelling out how the internal quantifier increases the expressive power of the language with an internal quantifier in it. But of course, the claim I am making here isn't that natural language contain infinitely long expressions, or anything like that. All I am claiming is that by taking recourse to certain formal languages we can spell out the truth conditions of certain natural language expressions. The formal languages are only tools taken recourse to to make the truth conditions of natural expressions more explicit. Speakers of the language of course don't have to have access to this. Just like ordinary speakers don't have to have access to Montague's type hierarchy for Montague semantics to be correct.

To sum up we can say the following. Ordinary quantifiers like "something" are semantically underspecified with respect to whether or not they impose domain conditions, or whether or not they have a certain inferential role. In languages like ours these two do not coincide with respect to truth conditions, but we nonetheless have a need for both of them. Whether or not a concrete occurrence of "something" is one where it is used in its inferential role reading or in its domain conditions reading is dependent on subtle features of the context of the utterance. This is just as in the case of whether or not s plural phrase is used in its collective or its distributive reading. There is no easy test for this. It depends on the details.

2.3.7 Substitutional quantification

What I said about quantification above has certain similarities and important differences with substitutional quantification. This section will spell out what the similarities and differences are.

Substitutional quantification is based on an alternative way to give a semantics for first order quantification. The standard way of doing this, also called objectual quantification, goes as follows:[4]

(54) "$\exists x \Phi(x)$" is true in model **M** iff there is an object o in the domain of **M** such that o satisfies Φ

An alternative way of giving a definition of when such a first order quantified statement is true in a model can be given by taking recourse to when substitution instances are true, rather then when an open formula is satisfied by an object from the domain of a model. By a substitution instance I mean the result of substituting the variable bound by the quantifier by some term in what is called the substitution class. For simplicity reasons, lets take the substitution class to be all the names (constants) in the formal language we are talking about. $\Phi[c/x]$ is the result of substituting c for x in formula Φ. The alternative definition is:

(55) "$\exists x \Phi(x)$" is true in model**M** iff there is a substitution instance $\Phi[c/x]$ which is true in **M**.

It was a debate in the philosophy of logic some time ago which one of these two should be considered the correct semantics for first order quantification. And it was a related debate what role quantification should play in determining ontological commitment.[5]

This debate strikes me as quite confused. The semantics for first order quantification does not have any direct implications for ontological commitment. The relevant question isn't how we should give the semantics for first order logic, but what we are doing when we utter ordinary English sentences that involve quantifiers. Why would it matter which one the correct semantics is for a formal language? What would it even mean for a formal language to have the correct semantics? This might make more sense if one takes it for granted

[4] I simplify and won't consider the case where Φ contains free variables. This is taken care of as usual by taking recourse to formulas being satisfied by sequences of object of the domain.

[5] See (Marcus 1993).

that the ordinary English quantifier "something" should be understood as, or modeled as, a first order quantifier. Then one might think that the semantics for the first order quantifier has to be such that it gives the right results about the ordinary English quantifier. One can then ask whether or not the ordinary English quantifier behaves like a substitutional quantifier, or like an objectual quantifier. And in this debate either side seems to have good reason to reject the other one. The substitutional quantifiers side had good cases coming from quantifying into intensional contexts, the objectual quantifiers side had good cases involving quantification over uncountable domains, over things with no name, and the like. But really, I think, this debate is based on false options. The debate assumed that either the ordinary English quantifier is always like a substitutional one, or always like an objectual one. And, of course, if one thinks that natural language quantification is like first order quantification than that is not an unreasonable assumption. But once we notice that natural languages involve issues like semantic underspecification, subtle contributions of context to content, and the like, things look differently. Still, though, it is important not to make the mistake to think that given this we can just pick on the individual occasions of an utterance how a quantifier should be understood, and pick so in accordance with our ontological persuasions. It is not up to us, as the people doing the theory about quantification to pick what is the right interpretation of an occurrence of a quantifier. It is a fact of the matter about the utterance how it has to be correctly interpreted. If we conclude that quantifiers have to be interpreted differently on different occasions, as I have argued we should, then this should not come from ontological persuasions, but from a general story about the communicative needs that speakers have in ordinary situations of communication. And this, I have also argued, is the case.

Believers in the view that quantification in English should be understood as substitutional quantification thought that this view would have implications for ontology and what ontological commitment comes down to. From the fact that some theory implies "There is an F" one couldn't directly conclude that the theory has F's among its ontological commitments, or so the story goes. But that seems hard to believe. How could an ordinary utterance of

(56) There is a cow in the back yard.

not commit us to the existence of a cow? Whatever the semantics of the formal language that is associated with English sentences is, it can't have anything to do with that fact that an ordinary utterance of (56) commits me to the existence of a cow. Without a doubt, I think, some uses of quantifier directly carry ontological commitment. The question can only be whether or not all do. (Gottlieb 1980) attempts to deal with something like this and proposes as a criterion for ontological commitment, in effect, that a quantified statement only carries ontological commitment if the quantifier in it has to be interpreted as an objectual quantifier. "has to" here has to be understood such that an interpretation of the quantifier as a substitutional quantifier would give us the wrong truth value because, say, the domain of quantification contains things without a name, as when we quantify over the real numbers.[6] This approach, however, gives the wrong results. It makes it too difficult for us to commit ourselves to things that we want to commit ourselves to. How could I, for example, commit myself to the existence of my mother given that my mother has a name? Any quantifier over my mother can be interpreted as a substitutional quantifier without any problems, since my mother has name. The question isn't whether or not a quantifier can be interpreted this way or that way, but how it should be interpreted. If one thinks that quantifiers are always used the same way then one can only do what Gottlieb does and try to build substitutional quantification somehow into an account of ontological commitment. But we have seen that ordinary English quantifiers really have two functions in communication. When we spell them out and see what contribution they make to the truth conditions in each of their functions we can see that in one of them it is a lot like a substitutional quantifier (I will elaborate on this in a second). But the important consideration is to show what these different functions are and why we need both of them. It isn't what semantic treatments one can give of quantification in formal languages.

Besides this, substitutional quantification has a close relation to the above internal quantification. The formal language that I took recourse to in giving a model of the truth conditions of the quantifier in its internal reading can be seen as a notational variant of a formal language with a substitutional quantifier in it. Substitutional quantifiers can be seen as devices for forming infinite disjunctions and conjunctions, and can be understood as giving a finite notation for certain simple infinite disjunctions and conjunctions. Understood thus, formal languages with a substitutional quantifier in them do provide the appropriate tools for modeling the truth conditions of quantifiers in their internal reading. But partly because lots of the discussion about substitutional quantification was in the ballpark of a

[6] See (Gottlieb 1980, 52).

discussion which one was the correct semantics for quantifiers, the substitutional or the objectual, I prefer to use the terminology of internal vs. external readings of quantifiers.[7]

2.3.8 An objection

Before we go on I would like to dispel an objection that I have encountered several times in personal communication and that must go back at least to the discussion of Wittgenstein's analysis of quantification as the conjunction or disjunction of its instances.[8] The objection is as follows:

(**O**) It can't be right that

> (57) Everything is F.
>
> (in any of its uses) is equivalent to an infinite conjunction
>
> (58) F(a) and F(b) and ...
>
> (57) implies (58), to be sure, but (57) in addition says that there is nothing that isn't F. (58) only says that a, and b, and ... are F. It doesn't say that nothing is not F. So, quantifiers can't be analyzed in terms of conjunctions and disjunctions.

This objection has some prima facie plausibility, but it is mistaken as an objection to what we are doing here. What we are trying to do here is to give an account of the truth conditions of utterances of sentences with quantifiers in them. Thus we have to find out under what circumstances a particular utterance is true. Now, *given* that the utterance of (57) was made in particular circumstances, or in a particular world, it will be true just in case each individual thing in that world is F. If a, b, ... are all the individuals there are then the utterance will be true just in case F(a) and F(b) and ... holds. That's exactly what is required for the utterance of (57) to be true in the situation where it was uttered.

To be sure, a quantified sentence can't be truth conditionally equivalent to an infinite disjunction or conjunction independently of the context of its utterance. At the level of sentence, not utterance, the truth conditions can be stated as something like: an utterance of this sentence in a world w is true just in case all the things in w are F. At this level no

[7] Also because this is Carnap's terminology. More on how the present account of quantification relates to Carnap's internal-external distinction later.

[8] In his (Wittgenstein 1984).

association with an infinite conjunction is possible. But given that the sentence was uttered in w, the utterance will be true just in case every object in w is F. And this is so just in case F(a) and F(b) and ... is true.[9]

2.3.9 Other quantifiers

So far I have focused on "something". This was partly for the reason that this quantifier is the one that is used in the trivial arguments that we have to try to understand. And it is partly because this quantifier best allows for the motivation that we have at least two uses for quantifiers. However, what I have said about "something" isn't restricted just to it. Quantifiers are connected to each other very closely, and what is true of one of them will get carried over to the others. For example, "something" is connected to "everything" in the well known way that

(59) Something is F.

is true if and only if

(60) Not everything is not F.

If "something" can make two different contributions to the truth conditions of an utterance then the above connection will motivate that "everything" can be used in two different ways, too. In this case, too, we will have two uses for this quantifier, one that comes from imposing domain conditions, the other that comes from the need for a certain inferential role. In the case of "everything" the domain condition is, of course, that every entity that is part of reality has a certain property. The inferential role is that from

(61) Everything is F.

it follows that

(62) t is F.

whether or not "t" stands for something real.

In general, internal quantifiers will give rise to an increased expressive power, whether or not they are closely associated with a certain inferential role. The quantifiers we have focused on do have such an inferential role, but this is not what is really essential. The increased expressive power is what is essential.

[9] For a more general and complete story on truth conditions given X, see (Perry 1997b), and (Perry 1997a).

2.3.10 Preliminary conclusion

Let's sum up what we have seen in this chapter. I have argued that we do have a need in ordinary communication for a phrase that imposes certain domain conditions, and for a phrase that plays a certain inferential role. The former simply to say what things the world is made of, the latter to communicate incomplete information, amongst possibly many other uses. In natural languages like our's these two semantic roles can't be played by one and the same phrase contributing in the same way to the truth conditions in each occasion. And I have argued that it is in fact played by a phrase that is semantically underspecified and that on some occasions plays one role and on others plays the other. In particular, it makes different contributions to the truth conditions of an utterance in which it occurs. In addition, I have tried to say a bit more about what the contribution to the truth conditions should be taken to be.

How a concrete utterance of a quantified statement should be understood is thus similar to how an utterance of a sentence containing a semantically underspecified phrase should be understood. It will be a subtle contextual contribution to content what the truth conditions of the utterance are in this particular occurrence. In particular, just as with plural, we should not expect for there to be a simple test or procedure to decide which one is the correct reading on an occasion. To decide what the correct reading is we will have to look at details of the context of the utterance. We will have to look at what the speaker is trying to do on this particular occasion. This will help to see how the underspecified part of the sentence uttered get's fully specified, if it get's specified at all.

In addition, we saw that the internal reading of the quantifier does not have a direct connection to ontology in the following sense. If the quantifier in

(63) Something is F.

is used internally this by itself does not settle the question whether or not there are Fs in the domain of discourse. This question would be settled if the quantifier would have been used externally (i.e. if you sincerely utter (63) and the quantifier in it is used externally in this utterance then your natural ontology will contain Fs. Your utterance will only be literally true if reality contains Fs.) But if the quantifier is used internally then this conclusion would be premature. All that we can conclude from that is that one of the disjuncts of the form 'F(t)' is true. It is then a further issue whether or not there also is a disjunct 'F(t)' such that it is true *and 't' stands for something real*. It might be that the only true disjuncts are ones that involve empty names or non-denoting terms. In this case we shouldn't draw any ontological conclusions from the literal truth of (63). So, if (63) is literally true and the quantifier is used internally then the ontological question, which is formulated using a quantifier externally, since it is a question about what things are out there in reality, like

(64) Are there Fs?

isn't thereby answered. More work has to be done to answer it.

2.4 Meinongean, Fregean and pretense theoretical alternatives

The discussion so far has touched on a number of issues that have been discussed elsewhere in quite some detail, and so far I haven't really said anything about why I think other approaches that try to deal with similar problems are mistaken. In this section I will discuss some of these approaches. The alternative approaches about dealing with quantified statements like the above, where we apparently get a separation of quantification and ontology, can be put in four different, but not necessarily exclusive, categories:

1. **Quineanism.** The above quantified examples have to be rejected since they can't be literally true. That is so simply because there is no Sherlock Holmes in the domain of quantification, since no such thing exists. However, that is not a bad thing, since these statements do not play a central or important role.

2. **Meinongianism.** The above quantified statements are to be accepted and have to be taken as literally true. However, they are not independent of ontology, only independent of what exists. Thus in particular, and surprisingly, reality contains such an object as Sherlock. He is a non-existent object, but nonetheless he is part of reality.

3. **Fregeanism.** Non-denoting terms like "Sherlock Holmes" can occur in true sentences only in intensional contexts, in which they refer to their sense. Quantified statements have to be explained with this in mind. Either they range over senses, or they themselves occur only in intensional contexts, or they are to be rejected.

4. **Pretense theory.** The above statements are not literally true, but true in a game of pretense in which we pretend, for the moment, that there are such things as Sherlock Holmes.

As I said, these categories are not exclusive. Quineanism might be combined with Fregeanism to have a position that allows for some quantifiers to (apparently) range over, say, the Fountain of Youth, but only in intensional contexts. (We will look at this in more details in section 2.4.2.)

Proponents of these views might and will claim that I have misanalyzed the situation and ignored lots of other work that has been done to deal with similar concerns. I will in the following discuss these alternatives. I'd like to point out, though, that I don't think that any one of them can be refuted in the sense that one can show they are inconsistent. I doubt this. But I do believe that they are, as it turns out, false. I will thus not proceed to show that Meinongianism, say, is inconsistent, and the above proposal is the only other option. I will rather point out why I think Meinongianism is mistaken as a theory about us, our language, and what we do when we talk in a certain way. Overall, the above proposal is the most plausible one to deal with these cases.

2.4.1 Quantification and non-existent objects

Meinongianism is the position that reality consists of more than the things that exist. It also consists of things that don't exist. In particular, Meinongians deny that from the literal truth of a quantified statement one can draw any conclusions about what exists. However, one still can draw conclusions about ontology. It does remain a further question, though, whether or not the thing that one has concluded is out there is an existing thing or a non-existing thing.

There is a bit of a puzzle about non-existent objects that is quite similar to the puzzles we started out with in chapter 1. There is a rather innocent way to believe that there are non-existent objects, and that there is a rather substantial view. According to the innocent view, to belief that there are non-existent objects doesn't come down to much more than believing that, say, Santa doesn't exist. This is so because to say that Santa is a non-existent object isn't to say anything substantially different than to say that Santa doesn't exist. Consider the following ways of basically saying the same thing:

1. - Fred doesn't cooperate.
   - Fred is non-cooperative.
   - Fred is a non-cooperative person (or thing, or object).
2. - Santa doesn't exist.
   - Santa is non-existent.
   - Santa is a non-existent person (or thing, or object).

It seems that to accept that there are non-existent objects in this sense is not the acceptance of anything philosophically substantial, and comes down to no more than accepting that, say, Santa doesn't exist, just like the acceptance of there being non-cooperative people comes down to nothing more substantial than the acceptance that, say, Fred doesn't cooperate.

However, there is also a more substantial way of accepting that there are non-existent objects. According to this way reality is made up of two kinds of things: the existing and the non-existing. Sure enough, we are only in causal contact with the former, but still, our domain of quantification contains both of them. Furthermore, the domain of non-existent objects has to be taken to be mind and language independent. Since, after all, how could they be mind and language dependent? How could something non-existent be created by or be dependent on language use? And since they are mind and language independent, the particular non-existent objects that we pick out with some of our terms aren't in any way distinguished. Any terms we would have used would have picked out non-existent objects, too. Thus there has to be a plenitude of non-existent objects. Thus the world is much more populated than we thought. And this is certainly a substantial way to accept that there are non-existent objects.

But how do those two ways of accepting that there are non-existent objects go together? Could it be that accepting in the apparently innocent way that there are non-existent objects is really accepting a substantial metaphysical thesis. It might be said that the above alone is not enough to accept that there are non-existent objects, since it might be one thing to accept that Santa is a non-existent object, but quite another that there are non-existent objects. It might be said that ontological commitment is carried by quantifiers, not by terms. Sure enough, but it still seems that there is a rather innocent way to accept quantified statements that apparently range over non-existing objects. Consider, again, the analogy between the following:

1.
   - Fred doesn't cooperate.
   - There is someone who doesn't cooperate, namely Fred.
   - There is someone who is non-cooperative, namely Fred.
   - There is a non-cooperative person (or thing, or object), namely Fred.
2.
   - Santa doesn't exist.
   - There is something that doesn't exist, namely Santa.
   - There is something that is non-existent, namely Santa.
   - There is a non-existent person (or thing, or object), namely Santa.

Similarly, when I say:

(65) There are three people that I admire, Sherlock, Homer and Clinton.

it seems that I do not say anything substantially more than

(66) I admire Sherlock, Homer and Clinton.

which isn't to say anything over and above that

(67) I admire Sherlock, I admire Homer, and I admire Clinton.

How come that these apparently innocent ways of speaking are supposed to lead to the acceptance of a substantial metaphysical picture?

One way to account for this is to say that such apparently innocent statements as

(68) I admire Sherlock.

directly imply that Sherlock is part of reality, since admiring is a relation between two things. Maybe (68) has ontological presuppositions about Sherlock besides me for its objective truth. But that doesn't seem to be right. It seems that what makes it true that I admire Sherlock is something about me, not something about me and some other thing. It seems that my psychological states alone should be enough to make it true. There really being a Sherlock doesn't seem to be required. At least not on the face of it. But what does seem to make it explicit that the reality of Holmes is required for the truth of (68) is the validity of certain quantifier inferences. Or at least so a standard line of reasoning goes. To see how our talk about what doesn't (seem to) exist functions we will have to look more closely at uses of quantifiers in the occasions where we appear to be quantifying over non-existent objects, and in particular how such quantification relates to quantification on other occasions.

Quinean and non-Quinean quantifiers

The standard debate about non-existent objects is roughly the following: first and foremost there is a debate whether or not we should or have to accept statements with quantifiers that apparently range over non-existent objects. It is commonly assumed that if we do accept such quantification than we have to accept a substantial metaphysical picture. It is partly this that leads people to reject such quantifiers, and insist that sentence in which they occur are not literally and strictly speaking true. Sure enough, there might be some other use to them, but literal and objective truth can't be given to them unless we accept the substantial metaphysical picture.

To evaluate whether or not acceptance of such quantification indeed leads to the acceptance of the substantial metaphysical view let's see how this transition is usually made. First, an observation that leads to some useful terminology. Those who accept quantification over non-existent objects basically accept that quantifiers come in two kinds. There are ordinary quantifiers, like

(69) Someone ate my sandwich.

and there are those that apparently require an ontology of non-existent objects, as

(70) There is something that Ponce de Leon is searching for, but will never find, since it doesn't exist.

The first kind can be explicitly modified with "which/who exists" without change of truth conditions. Only things that exist are relevant for the truth of the sentence. After all, to say:

(71) Someone ate my sandwich, but he doesn't exist.

is more than odd. When quantifiers are used in the way in which they apparently range over the Fountain of Youth and the like, such an explicit modification without change of truth conditions does not seem possible. (70) modified this way seems non-sensical.

Let's call the occurrence of a quantifier in an utterance that is such that we can explicitly modified with "which/who exists" without change of truth conditions of that utterance a **Quinean quantifier**. Let's call those occurrences where such a modification is not possible without change of truth conditions a **non-Quinean quantifier**. I'd like to stress here that the distinction between Quinean and non-Quinean quantifiers is one that applies at the level of individual utterances of quantifiers. It is not a distinction at the level of language. So, even if this distinction is legitimate and not empty, it is a further question whether or not our language has two kinds of quantifiers in it, or whether or not to account for the difference between the two kinds of occurrences of quantifiers in another way.

We can now distinguish two core issues in the debate about non-existent objects:

a) Are there any legitimate uses of non-Quinean quantifiers? Is there a need for us to take recourse to them when we try to make an objectively true statement?

b) If yes, how should we understand the function of the non-Quinean quantifiers in these uses?

Quineans and Meinongeans disagree about whether or not there are legitimate uses of non-Quinean quantifiers, but by now we have seen that Quineans are wrong.[10] There are such uses, and they play an important role in communication. But how about b)? Meinongeans would disagree with the account given above. But they would agree that there is an important role that context has in determining the truth conditions of a sentence with a quantifier in it. However, their understanding of the role of context is very much within the extended simple picture of the role of context, as spelled out above. The Meinongeans reasoning about how to understand non-Quinean quantifiers is as follows:[11]

- An expression like "something" can be used both as a Quinean and as a non-Quinean quantifier, as in

  (72) Something is eating my cheese, probably a mouse.

  and in

  (73) Something is keeping me awake at night, namely the monster I dream about.

[10] Unless they combine their Quineanism with a form of Fregeanism. This is still open for now, and will be discussed in section 2.4.2.

[11] See, for example, (Parsons 1980).

- Thus there has to be some difference in the particular occasions of the utterance which makes it that one can be modified with "which exists" without change of truth conditions, whereas the other one can't.

- The way to understand this is simply the following: Quinean quantifiers are a case of a well known way in which the context of utterance of a sentence with a quantifier contributes to the truth conditions, or what is said with the utterance, namely contextual restriction of the quantifier. Quinean quantifiers are implicitly restricted to what exists. Non-Quinean quantifiers don't have this restriction. This phenomenon is just like contextual restrictions of quantifiers in standard examples of utterances like:

  (74) Everyone has to die.

  (75) Everyone is hungry, let's take a break.

  The second occurrence is contextually restricted to the group of people in the room of the utterance (or some group like that), whereas the first one doesn't have such a restriction.

- Since non-Quinean quantifiers don't have such a restriction this shows that the domain of quantification really is what the non-Quinean quantifiers range over. Quinean quantifiers range over a subdomain of this domain, namely over all those things that exist. Thus the true domain of discourse contains non-existent objects, and thus the substantial picture of reality is justified.

And once one believes that it is no big step, and in fact apparently inevitable, to believing that there is a plenitude of non-existent objects that are part of reality. Since how could it be that our thoughts and our speech has anything to do with what non-existent objects there are? It is reasonable to assume that what we do has an effect on what exists, but how could it have an effect on what doesn't exist? Thus is seems that if one believes in the restricted quantification view of non-Quinean quantifiers then one is bound to believe that there is a plenitude of non-existent objects out there in reality waiting to be talked about.

This view strikes many people to be absurd, and I have my sympathies with their feelings. But being absurd is not being inconsistent. Meinongianism might be perfectly consistent. It just would be so miraculous if anything that we can possibly conceive and introduce a term for is already out there in reality, independently of our conceiving or introducing a term for, waiting in the realm of non-existent objects to be talked about. This commonly shared attitude has led to the following reaction to the considerations in this section. The reaction is to concede that if there are true utterances with non-Quinean quantifiers in them then the substantial picture of non-existent objects has to be correct. This persuasion has widely led to an insistence that non-Quinean quantifiers can never occur in literally true utterances, even though there seem to be a number of good cases of this.

I would rather like to look at things differently here. We saw above that we have a need in communication to take recourse to non-Quinean quantifiers. And we saw what one of these needs is. What I think is mistaken is rather the Meinongians reasoning from the acceptance of non-Quinean quantifiers to the acceptance of an ontology of non-existent objects. It is a central part of their reasoning that the difference between the two occurrence of quantifiers is a difference of contextually restricted quantification. And that is a reasonable assumption given something like the extended simple picture of the role of context in determining content (see section 2.2.1). We saw that the difference between Quinean and non-Quinean quantifiers is one of a difference in occurrences of the quantifier. And the only way in the extended simple picture how one occurrence can be different from another occurrence of the same quantifier phrase is by contextual restriction. But without the extended simple picture we have no reason (except maybe one to be discussed shortly) to assume that the difference between Quinean and non-Quinean quantifiers is one of contextual restriction. Once we realize that the extended simple picture of the role of context is too simple we can more freely look at what the contextual difference between these two occurrences of quantifiers really is. I have argued above that what really matters (at least in the cases we have looked at) is for the quantifier to have a certain inferential role. And when we look at the examples that are supposed to motivate Meinongianism we can see, I think, that this is what these examples also point to. Meinongians usually don't stress the inferential role of non-Quinean quantifiers as their important aspect, but what Meinongians in fact do with their account of non-Quinean quantifiers is to try to give them this inferential role via their domain conditions. As we have seen above, a quantifier in its external reading would have the inferential role that we want the quantifier in its internal reading for if certain conditions about our language would hold. What Meinongians, in effect, claim is that these conditions really do hold in our language, and that thus the inferential role of the quantifier in its internal reading is in fact had by the quantifier in its external reading.

They maintain that natural languages are like certain formal languages in the sense that there really are no terms or names that stand for nothing in the domain of discourse. Every term and name does stand for some thing in reality, most of which don't exist, or so their story goes. But once we notice certain more large scale features of natural languages, like semantic underspecification, and once we consider initial implausibility of an ontology of non-existent objects and initial plausibility of creatures like us will having a language that exhibits partiality (as spelled out in section 2.3.3) then it is fair to say that the semantic underspecification account is to be preferred over the contextual restriction account and an ontology of non-existent objects. But, as I said, both views might be perfectly consistent, and in the end it will come down to an empirical claim about what we are doing when we speak a certain way. I know where I would put my money if I had to bet on one side or the other.

The role of "exists"

There is, however, one consideration that might be taken to speak in favor of Meinongianism. It is the role of the word "exists". We were able to distinguish Quinean and non-Quinean quantifiers by saying that the Quinean's are the ones that can be explicitly modified with "which exists" without change of truth conditions. The Meinongians can understand this as making a contextual restriction of a quantifier explicit. They can say that this is just like if I would say

(76) Everyone (who is in this room) is hungry. Let's take a break.

Here "who is in this room" makes a contextual restriction of the quantifier "everyone" explicit. It is not necessary to make such restrictions explicit. But if a quantifier is contextually restricted then we can make the restriction explicit without change of truth conditions. And this, the Meinongians can say, is what is going on when we make the difference between Quinean and non-Quinean quantifiers explicit.

However, the believer in semantic underspecification has a nice story about this, too. It is based on noticing that there is an interesting parallel between the above situation and other cases of semantic underspecification. Consider plural, again. An utterance of

(77) Four philosophers wrote a book.

can be uttered in such a way that the plural phrase is either used collectively, or distributively. However, one can add more words to this sentence such that it then allows only one of these readings, but that also doesn't change the truth conditions of the sentence in that reading. For example, the sentence

(78) Four philosophers together wrote a book.

can only be uttered to have a collective reading of the plural, and the sentence

(79) Four philosophers each wrote a book.

can only be uttered with the plural phrase having a distributive reading. In other words, there are words in our language such that if we expand a semantically underspecified phrase with these words then we force a certain reading of this phrase (or we complete the underspecification). So, with more words we can force a certain reading, or fully specify what was left underspecified. And that this is so should not be surprising. After all, we don't always want to rely on the subtle features of context to determine what was left underdetermined. Sometimes we want to determine it without a role for context to play. For example, when we try to make explicit what we or someone else said, in particular how it is supposed to be understood in detail. And the same happens, I think, when we use the word "exists" right after a quantifier. The difference between "something" and "something, which exists" is that the latter forces that the quantifier is used in its external use. It shouldn't be understood as an explicit restriction of the scope of the quantifier, in the sense in which the Meinongians would want it.

2.4.2 Quantification and intentional verbs

One might think that the examples of internal uses of quantifiers that we saw above essentially involve intentional verbs, and that I made a mistake in how I used the interaction between intentional verbs and quantifiers. This line of reasoning that I have in mind here involves a Fregean attitude towards the role of names and the like in intensional contexts. According to Frege a name or other term does not refer to or stand for its referent in an intentional context, but rather refers to some other thing, its sense. Senses are supposed to be necessarily existing entities, and it is because of empty names standing for their sense in intentional context that such sentences as

(80) Fred admires Sherlock.

can be uttered truly. A Fregean, as I imagine them, also believes that there is a very important difference about scope when it comes to the interaction between quantifiers and intentional verbs. Thus, according to the Fregean, and not necessarily only the Fregean, there is an important difference between

(81) There is something Ponce de Leon is searching for.

and

(82) Ponce de Leon is searching for something.

The former is false, the Fregean says, because there really is no thing he is searching for, but the second is true since he is searching for the Fountain of Youth. This difference is analogous to a *de re* and *de dicto* difference. In the first case we say that there is a thing such that Ponce is searching for it, and in the second case we only say that he is searching for something. A Fregean may argue that I did not pay sufficient attention to this difference in my above discussion.

A Fregean, for our purposes here, will believe the following

- All quantifiers that range apparently over things that aren't real are either illegitimate (i.e. the statements in which they occur are literally false), or they occur within the scope of an intentional verb.
- There is an important difference between sentences like (81) and (82). The difference is a difference of quantifier scope, analogous to a *de re* and *de dicto* difference.

These are important issues that I have postponed so far. We have to address them now. I will argue that both of these points are in fact false. First, the need we have for internal quantifiers is not dependent on their occurrence with intentional verbs, and secondly, that the difference between the *there-is* construction and its counterpart is not at all a difference about scope analogous to a *de re* and *de dicto* distinction. It is rather something quite different. I'll elaborate on this in this order.

Intentional verbs

All that was required for the examples that motivated the need for an internal reading of quantifiers was that there were true sentences that contained empty names or non-denoting descriptions. The Fregean can concede this, but require that all such examples will involve empty names or non-denoting descriptions in intentional contexts. However, this is false. There a a variety of examples that work just as well as the one involving "admire" that I used above to make the case I was making there. If this is correct then the examples are not essentially connected to intentional verbs, and the Fregean strategy of dealing with them is a red herring. Thus there are examples that do not involve intentional verbs but where we also learn valuable information that we then can partially forget, leaving us in the same situation as we were left above using the example involving "admire". Let me give some such examples.

- **Negative existentials.** One classic example of true sentences involving empty names are negative existentials, sentence with which one claims that such a name is empty, that is, stands for nothing. For example

  (83) Sherlock Holmes doesn't exist.

  (84) Sherlock Holmes was only made up by Conan Doyle and doesn't really exist, even though it would be nice if he did exist.

  Knowing this, again, might be very useful information. It might refute a conspiracy theory according to which Sherlock Holmes shot JFK.

- **Comparison.** Even though Sherlock doesn't exist, we can compare him to real people, and we can say a number of things that are true with such comparisons. Consider:

  (85) Sherlock is smarter than Clinton.

  (86) Sherlock is more famous than Madonna.

  Note that these examples are not true within the Sherlock fiction. After all, Madonna isn't part of the Sherlock fiction.

- **Intentional adjectives and adverbs.** This category might seem to be not much different than intentional verbs, but without going into the details, it isn't clear what a Fregean should do with them, and how they might fit into a Fregean theory:

  (87) James Bond is supposedly irresistible.

  (88) Santa allegedly lives on the North Pole.

Of course, I have not offered an account of what is going on in these examples, I am merely mentioning them to diffuse the impression that everything hangs on the use of intentional verbs, and that the correct way to go is using a Fregean strategy with different referents in different contexts, rather than two readings of quantifiers. My approach is also supported by the following consideration.

There-is constructions and scope

I take it that only radicals would deny that Ponce de Leon is searching for something when he is searching for the Fountain of Youth. After all, whether or not you are searching for something depends only on you and what you do. It doesn't depend on whether or not you will or can succeed in what you are doing. In a word, you can search for something that doesn't exist. And it sure seems that Ponce would truly say of himself that he is looking for something if we would ask him what he was doing in the swamps of Florida on such a hot day. However, it is usually taken to be quite controversial whether or not it can be truly said that there is something Ponce de Leon is searching for. In short, the Fregean believes that there is a substantial difference between

(81) There is something Ponce de Leon is searching for.

and

(82) Ponce de Leon is searching for something.

In the first case Ponce has to look for something real, in the second he doesn't.

I doubt that this is correct. I rather think that the following is correct:

- Both (81) and (82) can be uttered such that for the utterance to be true Ponce either has or doesn't have to search for something real.

- The difference between (81) and (82) is one of presentational force or topicalization rather than of semantically relevant scope.

Let me explain.

The pair (82) and (81) is a special case of a pair of sentences that arise from the so-called existential construction, or there-is construction.[12] With this construction we take take a sentence, like

(89) A mouse ate my cheese.

and say apparently the same thing with another, a bit more complicated sentence:

(90) There is a mouse who ate my cheese.

And it seems that there is a subtle difference between these two. In (90) you make it clear and bring it out that it is the mouse that you now want to talk about, whereas in (89) it might be the mouse or the cheese. With (90) the mouse is given a bit more special place than with (89). This difference comes out nicely, I thing, in the examples we are most concerned with here. When we think about what the difference is between (81) and (82) we realize that they would lead to different ways how a conversation would be continued. For example, when I start with (82) I would usually continue talking about Ponce, as in:

(91) Ponce de Leon is searching for something, but he has failed to find anything of interest. If he doesn't have more success soon he will be in big financial trouble. I wonder why he continues on this mission for so long...

On the contrary, (81) will usually give rise to talk about whatever Ponce is looking for. So, it might continue as follows:

(92) There is something Ponce de Leon is searching for, and it must be really hard to find and of great value since he and many others have tried for years to find it. I wonder what it is.

If Ponce is supposed to be the topic[13] of the conversation then I am well advised to utter (82). However, if what he is looking for is supposed to be the topic of the conversation

[12] See (Kim 1996) for a whole dissertation on it.

[13] I use the term "topicalization" rather broadly. The topic is what intuitively is the main thing I talk about.

then I am better of uttering (81). Different syntactic construction can lead to different topicalizations, but not necessarily with a change of truth conditions.

I'd also like to point out that "there is" isn't the quantifier in a sentence like (90). "a mouse" is the quantifier. In particular, there is no difference in which and how many quantifiers occur in pairs like (90) and (89), or (81) and (82). That "there is" is a quantifier is something that I have heard in philosophical conversations several times, but is clearly a mistake. This is particularly relevant to understanding the overall view I am defending. The claim there is not that there are two quantifiers, "there is" and "there exists". Neither one of them is a quantifier. However, they together with a quantifier in the rest of the sentence can give rise to a difference related to quantification. In particular, after "there exists" the quantifier will be forced to be used in its external reading.

In saying that there is a difference of something like topicalization here I do not mean to deny that there can't a difference in scope also in these examples. However, I would like to claim that the scope difference isn't that (81) always is used such that the quantifier has wide scope, and (82) is always used such that the quantifier has narrow scope. I think that each one of them can be used such that we get something very much like this Fregean scope difference. Each one of these sentences can be uttered such that the utterance will only be true if Ponce is searching for something real, and they can be uttered such that it can be true even if he is only looking for something that isn't real, like the Fountain of Youth. In a certain context of an utterance it might be clear that we are only interested in real things, and when I utter (82) in this context it might only be true if Ponce is searching for something real. Similarly, in certain other contexts we might not care whether or not what he is looking for is real. In these context (81) can be uttered such that it is true even though Ponce only looks for something that isn't real. In fact, we have seen one example of this above, where someone said that there is something Fred is looking for, namely Sherlock. But the two readings that these sentences have don't come from a difference in scope, but from a difference in what the quantifier can contribute to the truth conditions.

2.4.3 Quantification and non-literal discourse

Sometimes we pretend that something is the case, and we then go from there. We act within the pretense that this is the case. This can happen in communication, as well as in other activities, like playing. One might think that in the above examples where we used non-Quinean quantifiers that these should better be understood as quantifiers within such a pretense. In particular, such quantified statements might not be literally true, but only true within a certain pretense.

I think the phenomena that we pretend certain things and then talk within that pretense, just like it were real, is an important phenomenon. In fact, I will later take recourse to something like it. But I don't think it can apply to our case here. I might very well know that Fred admires Sherlock, and that Sherlock doesn't exist. Still, I might not give you all the information I have. I want you to find out for yourself. I will only give you hints, not the full truth. I tell you

(93) There is someone Fred admires very much, and whoever that is, he is also admired by many detectives, and is smarter than Clinton. Come on, you should know by now....

I would be a mistake to think that in this situation I pretend that Sherlock is real, and then within that pretense I use a domain conditions quantifier to quantify over him. Pretense theory doesn't help us here, but it will help us later.

2.5 Conclusion

Let's sum up what we have seen in this chapter.

- We have seen that ordinary quantifiers like "something" are semantically underspecified with respect to whether or not they impose domain conditions or play a certain inferential role. This is the most plausible, but not the only possible, way of dealing with a number of examples where we use a quantifier apparently differently than on other occasions. I have argued that (at least some) of these uses arise from a need to communicate incomplete information.

- We have seen that the quantifiers in their internal, inferential role, use increase the expressive power of the language in which they occur, and that the simplest way for them to contribute to the truth conditions to get the inferential role for which we need them is to be truth conditionally equivalent to certain infinite disjunctions or conjunctions.

- In addition, we have seen that in the internal uses quantifiers don't have a direct connection to ontology. It requires further work to determine whether or not our ontology has to contain certain entities for our talk to be literally true. For example, we will have to further see if what Fred admires is something real or not.

- Whether or not a quantifier is used in its internal or external reading depends on the details of the particular occasion of its use. There is no simple syntactic test, nor any other simple way to determine this. In particular, one can't simply ask the speaker which one of the two ways they intend to be the right one on this occasion. Just like in the case of reciprocals or plurals, the speakers might not be aware of the fact that there are different reading of the sentences they utter.

This general story about quantification gives us a first look at how neo-Carnapianism about ontology could be true. By neo-Carnapianism about ontology I mean the view about ontology that is inspired by Carnap's internal-external distinction about questions about what there is.[14] Carnap thought that the question "Are there Fs?" could be understood in two ways, internally and externally. If it is understood internally then it has a trivially correct affirmative answer. If it is understood externally then it is, according to Carnap, meaningless (i.e. without cognitive content). Given the above account of of natural language quantification, we can see how something like this is true. If quantifiers can be used both internally and externally (in the above sense, not Carnap's) then a question "Are there Fs?" can be understood in two ways, depending on how the quantifier is understood. If understood internally then a true statement of the form "t is F" will answer the question affirmatively, whether or not "t" is a referring expression that stands for something real. If understood externally then an affirmative answer will require that reality contains an F, that there is an F out there as part of what is real.

Now, there is an important difference between this way of understanding that there are two questions that one can ask with the words "Are there Fs?" and Carnap's way. According to Carnap, one of the two questions one might be asking is really meaningless. According to the above way of understanding it, both ways of asking the question are completely meaningful and factual. In both ways the question will have an objective answer.

One important similarity with Carnap is, however, that internal questions about whole kinds of things, like "Are there numbers?", will have an affirmative answer, and they will have such an answer trivially. The positive answer to this question, understood internally,

[14] See (Carnap 1956).

follows immediately from "2 is a number.". But this by itself doesn't answer the question whether or not there are numbers, understood externally. The fact that the internal questions have trivial affirmative answer doesn't mean that the external questions have trivial negative answers, nor that they, too, have trivial affirmative answers. To decide this question we will have to look at what we are saying when we are saying that 2 is a number. Are we referring to some object when we say this, or are we doing something else? And if the second, what are we doing?

Quantification does not provide a cheap route to ontology. Simply because a certain quantified statement is true one can't conclude that our natural ontology contains entities of a certain kind. But quantifiers in their external use do provide a direct way to ontology. So, how are we using certain quantifiers when we talk about, say, numbers, or properties? To decide this we will have to have a closer look at what we are doing when we talk about them. This will determine whether or not properties, propositions and natural numbers are in our natural ontology. We will have to see how and for what purpose we quantify over them, and we will have to see what we are doing when we talk about them with quantifier free statements. And this will bring us back directly to the original puzzle we started out with, and conveniently leads us directly into the next chapter.

Chapter 3

Nominalizations

3.1 Introduction

In chapter 1 we started out with some general puzzles about ontology, in particular about the ontological presuppositions that our ordinary talk has for is literal and objective truth. In chapter 2 we have seen that quantification by itself does not give a cheap route to ontology. Further considerations are necessary to see whether or not a particular use of a quantifier phrase has ontological relevance. This was related, in a way still to be determined precisely, to the trivial arguments that we saw in chapter 1. But it alone doesn't give us an account of whether or not the trivial arguments are valid, and whether or not they show anything about ontology. The main part that is missing to determine this is to see what is going on in the first step of these arguments. Understanding the first step in the trivial arguments will get us closer to an account of what the function is of quantifier free talk about properties, propositions and natural numbers. This is closely related to whether or not the quantified statements that seem to follow from ordinary, apparently non-metaphysical, statements in the trivial arguments should be understood internally or externally. Consider, again

(94) Fido is a dog.

(95) Thus: Fido has the property of being a dog.

(96) Thus: there is a property Fido has, namely being a dog.

If (95) indeed follows from (94) (we will look at this in detail in this chapter), then it will depend on what the expression "the property of being a dog" is doing on this occasion whether or not (96) follows from (95) if the quantifier in (96) is understood externally. To be sure, we have seen that on the internal reading of the quantifier the inference will always go through. The question we will have to look at is whether or not the inference also goes through if the quantifier is understood externally. And that will depend on what the expression "the property of being a dog" does on this occasion. What the function of this expression is on this occasion will be the main topic of the first part of this chapter.

In the remainder of this chapter we will look at a certain general view about the function of noun phrases that misleads people into having a certain view of expressions like "being a dog" or "that Fido is a dog". I will discuss this general view in the second half of this chapter.

Expressions like "being a dog", "that Fido is a dog" and "the number of men" have been called nominalizations in the philosophical and some of the linguistic literature. They are called this way because they can be understood as being the result of some process that takes an expression that isn't a noun phrase and canonically turns it into a noun phrase. So, any sentence "S" can be transformed into a that-clause "that S". Any verb phrase of the form "is F" can be transformed into "being F", etc.. These resulting expressions syntactically then behave like terms, or noun phrases. In the linguistic literature it is a bit controversial whether or not these expressions are properly called noun phrases, and thus whether or not they should be called nominalizations,[1] but we shall not be concerned with this. This is merely a terminological matter, and should lead to no confusion as long as the use of this phrase is clarified. In the following "nominalization" can be understood as an expression of the following kinds: "that S", "being F" and "the number of Gs". We don't have to get into subtleties about whether or not that-clauses are noun phrases.[2]

This chapter will only focus on the function of these expressions in a very limited number of cases in which they can occur. It would be a mistake to try to conclude from the discussion of this chapter what their function is on other occasions. I do not attempt to make a claim of this generality in this chapter. We will have a look at the more general case in the next chapter.

[1] "Nominalization" can be understood as either the process of transforming a phrase that isn't a noun phrase into a noun phrase, or as the result of this process. It should be clear in context which one is the relevant one.

[2] There is a substantial literature about nominalizations in linguistics, which deals with a much wider class of cases than we do here. A classic book in this area, from a syntactic point of view, is (Lees 1960). More relevant to the present discussion are (Asher 1993) and (Zucchi 1993).

3.2 Innocent statements and their metaphysically loaded counterparts

The first part of the trivial arguments, the inference from (94) to (95) gives rise to a puzzle all by itself. It might be called **the puzzle of the innocent statements and their metaphysically loaded counterparts**. It is based on the strange feature that sometimes we seem to introduce talk about properties, propositions, and the like, apparently without change of truth conditions, and apparently so that we notice right away that this is without change of truth conditions. But how such innocent statements relate to their metaphysically loaded counterparts is not at all clear, and in fact quite puzzling. Let me explain.

By an **innocent statement** I mean a simple everyday statement that seems to have as little to do with metaphysics as anything. Here are some examples of innocent statements:

(97) I am hungry.

(98) Fido is a dog.

(99) Three men play soccer.

(100) John kissed Mary at midnight.

It prima facie sure seems surprising that from any one of them anything of metaphysical interest should follow, much less follow almost immediately. However, just that seems to be the case. All these statements have one or more of what I'll call **metaphysically loaded counterparts**. These are statements that can be obtained from the former by some kind of canonical procedure, some kind of transformation, and that are apparently equivalent to them. I'd like to point out right away that I am not claiming nor assuming that there is some kind of asymmetry between them, for example that the innocent statements are in some sense more basic. The procedure goes both ways. From the metaphysically loaded statements you can obtain the innocent ones by inverting it. The above statements have at least the following counterparts:

(101) It's true that I am hungry.

(102) Fido has the property of being a dog.[3]

[3]Or equivalently and more colloquially: Fido has the feature, or characteristic, of being a dog.

(103) The number of men who play soccer is three.

(104) John kissing Mary happened at midnight.

It seems that they are truth conditionally equivalent to the innocent statements simply because it seems it can't be that it's true that I am hungry, but I am not hungry. Or that I am hungry, but it's not true that I am hungry. And it seems that it can't be that the number of men who play soccer is three, but no three men play soccer. Or that three men play soccer, but their number isn't three. However, these statements don't seem to be metaphysically innocent any more. Whereas the innocent statements only were about people and dogs, their metaphysically loaded counterparts are about truth, propositions, numbers, properties and events. And that's hardly metaphysically innocent. After all, we explicitly talk about things that seem to be the proper domain of metaphysics.

Furthermore, it seem that we can start with necessarily true or even analytic, if I may use the word, premises. We could start with

(105) There are just as many bachelors as unmarried men.

(106) Fido is a dog or he isn't a dog.

and go to their loaded counterparts:

(107) The number of bachelors is the same as the number of unmarried men.

(108) Fido has the property of being a dog, or he doesn't have that property.

And again, we explicitly started talking about properties, propositions and numbers. This talk seems to presuppose the existence of properties, propositions and natural numbers. Thus any worries we might have had about their existence should be answered.

But how could we have gotten that so easily? Weren't metaphysics and ontology supposed to be difficult? How could we get results in ontology with arguments that start with completely innocent premises, even analytic premises, in just three simple steps? After all, how could conceptual truths have any implications about what there is? It seems intuitively right that to find out what reality is like you have to go beyond conceptual truths and confront reality.[4] But the above arguments seem to obtain ontological results out of

[4]As we will see, contrary positions have been defended by Frege, in (Frege 1988), and neo-Fregeans, like Wright, in (Wright 1983), and Rosen, in (Rosen 1993), amongst others.

nowhere. Our talk about properties, propositions and numbers has what has been called the something-from-nothing feature.[5] Talk about them can be introduced apparently without change of truth conditions. And it leaves us with the question: how can we get something from nothing?

Perhaps even more puzzling, and even more important, than this is the fact that if all of the above is true then it seems that the truth of our ordinary innocent statements depends on there being properties, propositions, numbers and the like. If (98) implies the existence of propositions then (98) can't be literally and objectively true if their aren't any propositions. But it seems that whether or not there really are propositions is an open question. At least it isn't obvious that there are. But if it is in doubt whether there are propositions, shouldn't it be in doubt just as well whether Fido is a dog? Sure enough, there is Fido, and he looks and barks just like a dog. But how can I be sure he is a dog if I am not sure whether there are propositions? After all, the above argument seems to show that we can be sure that if Fido is a dog then there are propositions. So, we can be sure that if there aren't any propositions then Fido isn't a dog.

This seems to be a general situation. If a certain statement (like (98)) implies the existence of certain entities (like propositions) then I can use this either in an argument for the existence of these entities, or as a reason casting doubt on the (objective and literal) truth of the statement, taking recourse to doubt about the existence of these entities. But it isn't clear which one. Again, one person's modus ponens is another person's modus tollens.[6] There is a way out of this dilemma, but it doesn't always seem possible to take recourse to it. For example, if we have independent evidence for the existence of the entities that are implied to exist then we can break the circle. And this seems to be just the case with ordinary statements like:

(109) John smokes.

Sure enough, (109) implies the existence of John, and any doubts about his existence will give rise to doubts about the truth of (109). But in the case of John we have other means to persuade ourselves that there is such a thing. Not so, it seems, for numbers, properties and propositions. The only way for us to persuade ourselves of their existence is through the truth of statements in which we talk about them, or perhaps through their role in the

[5] See (Schiffer 1994).

[6] This is the ontology-objectivity dilemma that we discussed in chapter 1.

grand theory. But do we have to wait for the grand theory before we can be sure that Fido is a dog?

Whatever one wants to say here, the apparent equivalences between the innocent statements and their metaphysically loaded counterparts give rise to a puzzle in metaphysics. If the innocent statements are indeed truth conditionally equivalent to their loaded counterparts then it seems that anything trivially implies that there are properties and propositions. This might be taken to show that the above conception of ontology is wrong, or that there is something about the special nature of these entities. Or, radically, it might be taken to cast doubt on the truth of any ordinary statement. If they are not truth conditionally equivalent then we should have an account of why we take recourse to the loaded statements so easily, and why we judge then to be truth conditionally equivalent with the innocent ones. And we have to make sense of such apparent absurdities as claiming that even though Fido is a dog, it is still an open question whether or not it's true that Fido is a dog. After all, it seems to be a conceptual truth that Fido is a dog iff it's true that Fido is a dog.

3.3 The standard accounts

3.3.1 Frege

The modern debate about the relation between the innocent statements and their metaphysically loaded counterparts goes back to Frege's *Grundlagen*, (Frege 1988). Frege there was concerned with the question whether or not numbers are objects. He observed that in uses like

(110) Jupiter has four moons.

numerals like "four" apparently were not used in the way that expressions that stand for objects are used. Their use seems more similar to that of adjectives, as "green" is in

(111) Jupiter has green moons.

However, in other uses "four" does seem to be used like expressions that stand for objects, namely

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

This seems to be similar to

(113) The composer of *Tannhäuser* is Wagner.

It seems that both (112) and (113) are identity claims with which it is said that what two singular terms stand for is identical. And thus numbers are objects, since objects are what singular terms stand for.

Frege wasn't too clear about how we should understand the uses of numerals as in (110), in which they apparently aren't singular terms. Or at least it isn't clear to me what Frege thought. Are we to believe that English contains two words spelled "four" that belong to two different semantic categories? If not, how do the two uses relate to each other? One might be inclined to understand Frege as claiming that we really always refer to numbers when we use numerals, even when we use them like adjectives. Such uses might be based on a confusing way of speaking that occurs in an imperfect natural language.[7]

Frege did accept that (110) and (112) are truth conditionally equivalent. Let's call a statement affirming the truth conditional equivalence between an innocent statement and one of its metaphysically loaded counterparts a **Frege biconditional**.

(114) Necessarily, Jupiter has four moons iff the number of moons of Jupiter is four.

is a Frege biconditional. Frege did accept Frege biconditionals, at least for the case of numbers. Anyone who accepts the truth of Frege biconditionals apparently has to accept the trivial arguments we started out with. Thus anyone who accepts them has to deal with the apparent puzzles we started out with, in particular how we could have gotten substantial ontological results from metaphysically innocent premises. One way of understanding Frege is that we can do this because we never started out with metaphysically innocent statements. If "four" is always a referring expression then it should be no wonder that (110) implies the existence of numbers.

[7]See (Frege 1988, §57). Frege, of course, placed much more focus on the examples of a statement saying that there are just as many As as Bs, and noting that this is equivalent to the number of As being identical to the number of Bs. The results of this paper will be most relevant for this, too, but we will not deal with Frege's overall strategy in any detail here.

3.3.2 The standard alternatives to Frege

One alternative to Frege[8] is a radical proposal by Schiffer.[9] According to him we should take it quite seriously that we get something from nothing in the trivial arguments. This is so because through our transition to the metaphysically loaded counterparts we, in a sense, create the new entities they are about. Through our practice of talking about "that Fido is a dog" and "being a dog" we create properties and propositions, just like certain practices of speaking can create fictional characters.[10] According to this view, Frege biconditionals are true, and there is no threat from ontology to the truth of ordinary claims about Fido, since the existence of properties and propositions is guaranteed (they are created).

A further alternative, endorsed by Field[11], is to deny the truth of Frege biconditionals, and offer an account of why they appear to us to be true. One might think that they are only true assuming that there are numbers, properties and propositions. However, they are not true simpliciter.

We thus have three answers to the question how we could have gotten something from nothing with the trivial arguments:

A1 We don't get something from nothing, because we never started with nothing. We started with something, even though it was hidden.

A2 We do get something from nothing. This, however, reflects the nature of the entities we get. They are second class entities that are created by a certain practice of speaking.

A3 We don't get something from nothing because Frege biconditionals are not true and thus the innocent statements don't imply their metaphysically loaded counterparts.

3.3.3 Some problems with the standard accounts

All of these proposals have certain implausibilities of their own, but more importantly, they seem to me to overlook a central aspect of what is going on in the relation between the

[8]Or better: to Frege as described above. It might be that Frege should really be understood along the following lines, but since Frege isn't the issue here, let's not worry about this right now.

[9]Schiffer's proposal isn't very standard, but I call it that way anyway because he agrees with the other standard proposals about the basic setup. An earlier view that is different from his present one, but in certain respects closer to the one defended in this paper, can be found in (Schiffer 1987). Both of these positions will be discussed in some detail in section 6.9 below.

[10]See (Schiffer 1996b).

[11]See (Field 1989a).

innocent statements and their loaded counterparts.

Some implausibilities

For individual implausibilities, let me briefly mention, but not discuss, the following, listed in the order of the different standard accounts as above:

1. What is the basis for claiming that in ordinary uses of (110) we use “four” as a referring expression? It sure seems that the analysis this sentence should be given takes “four” to be a determiner that together with “moons” forms a quantified noun phrase. Determiners include “some”, “many” and the like, which sure don’t seem like referring expressions.

2. Can we really accept the claim that numbers, properties and propositions are created by what we do when we talk about them?

3. Is there any basis to deny the truth of Frege biconditionals other than ontological worries? After all, intuitively we would certainly judge them to be true.

A suspicious notion

All of the standard accounts agree with Frege that

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

is just like

(113) The composer of *Tannhäuser* is Wagner.

Both of them are identity claims involving two *singular terms*. And whatever the singular terms stand for has to exist, in order for these claims to be true. However, I think that taking recourse to the notion of a singular term is not helpful here, and in fact misleading. Even though this notion is intuitively appealing, it does not seem to pick out a class of expressions that have all that much in common, other than being, in a sense, about a single thing. Even the prime examples of singular terms, referring expressions and descriptions, are really of two quite different semantic categories. Both can be said to “be about a single thing”, but this is true only in two different sense of this expression. Referring expressions are about a single thing in the sense that they have it as their semantic function to refer to whatever they refer to. Descriptions, on the other hand, are quantifiers that have a certain uniqueness build into them. The are about a single thing in the sense that there has to be a unique thing having a certain property for a statement involving a description to be true. In this sense the description can be said to be about that thing. But this is really quite different.

Furthermore, there are a number of expressions that might also be seen as being singular terms, but that intuitively have nothing to do with the existence of certain entities. Consider, for example,

(115) What Mary likes to do on Sundays is fishing.

(116) Softly is how Mary spoke.

(117) How Mary made the chocolate cake is with a lot of care.

These can be seen as singular terms because they seem to satisfy what is commonly believed to be required to be a singular terms, like occurring in true "identity" claims, as in:

(118) How Mary makes a chocolate cake is identical to how my grandmother used to make it.

Does this show that "how Mary makes a cake" refers to some entity? Does this really have any metaphysical or ontological consequences?

It might seem that taking recourse to the notion of a singular terms is a red herring, and we can simply talk about referring expressions and descriptions instead. This seems to be all that is required when one claims that

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

is just like

(113) The composer of *Tannhäuser* is Wagner.

However, "the number of moons of Jupiter" is not a description, and is in important respects quite different than "the composer of *Tannhäuser*". A description "the F" has a close relation to another quantifier, namely "a F". With "a F" we claim that there is at least one F, with "the F" we also claim that there is at most one. In particular, "the F is G" implies "a F is G". Thus

(113) The composer of *Tannhäuser* is Wagner.

implies

(119) A composer of *Tannhäuser* is Wagner.

and (113) further claims that *Tannhäuser* had one composer. However, the relation between "the number of F" and "a number of F" is quite different.

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

doesn't imply the nonsensical

(120) A number of moons of Jupiter is four.

However,

(121) A number of moons of Jupiter are covered with ice.

makes perfect sense, and says that several of Jupiter's moons are icy, whereas

(122) The number of moons of Jupiter is covered with ice.

is again nonsensical. Note in particular that "the number of moons" is singular, whereas "a number of moons" is plural.

So, perhaps the Fregean analysis of these examples missed some important point. Perhaps (110) isn't really closely related to (113). But if not, what is going on in this example? And how does it relate to (110)? Before we look at a different account, we should have a look at a rather general worry with the standard accounts, which I think will get us in the right direction.

A general worry

A key to understanding what is going on here seems to me to have a closer look at the following puzzle, which I'll call **the puzzle of extravagance**. It becomes especially clear if we assume for the moment that Frege biconditionals are true.

It should be beyond dispute that one of the core functions of language is to communicate information. There are other things we do with language, like joking and flirting, but these uses of language seem to be derivative on the main function that it has, namely to get information from the speaker to the hearer, or to request information from the hearer. With this in mind, we can ask ourselves, why should it be that our language (and many, perhaps all, others) have systematically two ways to say the same thing? If I want to get the information across how many men play soccer, why do I have two ways to do it, by either saying that three men played, or that their number is three? It seems that if the Frege biconditionals are true then I systematically have two or more sentences that I could utter to communicate the same information. Why would our language give us these extra tools?

We can look at the same puzzle from a more Gricean / pragmatic point of view. All the metaphysically loaded counterparts of the innocent statements involve more words and more elaborate sentence structure. They are, in a word, more complicated. But it is one of Grice's maxims that I should be as cooperative as possible, and communicate the information I have in an as simple, relevant and elegant way as I can. If I don't do this, and if I use more complicated ways to communicate than necessary, then I will often try to do something else besides communicating the information I have. In particular, the hearer's recognizing that I violate one of the maxims will draw the hearer's attention to that I do more than just try to communicate what the sentence uttered literally means. If that is so, I might be trying to do something over and above just communicating a certain information when using the loaded statements instead of their innocent counterparts. But what?

It might be that the loaded counterparts are merely an accident. For example, the conjunction of two sentences is always a sentence, even if the two sentences are the same. Therefore

(123) Water is wet and water is wet and water is wet and water is wet.

is a well-formed sentence that is truth conditionally equivalent to

(124) Water is wet.

Thus a general fact about the language accounts for this case of extravagance. But note that in this case (123) does not play any role in communication. Sure enough, it is a meaningful sentence, but not one we would ever have any need to use it in communication. The puzzle of extravagance only really arises in cases where we do have a function for the sentences in communication. Do we have one for the metaphysically loaded counterparts?

In the philosophy of mathematics or metaphysics literature, where Frege biconditionals are discussed, the authors usually take sentences like (110) and (112) and reflect on their truth conditional equivalence. However, they usually don't discuss whether or not these have a different use in actual communication, and if so, what this difference is. This is a bit surprising, I think, because it seems to be quite obvious that their use in communication is quite different, and that seems to be so independently of the issue of their truth conditional equivalence. To illustrate this, let's consider the following example.

I visit your town for the first time, don't know my way around well, and would like to get a quick lunch. You suggest a pizza place and a bagel shop that are close. Half an hour later you see me again and ask me what I had for lunch. I reply

(125) The number of bagels I had is two.

This seems quite odd. To be sure, it will be clear to you now that I went to the bagel shop, and had two bagels there. But the way I did put this is strange. It seems that to put it this way something else must have happened in the conversation first. We'll see more about this in a second. Note that it is perfectly normal to say

(126) I had two bagels.

This isn't odd. I told you I had bagels, and furthermore that I had two of them. Intuitively, in this situation is is more important what I had, not how many I had. And in (126) this is communicated just fine, but in (125) I somehow give too much prominence to how many I had, not what I had. Similar oddness in this situation is attached to

(127) It's true that I had two bagels.

but I will concentrate on the other example for now.

We can see quite easily that the role in communication that the innocent statements and their loaded counterparts have is different in certain respects. But what is their difference, and how does this relate to our puzzle in metaphysics we started out with?

3.4 A new account

An account of the relation between the innocent statements and their loaded counterparts has to include the following:

1. An account of our intuitive judgments of the their truth conditional equivalence.
2. An account of whether or not the meaningfulness of one of the counterparts is an accident, or whether they have a function in communication. This should solve the puzzle of extravagance.
3. An account of what is going on in the something-from-nothing arguments we started out with.

As it turns out, there are a number of very instructive cases where the first two of these issues arise. They play no role in philosophy, but, as I will argue, are most relevant here for us to see what the relation is between the innocent statements and their loaded counterparts.

3.4.1 The function of the cleft construction

A good example of this is the so-called cleft construction: “it is X that Y”. With it one can say what one says with

(128) Johan likes soccer.

also with two different, but truth conditionally equivalent sentences, namely

(129) It is Johan who likes soccer.

(130) It is soccer that Johan likes.

Here, too, we have (apparently) truth conditional equivalence, but we use more words and a more complex sentence. So, what is the difference? When would we use one, but not the other? Is the meaningfulness of these sentence an accident, or do they play a role in communication? What is the function of the cleft construction? Why would a language want to have it?

The answer is quite straightforward. Even though we communicate the same information with (128), (129), and (130), we do so in a different way. In an ordinary utterance of (128) the information that Johan likes soccer is communicated neutrally. No particular aspect of the information is given a special status. In an ordinary utterance of (129) or (130) this is not so. Some aspect of what is communicated is given a special status. It is stressed. The common term for this phenomenon is **focus**. The focus of an utterance of a sentence is the aspect of what is said that is given a special stress or importance. The clefted sentences do just that. The above ones focus either on Johan, or on soccer, respectively.

Clefted sentences are by no means the only way in which one can achieve a special stress. An utterance of (128) can be used to communicate the information in a non-neutral way, too. A speaker could phonetically stress one aspect or another, and this way of doing it is certainly the most common. A speaker could utter

(131) Johan likes SOCCER.

or

(132) JOHAN likes soccer.

The capitals here stand for a rise in one's voice. An utterance of one of (131) or (132) would not present the information neutrally. It would rather stress the fact that it is Johan who likes soccer, and not someone else, in the case of (132). Or that it is soccer that he likes, and not something else, as in the case of (131).[12] Thus (131) is a lot like (130), and (132) is a lot like (129). In each case, they both can be used to communicate the same information, and in addition, they both have the same focus. So it seems that in English we have at least two ways to achieve focus. We could raise our voice, or we could us a clefted sentence, where the focused item is put in a distinguished position. In written English, however, we only have the cleft construction and not the focus results through intonation (unless, of course, one introduces representations for phonetic stress in written English, as the capital letters do). In addition, these two ways can be combined, to achieve even more effective focus. When I say

(133) It is SOCCER that Johan likes.

[12] This kind of focus is called contrastive focus, and is just one kind. See (Lambrecht 1994), (Rochemont and Culicover 1990), or (Rochemont 1986) for much more on focus and its relation to syntax and semantics.

I draw even more attention to what Johan likes than with (131) or (130).

In communication we take recourse to the cleft construction for that purpose, to achieve a certain focus effect. And we do have a use in communication to do so (I will elaborate on this in a minute). We can thus say:

F1 The function of the cleft construction is to focus on the clefted item.

The clefted item is the one that was extracted and given a distinguished position, for example "Johan" in (129). The cleft construction is not the only construction we have for doing this. There are a number of constructions where a certain item in a sentence gets extracted and put in a distinguished position. The truth conditions of the sentence are unaffected by this, but a certain focus affect is achieved with it. Consider, for example, the pairs:

(134) John ate the beans.

(135) The beans is what John ate.

and

(136) Mary spoke softly.

(137) Softly is how Mary spoke.

Cleft constructions and the like show that there are certain syntactic constructions that have as their function to present information in a certain way. That is why we take recourse to them in communication. There are other, more simple, sentences that could be used to communicate the same information. But we don't take recourse to them because we want to achieve a certain focus, and using an equivalent clefted sentence is one way to achieve this.

3.4.2 Focus and communication

In communication we transmit information from one person to another, either through making an utterance with the right truth conditions, or by requesting certain information from one of the participants in the communication. But for this to work effectively not any sentence with the right truth conditions will do. A number of aspects have to be considered that will be relevant to choose among different sentences to utter:

- What is the shared background knowledge of the participants in the communication?
- What will be new information, and what is old information?
- Is there any misinformation held by one of the participants in the communication?
- What is important and what isn't for the present purpose?
- and so on and so forth.

To communicate effectively it is important to make clear what is new, what is important and what is supposed to be a correction of an earlier misunderstanding, for example. We do this through communicating information in a certain way, and this, amongst others, is what focus contributes to. Focus is an important aspect of communication. We will see examples of focus at work below.

3.4.3 The function of the loaded counterparts

We have seen that natural languages contain syntactic constructions that have as their function to communicate information with a special focus. The cleft construction gives the focused item a syntactically special position, which makes sure that it gets special attention. That is what this construction is good for and why we use it. Even though we can say what we say with a clefted sentence uttering a simpler sentence, we use the clefted sentence because using it we achieve a certain focus. As you can guess, it is this that I claim is the reason we use the metaphysically loaded counterparts. It, I claim, is their function. To see this, I would like to point out when we actually use these sentences in communication.

Starting with Frege, the tradition in metaphysics restricted itself to noting that the metaphysically loaded counterparts are true just in case the innocent statements are true. Philosophers in that tradition then were quick to conclude that we have two "singular terms" here that figure in an identity statement. Very little, however, is ever said about when we actually use the metaphysically loaded statements in ordinary communication, in particular when the innocent statements would not serve our purposes just as well. It seems to me that looking at this is most instructive to see what the function of these loaded statements really is. After all, it would be surprising if their only real function in communication is debates in metaphysics. So, lets see when we use them in ordinary everyday communication. As we have seen above, in the bagels example, they sometimes certainly have a different overall effect. Sometimes it is weird to use the loaded counterparts, and sometimes it is appropriate and even preferable to do so. Looking at an example should make this clearer. It seems to me that a good example of when we use the loaded statements is certain cases of miscommunication, or rather, attempts to correct miscommunication. And these are exactly the situations when we use cleft constructions. Let me explain.

Suppose you are a secret agent on an important mission, and your job is to observe a street corner and communicate to your superior on intercom what is going on at the street corner. As you observe the corner three men enter the bar that is located at that corner. You report:

(138) Three men entered the bar.

you superior is confirming with

(139) OK, got it, three men entered the bus.

You see that you didn't get it across and try again

(140) No, three men entered the BAR.

or alternatively

(141) No, it's the bar that three men entered.

This should not only make clear that there is miscommunication, but also which part of what you said you didn't get across. Your superior tries again:

(142) Oh, my mistake. Now I got it: three men entered the BART.[13]

You have to try again, this time even more forcefully:

(143) No, it's the BAR that three men entered.

Taking recourse to a clefted sentence or to intonation have the effect of focusing on a certain aspect of the information you try to communicate. However, using the cleft construction

[13] BART is a local train in San Francisco, which is always written in capitals. No special intonation here.

and, at the same time, a certain intonation is most effective in pinpointing the particular part of what you try to communicate and which your superior doesn't get.

And similarly if your superior get's it wrong in a different way. Suppose things would have gone as follows. Again, you report

(144) Three men entered the bar.

You superior get's it wrong:

(145) Ok, got it, free men entered the bar.

you try to correct

(146) No, THREE men entered the bar.

or alternatively

(147) No, three is the number of men who entered the bar.

Again, your superior:

(148) Oh, my mistake. Now I got it: thirteen men entered the bar.

At that stage you will say:

(149) No, THREE is the number of men who entered the bar.

If you have difficulty communicating how many men entered the bar you will use "the number of" construction, or to intonation. If the miscommunication continues you will use a combination of the two. It is these situations where we use this construction in regular communication. And in these situations we are interested in stressing certain aspects of what we try to communicate. Namely the aspect that the hearer misunderstands. Note also that if your superior would repeatedly misunderstand you as talking about a kind of men, and not how many men, you could use the "number of" construction effectively to make this clear that this is the misunderstanding. Suppose your superior would have answered, after you tried to correct the first understanding, with

(150) Oh, my mistake. Now I got it: tree men entered the bar.

You you would be advised to use

(151) No, the NUMBER of men who entered the bar is THREE.

Note further that the cleft construction is not available in this case. One can't say, for some reason or other,

(152) * It is three that how many men entered the bar.

or anything like that. The cleft construction allows us only to focus on certain aspects of what is said, namely something that is expressed with a full noun phrase, a noun, an adjective and a few more. If we want to focus on other aspects we have to take recourse to other means, like the "the number of" construction.

This gives rise to the following thesis about the function of the "the number of" construction, as it occurs in the loaded counterparts.

F2 The function of the "the number of" construction, as it occurs in the loaded counterparts, is to focus on how many things of a certain kind one is talking about.

This is limited to the loaded counterparts, because so far we have only talked about them. Also, so far I have only talked about the "the number of" construction. We will have to look at the other ones, namely the "has the property of", "it's true that" and the event construction separately. I don't have the space to go through all of them here, but I do want to make a proposal about their function, too:

F3 The function of the constructions that give rise to the loaded counterparts is not to present new or different information than their innocent counterparts. It is rather to present the same information in a non-neutral way. In using one of the loaded counterparts in communication we present the same information that we would present with the innocent counterparts, but we do so with a special focus, or stress, or the like.

Of course, I haven't said for each case what the focus or stress is that the loaded counterpart brings with it. But this we can see if we look at when we use them in communication, and when it is awkward to do so. For example, in order for one to utterer

(153) It's true that Johan likes soccer.

without awkwardness it has to be a prior topic of the conversation whether or not Johan does like soccer. In uttering this sentence rather than it's innocent counterpart you give it special force that is supposed to calm down and persuade people who thought otherwise.

Let's draw some conclusions.

- We have seen that focus is an important aspect of communication. In order to successfully deal with real life situations of communication we sometimes need to stress particular aspects of the information that we are communicating.

- We have seen that natural languages have syntactic constructions that have as their function to focus on a particular aspect of the information they are used to communicate. This is to say that a) these constructions do have a role in communication, as opposed to being accidentally well-formed, and b) the role they have is to focus on particular aspects of the information communicated. The cleft construction was one such example, and certain *how* constructions were another.

- We have seen that the "the number of" construction, as it occurs in the loaded counterparts, is used in communication to focus on a particular aspect of what said, namely how many of a certain kind of thing one is talking about.

- I have argued that the function of this construction, as it occurs in the loaded counterparts, is focus, and thus that it is analogous to the cleft construction. I have claimed, but not argued for in any detail, that this also holds for the other loaded counterparts.

We thus have two pictures of what is going on with the loaded counterparts. One is the picture that is shared by all standard accounts:

P1 The loaded counterparts contain singular terms over and above the singular terms that their innocent counterparts contain. For the loaded counterpart to be (literally and objectively) true the entities that these singular terms stand for have to exist.

The other picture is the one I have been pushing in this paper:

P2 The loaded counterparts communicate the same information as the innocent statements. However, in them the words are shuffled around a bit so that one of them, which is given special importance, gets a distinguished place in the sentence. This is done to achieve a certain focus.

Now, one might think that this is all fine, and that it is still a further issue whether or not these two views are in fact in conflict. Couldn't it be that we achieve the focus effect because we talk about certain entities? Can't a Fregean endorse everything I have said? To be short: it's not clear. One reason is that utterances of sentences like

(154) The president of the US is Clinton.

do not achieve a focus effect, unless, of course, through intonation. However, utterances of

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

always has a focus effect, independent of intonation. This is not to say that a Fregean can't come up with some clever theory that accounts for the focus effect after all. I haven't seen one, but I don't claim to have shown that there isn't one. It really all comes down to the big picture, to what we do when we talk about numbers, properties and propositions. According to the standard accounts we do, or at least try to, talk about entities that make up reality. The view about the loaded counterparts I have been pushing here, however, fits into a quite different picture. Before we look at the big issues more closely, let me just state that the view defended here provides a further, fourth, answer to the question how we could have gotten something from nothing by making the transition form innocent statements to the loaded counterparts (see page 92 for the others):

A4 We didn't get something from nothing, even though Frege biconditionals are true. We started with nothing, and we ended up with nothing, appearances to the contrary.

3.4.4 The scope of the present account

So far I have offered an account of the function of certain nominalizations on particular occasions of their use. It would be a mistake to think that this is already an account of the function of that-clauses in general. It only applies to particular tokens of that-clauses, to particular occurrences of them in certain utterances. This thus leaves the following questions open:

1. Are all uses of that-clauses and the like the ones we just discussed? Does the story about their function in these specific uses carry over to all of their uses?

2. If not, what is their function in other uses?

It is clear that the answer to the questions in 1. is “no”. There are many occurrences of that-clauses where they do not have as their function to focus on a particular aspect of the information communicated. Take attitude ascriptions as one example. When I say that

(155) John believes that surfing is fun.

then I am not using the that-clause to focus. Partly because there isn’t another way for me to say what I say with (155) other than using a that clause. I will have to use a that-clause to say what I say with (155). Thus we are left with answering question 2. The main question there is to say whether that-clauses function as referring expressions, or as something else. And this is closely related to asking what the function of our talk about propositions is. Is it to talk about some domain of entities, or something else? We will investigate this in the next chapter. But before we get to that, let’s have a look at what we now can say about the first of the two puzzles we started out with in chapter 1.

3.5 Towards solving the first puzzle

The first puzzle arose from there being apparently two correct but conflicting answers to the question: “How hard are ontological questions?”. One of these answers was that since ontological questions are questions about what the building blocks of reality are they are very difficult questions. The other answer was that basically everything we all accept trivially implies that there are properties, propositions and numbers, thus at least ontological questions about them are trivial. The reasoning for this second answer went via the trivial arguments. With these one could start from a metaphysically innocent statement, infer one of its metaphysically loaded counterparts, and then trivially infer from that a certain quantified statements which appears to claim that properties, propositions or numbers exist. After what we have seen so far, however, we can view these trivial arguments in a quite different light.

The first leg of the dilemma is unaffected by what we have seen so far. Ontological questions are difficult questions, or at least, we have seen no evidence to the contrary. But the second leg, the one pushing that ontological questions are trivial to answer, is very much affected by the above. The trivial argument proceeds in two steps. Step one is the transition from an innocent statement to one of its metaphysically loaded counterparts. In this chapter we have seen that such transitions are perfectly valid. Innocent statements are truth conditionally equivalent to their loaded counterparts. They only differ in how they present information. In particular, each one implies the other. However, the loaded counterparts do not talk about any more entities than the innocent statements. The extra noun phrases that occur in them do not function as referring expressions, but rather give a particular aspect of the information communicated a syntactically special position, to give it special prominence. Even though the noun phrase in the metaphysically loaded statements is not a referring noun phrase, the quantifier inference still goes through. We saw in chapter 2 that there is a need we have in communication for quantifiers that have a certain inferential role, and have it whether or not the noun phrases that occur in certain premises refer. Thus in particular, the quantified statement trivially follows if the quantifier is understood internally. And it is because it will be used internally on this occasion that the inference is so trivial.

The trivial arguments thus are perfectly valid, and they do imply that there are properties, propositions and numbers. But they do so only if the quantifier is understood internally. But a quantified statement in which the quantifier is used internally is ontologically neutral. Thus even though the trivial arguments imply that there are properties, propositions and numbers, this conclusion does not answer the ontological question whether or not there are properties, propositions and numbers. For this to be truly an ontological question the quantifier has be understood externally. The trivial arguments only have implications when the quantifier is used internally. Thus the trivial arguments are perfectly valid, but they don't answer the ontological questions. This is, in outline, the solution to the first puzzle.

The question whether or not it is also true when the quantifier is understood externally is still open so far. What I have said so far neither shows that the answer to the external question is affirmative or negative. We will get back to this, and the puzzles, in the following chapters.

3.6 A dogma about noun phrases

Before we can look more closely at what function our ordinary, and theoretical, talk about properties, propositions and natural numbers has, we should look at a general issue about the semantic function of noun phrases. The above discussion did touch on it already, in particular in the discussion of singular terms in section 3.3.3, but it is instructive to look at it more closely before we go on. There is a not uncommon opinion about the semantic function of noun phrases that implicitly surfaced in the position of the people who endorsed the above standard approaches to Frege biconditionals. This opinion is based on a dogma about noun phrase semantics we have very little reason to believe to be true, but lots of reason to believe to be false. In the remainder of this chapter we will look at these rather general issues about the function of noun phrases. After that we are done with investigating general features of languages, as we did in chapters 2 and 3. We can then return to the specific discussion of talk about properties, propositions and numbers, and finally to ontology. This will be done, in this order, in chapters 4, 5, and 6.

3.6.1 The dogma

What I said above about the function of particular uses of certain noun phrases goes contrary to what is often assumed in doing the semantics for noun phrases. This belief is based on a certain picture of the function of noun phrases which arises from a rather simplistic conception of what we do when we talk about objects. Because this belief isn't really based on any evidence I will call a dogma. The dogma is as follows:

(D) Every noun phrase is either a quantifier or a referring expression.

Here "either-or" has to be understood as an exclusive and exhaustive classification. Every noun phrase is either one or the other, and no noun phrase is both. The dogma can be restated in the following way, introducing some terminology. A noun phrase that would be a counterexample to this claim would be one that is neither a quantifier nor a referring expression. In other words it would be a non-quantificational, non-referential noun phrase. This can be abbreviated as NQ-NR noun phrase, and read out as **encuneral noun phrase**. Thus the dogma is equivalent to the claim that there are no encuneral noun phrases.[14]

[14] See (Hofweber and Pelletier) for an elaboration of the claim of this section, and many more examples of encuneral noun phrases.

Let me explain what this dogma comes down to, and then give some evidence why it is unjustified, and why we have reason to believe it is false.

(D) presupposes that referring expressions are different from quantifiers. We can happily grant this. We can concede that referring expressions are different than quantifiers with respect to what their semantic function is. What the semantic function of referring expressions is is intuitively quite clear. It is supposed to stick to an object an contribute it to the truth conditions of an utterance in which it occurs. What the semantic function of quantifiers is is a bit more difficult to state. Quantifiers of course include the traditional quantifiers, like "something", but also a much wider class that has been dealt with in generalized quantifier theory, like "many men", "at least three, but no more than twelve, men", etc.. It is not completely clear how to distinguish quantifiers from other expressions, at least not if one doesn't want to beg the question about (D). One way to distinguish them is to say that all noun phrases are quantifiers, and thus distinguishing quantifiers from the rest comes down to distinguishing noun phrases from the rest. But this certainly is unsatisfactory as a definition of what a quantifier is. It seems that quantifiers form a semantic category, and the question whether or not all phrases of a certain syntactic category belong to it shouldn't be settled by the definition of that semantic category alone. The above examples of generalized quantifiers certainly deserve that name because it is quite straightforward to show that they belong to the same semantic kind as the paradigmatic quantifiers. But that alone doesn't give rise to a definition of what a quantifier is. Here, too, we will have to leave this a bit vague. Quantifiers are expressions that belong to the same semantic kind as paradigmatic quantifiers. And generalized quantifier theory has shown that this class is a lot larger than one might intuitively think. What we will have to see is whether or not it includes all noun phrases besides the referring ones.

Furthermore, for (D) to have any real content, we have to distinguish an expression's being of a certain semantic category (like: being a quantifier) and that expressions being represented in a certain semantic theory of being of a certain semantic category. Expressions that are of one semantic category can be accommodated in certain semantic theories by putting them into a different semantic category than they are. This often has theoretical advantages, but shouldn't be taken to be in conflict with the fact that the expression really is of a different semantic category. A classic example of this is Montague's treatment of proper names. He assigned them semantic values of the same type that all noun phrases got assigned to, namely sets of properties. A name "Fred" has as it's semantic value the set of all properties that Fred has. That, of course, doesn't mean that "Fred" refers to that a set of properties, nor that "Fred" isn't a referring expression. Taking the set of properties that Fred has rather than Fred as the semantic value in this particular semantic theory has certain theoretical advantages, and that's why it is done. Doing this is in no conflict with he fact that "Fred" is a referring expression that refers to Fred. For referring expressions their referents don't have to be their semantic value. To be sure, their semantic value will be closely, or systematically, related to their referents. As in our case, the referent is closely related to the set of its properties. But semantic theory shouldn't be constrained by being allowed to take only one of them as semantic value. Sometimes taking the other might be more useful. We will see more on the use of semantic values in section 6.3.1 of chapter 6. All need now is that we have to distinguish an expressions being of a certain semantic category from its being treated to be of a certain semantic category in a certain semantic theory. (D) is a claim about the former.

It might seem that (D) is quite plausible, and that in fact it is supported by the success of generalized quantifier theory. Within that theory a variety of noun phrases have been accommodated in a truth conditional semantics. And thus they can be seen as being quantifiers (within that semantic theory, at least). However, generalized quantifier theory assumes (D) to be true. According to it a quantifier is merely an expression of a certain type, the type of noun phrases. It is trivial, given generalized quantifier theory, that every noun phrase is a quantifier. Generalized quantifier theory provides a very general and powerful way to accommodate noun phrases in a compositional truth theoretic semantics. But whether or not it will apply to all noun phrases is not at all clear. In fact, the examples given below seem to be a problem for (D) just as well as for the claim that generalized quantifier theory applies to all noun phrases.

It is not clear why we should believe (D). One reason is a certain simple picture of the function of noun phrases. They either stand for individual objects (and thus are referential), or they talk about a class of individual objects (and thus are quantifiers). But on reflection things seem to be more complicated. It in fact seems that there are a number of examples against this simple picture, and against (D). And, on reflection, it seems that the simple picture of the function of noun phrases should be abandoned for a more complicated one.

3.6.2 Evidence against the dogma

Before we can see how the simple picture of the function of noun phrases should be extended or modified, it will be helpful to see what apparent problem cases there are for (D), and what the function of these apparently problematic noun phrases is. I'd like to mention two kinds of such cases: generics and the average-construction.

Generics

Genericity is a common phenomenon in ordinary communication. It occurs when we make a general statement that seems to allow for certain not well defined exceptions. This occurs in two common kinds. First involving certain special noun phrases, as in

(156) a. The tiger is a fierce animal.

b. The potato contains vitamin C.

Secondly, by reporting some kind of regularity, as in

(157) a. John smokes a cigar after dinner.

b. Mary drinks red wine with dessert.

We will be concerned with the first kind only here.[15] The question we should focus on here is what the semantic function of the noun phrase "the tiger" is in (156a). To see this we should note the following:

- (156a) is true.
- (156a) is true, even though not every tiger is fierce all the time. Baby tigers and severely ill tigers aren't. Thus (156a) isn't equivalent to

  (158) Every tiger is fierce.

- "the tiger" isn't, or at least apparently isn't, a description. (156a) is in many important respects different than

  (159) The author of *Faust* is fierce.

[15] See (Carlson and Pelletier 1995) for a variety of articles on various aspects of generics, or (Pelletier and Asher 1997) for an overview of the debate.

With "the tiger" we don't seem to describe some thing and than say of it that it is fierce, contrary to "the author of *Faust*".

- "the tiger" isn't simply a referring expression that refers to a kind, the kind tiger. At least not if the verb phrase says that whatever the noun phrase refers to is fierce. That is so because the kind tiger isn't fierce. Only individual tigers are. So, if "the tiger" refers to a kind, then "is fierce" on this occasion doesn't say that what the noun phrase refers to is fierce, but rather must mean something more complicated, like that members of what the noun phrase referred to are fierce, or the like.

With these points in mind, let's ask ourselves, is the noun phrase "the tiger" as it occurs in (156a) a referring expression or a quantifier? And if neither one of them, what is it?

To be honest, we won't be able to decide this question here, and in fact, the semantics of generic expressions like these is quite tricky and has been discussed in an extensive literature.[16] However, I think it is quite plausible to say that these expressions are neither referring nor quantificational. They can't be taken to be referential without some extra bells and whistles. And there doesn't seem to be any obvious candidate for what quantifier they could be. A much more plausible account of their semantic function is based on the following.

It seems that even though generic expressions aren't exceptionless generalizations, they are *ceteris paribus* generalizations. All things being equal, a particular tiger will be fierce. This has a close connection to good reasoning with partial information. In real life we usually don't have enough information to let our reasoning be guided by anything as strict as logical implication. Simply because Mom said that dinner will be ready in 5 minutes doesn't imply that it will be. All kinds of things might go wrong. But it is nonetheless good reasoning to conclude that it will be. All things being as they usually are, it will be ready. If we would have much, much more information about the world, and much, much greater computing power in our brain we might be able to deduce that it will be ready. But we don't have that, and it seems it would be a waste of resources to try to have it. In real life we will have to do with less, and we can do with less. We will only have very partial information of the world around us, but still we can reason quite well with it. What we do is something like sticking to the simplest or most paradigmatic cases and assume that

[16] Again, see the articles in (Carlson and Pelletier 1995).

they hold and reason under this assumption. In addition, we only deviate from assuming that it is the most paradigmatic case if we have reason to do so. Then we take any special considerations into account, like that Mom hasn't been herself recently, but without such extra information we don't. This is what has been called *default reasoning* or *defeasible reasoning*. And generic sentences like (156a) seem to be closely related to this. If I know (156a) and I know that Fred is a tiger then it seems that I am justified to infer that Fred is fierce, unless I have overriding information to the contrary. So, without any special extra information I can conclude that Fred is fierce, even though all my premises might be true and the conclusion might turn out to be false because, to my ignorance, Fred is only a baby tiger. A true generic sentence is closely related to a valid inference rule in default reasoning.[17]

Generic noun phrases, like "the tiger" in (156a), help us to talk about tigers, but the information we communicate with them is much more loosely connected to individual tigers than with referring expressions, or quantifiers. They participate in making very precise and well cut claims about individual tigers. Generics don't do that. But that is also why we want them and why they are important. The figure in how we think, and have to think, about the world. The full complexity of when precisely tigers are fierce is way too much for us to handle in any effective way. We represent the lose generality that tigers are fierce with a generic sentence, and we use this in a very effective, but defeasible, way of reasoning about tigers. Nonetheless, such information is represented, and articulated, in subject predicate form. The fact that there is a singular noun phrase "the tiger" occurring in such a sentence gives rise to the impression that we are talking about some specific entity, but this is really a mistake. "the tiger" as it occurs in (156a) is a plausible candidate for a noun phrase that is neither referential nor a quantifier. I haven't shown this, to be sure, but I hope I made it plausible. Let's also, briefly, look at another candidate.

The average F

Another example that apparently is a counterexample to our above dogma (D) is the so-called average-construction, as in

(160) a. The average American has 2.3 children.

[17] See (Asher and Morreau 1995) and (Veltman 1996) for more on this. In particular also on how this relates to giving a truth conditional semantics for such sentences, a topic into which I won't get into.

b. The average American man prefers football and beer to port and ballet.

These two are different. In (160a) we say that the number of children by Americans divided by the number of (pairs of) parents is 2.3. In (160b) we make no such divisions, but rather say something like that this is what many or most American men prefer. Let's focus on the function of the expression "the average American" as it occurs in (160a) for now. Here, again, we can note the following:

- (160a) is true.
- "the average American" isn't, or isn't obviously, a referring expression. Partly because there is no such thing as the average American, and thus any attempt to refer would fail, and partly because speakers know very well that there is no such thing and thus don't even try to refer to it.[18]
- "the average American" isn't, or isn't obviously, a description, even though it is of the form "the F". Again, the description would fail to denote, and speakers would be well aware of this.
- It isn't clear what other quantifier this phrase could be.

However, it is completely clear what the truth conditions of these sentences are. In fact, or at least for (160a), this can be stated much more precisely as we were able to in the case of generics. "The average F has X G." is true just in case the ratio between the G's and F's is X. There is not much of a mystery about this expression, or how it works. It only becomes mysterious if one things that we are or try to refer or quantify with it. But why would we think that?

Encuneral noun phrases

Of course, I haven't shown that (D) is false. But I think I have made it plausible that

- it is not clear why one might want to think that (D) is true
- what the evidence for (D) is supposed to be
- and that there are a number of cases that seem to speak against (D).

[18] A different view about this will be discussed in chapter 6, section 6.9.

But these cases are not semantically mysterious. It is in fact quite clear what we are saying when we utter such sentences as (160a). To deny or at least be skeptical about (D) seems prima facie quite reasonable. When we look for the semantic function of noun phrases we shouldn't only look for referring expressions or quantifiers. In each individual case it will be tricky business to determine what the function of that noun phrase is on that occasions, just as it is with, say, generics.

3.7 Conclusion

In this chapter we have seen the following:

- The first step in the trivial arguments should be understood as making the transition to a focus construction. With the premise we communicate the same information as with the conclusion, but we do so in a different way. The conclusion focuses on a particular aspect of the information communicated.

- It is because of the above that the inferences are so obvious, not because speakers implicitly hold a certain substantial metaphysical picture of the world.

- This phenomenon can be seen as part of a more general picture of the function of noun phrases. It is plausible that their function is wider and more diverse than just to refer or to quantify.

- None of the above is a threat to the literal and objective truth of the utterances in which such non-referring non-quantificational noun phrases occur. All of the following are literally and objectively true

  (161) The tiger is fierce

  (162) The average American has 2.3 children.

  (163) Fido has the property of being a dog.

  but none of the noun phrases that occur in them have any direct ontological relevance.

In chapter 2 we have looked rather generally at the function quantification, and in this chapter we looked at the function of noun phrases. So far our results about ontology have been only rather limited. We have only seen that there is no cheap way to ontology through quantification or through noun phrases, even in literally and objectively true sentences. We have seen that these expressions also occur in uses that are not directly related to ontology, even though they, of course, are on other occasions. Referring noun phrases as well as external quantifiers have a direct relevance for ontology. But whether or not a particular occurrence of a noun phrase is referring, and whether or not a particular occurrence of a quantifier is external is often tricky business.

In the next two chapters we will have to deal with that tricky business in the case of our talk about properties, propositions and natural numbers. When we talk about them, both in everyday life and in theoretical enterprises, are we talking about some domain of entities, or are we doing something else? So far we have only seen that the answer isn't obvious even if such talk is literally and objectively true. Now we have to try to figure out what the answer is.

Chapter 4

Properties and Propositions

4.1 Introduction

In chapters 2 and 3 we have looked at quantification and noun phrases in general, and at how they relate to ontology. It is clear that both of them have a close connection to ontology in some of their uses. That is, there are uses of quantifiers and noun phrases such that if an utterance with one of them in it is literally and objectively true then certain ontological consequences follow. But we have also seen that not all of their uses have ontological relevance. There are uses of quantifiers that don't have any direct connection to ontology. And there are uses of noun phrases that don't have, or at least don't seem to have, them either. But so far we have not seen any positive account related to ontology, and we have not seen an account of what we do when we talk about properties, propositions and natural numbers. After all, it was talk about them that was supposed to be the focus of this dissertation. Does our ordinary or theoretical talk about them have ontological presuppositions for its literal and objective truth? What are we doing when we talk about them? Why do we talk about them in the first place?

We have to address these questions in this chapter. This will take us beyond merely arguing that the existence of properties and propositions isn't trivial to show, and towards a more positive view of what the function of such talk is. In the first three chapters of this dissertation we have dealt with rather general issues that are of relevance to many debates about ontology. In this chapter, however, we will focus on properties and propositions, and in the next chapter on natural numbers. Much of what has been said so far will also be relevant to understand what we do when we talk about events, facts, etc., but we will not get into this here. We will return to general issues about ontology in chapter 6.

4.2 Two views about (talk about) properties and propositions

To understand what we do when we talk about properties and propositions it will be helpful to ask why we talk about them in the first place. And here it will be helpful to first look at why we talk about them in ordinary everyday communication, communication that doesn't deal with ontological concerns directly, as some philosophical conversations do. Later, after we made some progress with this, we will have to look at why we talk about properties and propositions in more theoretical enterprises, like philosophy, or semantics. But to start, lets first ask: why do we talk about properties and propositions in ordinary, everyday life?

Do see what we are asking here, let's ask a different question: what are we doing when we talk about bagels, and why are we doing it? I take it that the answer to this question is quite straightforward, at least in its rough outlines. When we talk about bagels we talk about a certain kind of entity that is part of the world that we are part of, too. We talk about them because bagels have a certain importance in our life, and communicating information about them is helpful for us to do with bagels what we want to do with them. Bagels are of a special kind of entity, they are food. And we need food, and talk about food helps us to find it and eat it. Communicating information about food helps us get fed. That's why we talk about bagels, in rough outline. Of course, we also talk about bagels in more complicated ways. I just talked about bagels in a way that doesn't help you get fed. Poets talk about bagels in a way that doesn't help you get fed. Still, though, I think it is fair to say that these ways of talking about bagels are not why we talk about bagels in the first place. In the first place we talk about bagels in the way in which we communicate information about certain entities that have some importance in our lives, and in the second place we use them as examples in philosophy or poetry. So, roughly speaking, that's why we talk about bagels. Now, why do we talk about properties and propositions?

It seems clear that we don't talk about properties or propositions for a similar straightforward reason as the one why we talk about bagels. We don't have any similarly direct need to gain information about properties themselves. It rather seems that we talk about properties because we are really interested in the instances of these properties. We talk about bagelhood because we are interested in bagels, not because we are interested in bagelhood all by itself. It seems that the reason we talk about properties is more tricky and a bit less direct as it is with bagels.

To see why we talk about properties (I will focus on those rather than propositions for now) lets see what we do when we have to take recourse to such talk. Is it ever the case that we have to talk about properties in a ordinary situation of communication? Is it ever the case that we have a need to communicate certain information, but the only way we can do it is by taking recourse to talk about properties?

Yes, there are such circumstances. One such situation is a situation that is familiar from discussions in chapter 2. It is the case where we communicate in ignorance. Suppose, for example, that you know that both Nixon and Bush have something in common that Reagan doesn't share with them. If you remember what it is then you can just say

(164) Both Nixon and Bush like broccoli, but Reagan doesn't.

But if you can't remember what it is then you have to take recourse to property talk. If you only have incomplete information then you have to take recourse to talk about properties to communicate it. In this case you will have to utter something like

(165) There is a property that Nixon and Bush have in common, but Reagan doesn't have.

(166) There is something Nixon and Bush have in common, but Reagan doesn't have.

Talk about properties thus increases our expressive power. It allows us to say things that we can't say, in certain circumstances, without it. So, at least one of the reason why we take recourse to talk about properties is because it increases our expressive power in a way that we have to take recourse to to communicate certain states of incomplete information. But even if we take recourse to property talk because of the increased expressive power that comes with it, it is not clear how the increased expressive power arises. Is it from an internal, or an external, quantifier? Both internal and external quantifiers can give rise to increased expressive power. An internal quantifier does so, so to speak, directly, because it allows us to express certain infinite disjunctions and conjunctions. An external quantifier does so more indirectly. From a semantic point of view it only imposes a condition on the domain of objects it ranges over. What expressive power such a quantifier has will thus not only depend on its semantics, but also on what things there are in the domain it imposes a condition on. To illustrate this, consider an utterance of

(167) Every man likes beer.

This will imply the infinite conjunction

(168) Fred likes beer and John likes beer and ...

with one conjunct for every man, possibly adding names for unnamed men. But what disjuncts this disjunction contains will depend on which men are in the domain of discourse. Thus what (167) implies will depend on the size of the domain of discourse, and thus in a sense its expressive strength will depend on it. Not so in the internal reading of the quantifier. There expressive strength depends only on the language. What disjuncts are involved there is independent of the domain of discourse.

How should we understand an ordinary everyday utterance of (165)? Well, roughly there are two ways. There is one where the quantifier is read externally, and one where it is read internally. And these two ways can be seen as giving rise two two different views of what we do when we talk about properties in everyday life:

- **Externalism:** Talk about properties is talk about some domain of entities. Thus quantification over properties is external quantification over this domain of entities. And talk about properties with property nominalizations is reference to a certain entity.

- **Internalism:** Talk about properties isn't talk about some domain of entities. Rather talk about properties makes general, or otherwise complicated, claims about things that have properties. In particular, quantification over properties is internal quantification, and talk about them with noun phrases is non-referential.

These are two rough pictures of what we do when we talk about properties. According to externalism talk about properties will be a lot like talk about cars. They are things out there in the world and we make external claims about them. According to the internalist it will be quite different. Properties are not things out there in the world, and talk about them is quite different than talk about cars. When we talk about properties we rather make general claims about the things that have the properties, rather than claims about some other things. Of course, this is only a rough characterization of these views. And there are other views that fall, so to speak, in between. One of them is to say that properties are things out there, but they are language dependent entities, things that are closely related to the predicates we use. Such a view will have some similarities with externalism and some with internalism. One such view will be discussed in chapter 6. Similarly the view that we only pretend that properties exist and then quantify over them externally within that pretense. It also will be discussed in chapter 6.

Given certain assumptions, internalists and externalists will disagree on whether or not properties exist. These assumptions are that

- whatever properties are, they are the things that expressions like 'being a dog' stand for, and
- (certain) talk about properties is literally and objectively true.

If we assume this, and also assume that (165) is one of the true statements about properties, then externalism will imply that properties exist. In addition, if properties exist then externalism will be true. Since properties are nothing more that whatever expressions like 'being a dog' stand for or refer to, then these expressions indeed have to be referring expressions that succeed in referring. Thus externalism is true. Thus on can defend externalism by arguing that properties exist, assuming the above assumptions. This gives rise to another set of apparently easy arguments for externalism and against internalism.

Before I will get into a more positive account of what the function of talk about properties and propositions is, I would like to discuss some of these apparently easy or direct arguments for externalism. They either concern directly the existence of properties, or that internalism makes a too close connection between language and properties. All of them are mistaken. After we have discussed the apparently easy refutations of internalism I will propose that internalism is correct, try to make this plausible and show that externalism has a number of problems as an account of the function of our talk about properties and propositions.

4.3 Further arguments against internalism

We started this dissertation with some easy arguments for the claim that properties, propositions and numbers exist. We have by now seen that these arguments indeed were too easy and didn't really have any ontological conclusions. But there are also some other arguments, that don't have the form of the above trivial arguments, to a similar conclusion. We will look at these in this section.

4.3.1 Aboutness

One might think that it is rather trivial to see that phrases like "being a dog" or "that John kissed Mary" are referential phrases since they are about something. The first is about a certain property, the latter about a certain proposition or fact. So, they refer to these entities, and thus they are referential expressions.

I don't think this is right. Aboutness can't be used as a test for referentiality, nor as a test for existence. The first is clear from the fact that descriptions are about something, too. "The vase I broke" can be about a certain object, and a speaker uttering a sentence in which this phrase occurs will be talking about a certain vase.

Furthermore, one can speak about things that don't exist. Consider:

(169) Fred and I spent hours talking about aliens after watching the X-files.

This can be true whether or not aliens exist. Aboutness doesn't settle any interesting issue about the function of a noun phrase.

"about" has at least two uses:

1. It describes the topic of a conversation, as in (169). This is a rather general use of "about". In this sense one can talk about Santa Claus, objects of any category, like surfing, Kennedy's assassination etc..

2. It describes the referent or denotation of a phrase. In this stricter sense one can only talk about things that exist. In this sense one can truly say things like

   (170) All these debates about Homer's intentions, his life etc.. It is amazing. But one day people will realize that there never was a Homer and that all these debates will have been about nobody.

Thus we have to distinguish **topical aboutness** and **referential** or **denotational aboutness**. Noun phrases are only trivially about something in the topical sense. Only some noun phrases are about something also in the referential/denotational sense.

4.3.2 Explanation of truth

A further reasoning is that property nominalizations have to be referential because otherwise we would not be able to explain the truth of the sentences in which they occur. Take a sentence like

(171) Humility is a virtue.

How can we explain its truth? Well, so the story goes, because "humility" refers to a certain property, and "is a virtue" expresses another (second order) property, and there is a relation (exemplification) that holds between the first and the second. More generally,

(172) The truth of a sentence of the form "F(a)" can be explained by saying that "a" refers to for some object, and "F" expresses some property and that this object exemplifies that property.

But for that explanation to be true in a case where "a" is a nominalization, the nominalization has to refer to a certain object which has to exist.

I doubt it. It seems to me that such explanations are vacuous. What explains that "Fido is a dog" is true? Not that "Fido" refers to Fido and "is a dog" expresses doghood, and Fido instantiates doghood. What explains it is first that "Fido is a dog" means that Fido is a dog, and secondly that Fido has certain features that one has to have in order to be a dog. So, such an explanation comes in two parts. "Fido is a dog" is true because of what it means, and secondly because of how the world is. When you ask: Why is "Fido is a dog" true? then one could either ask: Why does it mean what it means? or: Why is the world as it is? The first one, which is usually not the one that is intended, has to do with conventions, communicative intentions and the like. The second one comes down to: Why is Fido a dog? And to answer this is a job for biology. Fido is a dog because he has such and such genes, or whatever else is constitutive of doghood.

(172) is close to being true, but not explanatory. That is, the following is true, even though (172) isn't.

(173) A sentence of the form “F(a)” is true iff “a” stands for something, “F” expresses some property, and what “a” stands for has that property.

But this is not a substantial explanatory story about truth. As we will see, talk about properties arises naturally when we try to make very general claims. And it is the reason why we take recourse to this talk on these occasions. We don’t talk about properties to talk about a certain class of objects, whatever they may be, but to make certain general claims. I think what is said with (173) is true, but does not imply the existence of properties nor that property nominalizations refer. In particular, it has no explanatory force.

4.3.3 Paraphrases

A closely related issue is the unavailability of paraphrases for certain sentences that contain noun phrases like “humility”. Paraphrases have played quite an important role in the discussions about ontology, and do so up to this day. Originally paraphrase was a major tool for nominalists. They said that we don’t have to accept that

(174) Fido has the property of being a dog.

commits us to properties, because (174) can be paraphrased away as

(175) Fido is a dog.

It was pointed out by Alston in (Alston 1958) that this doesn’t work, in particular if one views ontological commitment as what is implied to exist by what one believes or accepts. Simply because paraphrases have to be truth conditionally equivalent, and thus either one implies the other. So, if (174) implies that properties exist then (175) does so, too, because it implies (174). So, paraphrases have the same ontological commitments. And this maneuver is often used to argue that a sentence as innocent as (175) implies that properties exist. Furthermore, there are sentences that apparently talk about properties and the like that seem to have no paraphrase (see (Jackson 1977)). In any case, they imply that properties exist.

I see things somewhat differently. Paraphrases have nothing directly to do with ontology. The real issues are not whether or not every true sentence that contains a property nominalization has a paraphrase without one, but whether these nominalizations have the function of referring to an object, or some other semantic function. But paraphrases still have something to do with the issue in question. When we take a true sentence with a nominalization of a certain kind and we ask why it is true, then we will take recourse to something like paraphrases in an account of why it is true. As I said above, an explanation of truth will in its main part down to an explanation of why the world (in a certain respect) is as it is. Consider

(176) Heavy drinking is a vice.

It is true. But why? The answer isn't because the property of being a heavy drinker exemplifies the second order property of being a vice. But it also isn't a brute fact. Heavy drinking is a vice because anyone who drinks heavily is likely to do some damage to himself or other people. That of course does not mean that "heavy drinking is a vice" can be paraphrased away as "Anyone who drinks heavily is likely to do some damage to himself and other people". But we naturally take recourse to sentences like the second to explain the truth of the first. And because our explanations go this way many people, rightly, I think, take this as evidence that really in the first we are not talking about some object or other that is referred to by "heavy drinking".

Compare this with a case where we have a referring noun phrase in the subject position, as in

(177) Fred is wealthy.

If we try to explain why it is the case we do it differently than with (176). Here we will talk about Fred and how things are with him.

It is a mistake to believe that sentences like (176) are either about some objects other than heavy drinkers, or can be paraphrased into a sentence with quantifiers ranging over heavy drinkers, or noun phrases referring to heavy drinkers. Something like this is behind the belief that if a noun phrase / verb phrase sentence isn't paraphraseable into a sentence that only refers to or quantifies over certain objects then it has to be about different objects. But that is a mistake. A sentence can make a complicated claim about one kind of object in a subject predicate form which can't be made with just referring or quantificational noun phrases. Think of generics or the average construction as guides into the right direction. Now, in their case we can say what they say in other ways, too, when we use ratios, or talk about what usually, or all things being equal, is the case. But that doesn't have to be so. We can imagine a language with generics in it, but without "ceteris paribus" or "all things being equal" or the like. Still, generics would make complicated claims about certain individuals, but the other resources of the language might not allow to express the same truth just with reference and quantification. If a sentence isn't paraphraseable it doesn't mean it talks about different or new entities.

4.4 Inexpressible properties and propositions

4.4.1 The problem

The most common and the most plausible argument against internalism is the argument that the internalist has to make the obvious mistake of claiming that there are no inexpressible properties and propositions. It is so important for our present discussion, and so much more involved to deal with then the above ones, that it deserves its own section. This objection to internalism arises from the following consideration. If internalism were right then quantification over properties and propositions would be equivalent to infinite disjunctions and conjunctions of sentences *in our language*. After all, that is what internalism is all about, that such talk is not about some external, language independent, entities, but rather a language internal source of increased expressive power. But it is a mistake, so the argument continues, to make such a close connection between our language and properties and propositions. After all, there are properties that are inexpressible in our language. Some properties can't be expressed in English. And sometimes we essentially take recourse to such properties when we make quantified statements. Sometimes we say

(178) There is a property such that Φ.

but the only properties that are Φ are ones that are not expressible in English. So, there will be no true disjunct among the infinitely many disjuncts in the analysis of (178). Even though (178) is true the internalist understanding of it will have to come out false. To sum it up, internalism mistakenly ties properties too close to predicates. An internalist will subscribe to the metaphor that properties are mere shadows of predicates. They are not language independent entities. But this is false, because there are inexpressible properties. Or so the argument goes.

I think this argument is mistaken. Properties are shadows of predicates, but this is not incompatible with there being inexpressible properties. And the truth conditions of quantification over properties is to be understood as being equivalent to infinite disjunctions and conjunctions of certain sentences of our language, but this, too, is compatible with there being inexpressible properties. In this section I will spell out how that can be so. I will show how inexpressible properties fit into an internalist framework. I will show how the above argument isn't a problem for internalism. But I won't do this by extending our language in any way with further predicates, or the like. It is perfectly consistent to claim that quantification over properties is merely a device that increases our expressive power in a "logical" way that is metaphysically thin, and to hold that there are inexpressible properties. Let me explain.

4.4.2 Inexpressible properties

I take it that we all believe that there are properties inexpressible in English. I don't want to challenge this, but I would like to point out that it is prima facie not so clear why we accept this. After all:

1. For a property to be inexpressible in a language means that no predicate (however complex) expresses it. Simply because there is no single word in a certain language for a certain property doesn't mean it isn't expressible in that language.

2. We can't be persuaded that there are properties inexpressible in English by example. One can't say in English without contradiction that the property of being Φ isn't expressible in English.

So, why again do we believe that there are inexpressible properties?

The best reason seems to me to be the following: Even though we can't give an example of an inexpressible property of English, we can give examples of properties not expressible in older, apparently weaker languages. For example, the property of tasting better than Diet Pepsi is not expressible in Ancient Greek. So, there are properties expressible in English, but not in Ancient Greek. In addition, we have no reason to believe that English is the final word when it comes to expressing properties. We can expect that future languages will have the same relation to English that English has to Ancient Greek. Thus, we can expect that there are properties inexpressible in English, but expressible in future languages. In short, there are properties not expressible in present day English.

Let's call this argument for inexpressible properties **the inductive argument**. It is a powerful argument. Can we accept the inductive argument, but still believe in an internalist view of talk about properties?

4.4.3 Internalism and inexpressible properties

What does it mean for a property to be expressible in English? Well, that there is a predicate of English that expresses it. But that could mean two things. One the one hand, it could mean that there is a predicate of English that expresses this property in the language English. On the other hand, it could mean that there is a predicate of English such that a speaker of English would express this property with an utterance of this predicate. Which one of these we take will make a difference for the issue under discussion. To illustrate the difference, consider:

(179) being that guy's brother

This predicate does not express a property simpliciter, it only expresses one on a particular occasion of an utterance of it, that is, in a particular context. In different contexts of utterance it will express different properties. However,

(180) being Fred's brother

expresses a property independent of particular utterances, or better, expresses the same one in each utterance.[1] If "that guy" in an utterance of (179) refers to Fred then this utterance of (179) will express the same property as any utterance of (180) will. However, there might be contexts in which an utterance of (179) will express a property that can't be expressed by an "eternal" predicate like (180).

So, when we ask whether or not a property P is expressible in a language L we could either ask

1. whether or not there is a predicate Φ such that in every context C, an utterance of Φ (by a speaker of L) in C expresses P, or

[1] I assume that "Fred" and "brother" are disambiguated, i.e. with respect to whether Fred is a monk or a sibling, and whether it's Fred Dretske, Fred Astaire, Fred Flintstone, or any other Fred.

2. whether or not there is a predicate Φ and a context C such that an utterance of Φ (by a speaker of L) in C expresses P.

Let's call expressible in the first sense **language expressible** and expressible in the second sense **loosely speaker expressible**. The latter is called loosely speaker expressible because it only requires for there to be a context such that an utterance of Φ in that context by a speaker of L would express P. Any context is allowed here, whether or not speakers of that language actually ever are in such contexts. We can distinguish this from what is **factually speaker expressible**. Here we allow only contexts that speakers of that language actually are in.[2] Let me illustrate.

Ancient Greek does not allow for the expression of the property of tasting better than Diet Pepsi in the sense of being language expressible. We can assume that. However, it seems that it is expressible in Ancient Greek in the sense of being loosely speaker expressible. In the context where there is Diet Pepsi right in front of a speaker of Ancient Greek he could simple utter the Ancient Greek equivalent of

(181) tasting better than this.

while demonstratively referring to Diet Pepsi. But since there was no Diet Pepsi around during the time when Ancient Greek was a living language, this context is not allowed when considering the question whether or not this property is factually speaker expressible. In this case, it seems that the property is not factually speaker expressible in Ancient Greek, just as well as it is not language expressible in Ancient Greek.

Being language expressible implies being factually speaker expressible, which implies being loosely speaker expressible, and none of these implications can be reversed (or so we can concede for now).

What all this shows is that both the inductive argument and the above account of the internalist view of talk about properties was too simplistic. In the latter it was simply assumed that contextual contributions to content do not occur and that the truth conditions of talk about properties can simply be given by infinite disjunctions and conjunctions of eternal sentences of the language in question. But that's not always so. Sometimes predicates express properties in some contexts that can't be expressed with eternal predicates.

[2] To simplify, we consider someone only as a speaker of their native language. This is also implicitly assumed in the inductive argument.

To say this is not to deny that properties are shadows of predicates, just that they are shadows of eternal predicates. But still, how can an internalist deal with this situation? How can they accommodate contextual contributions to what is expressed by a predicate? This seems to be a problem because of the following.

Suppose that the only property of beer that interests Fred is that it tastes better than Diet Pepsi. This seems to be incompatible with the view that talk about properties in Ancient Greek is internalist. The only way this property is expressible in Ancient Greek is as being loosely speaker expressible. But if the truth conditions of the Greek equivalent of

(182) There is a property of beer that interests Fred.

are supposed to be some infinitary disjunction of sentences in Ancient Greek then this does not seem to work. After all, the Greek equivalent of "Fred interests about beer that it tastes better than Diet Pepsi" won't be one of the disjuncts, since it is not language expressible in Ancient Greek. Also, the Greek equivalent of "Fred interests in beer that it tastes better than this" (whereby "this" refers to Diet Pepsi) won't be among the disjuncts either. If it were then the speaker of (182) would have to have referential intentions that fix the referent of "this", just like any speaker using a demonstrative would have to. However:

a) An ordinary speaker who utters (182) certainly won't have such intentions. After all, they might not know what this property is.

b) Such speakers won't be able to refer to Diet Pepsi demonstratively, simply because there was no Diet Pepsi in Ancient Greece.

So, is the internalist view refuted? No. Here's why.

4.4.4 The solution

We need to have a witness for every property that can be expressed by predicates containing demonstratives. Since we have to allow arbitrary contexts here we have to allow demonstrative reference to any object whatsoever. In addition, for any object there will be a context such that if a predicate with a demonstrative were to be uttered in that context then the demonstrative would refer to that object. The infinite disjunctions (and conjunctions) that model the truth conditions of property talk have to have a witness for every one of these properties. If there is a property such that Φ then there has to be a predicate with demonstratives in it and a context such that the predicate expresses that property when uttered in that context. The only contribution of the context that we have to consider here is what objects the demonstratives refer to. Thus overall, for every property P there have to be objects $o_1, \ldots o_n$ and a predicate with demonstratives $d_1 \ldots d_n$ in it such that the predicate expresses the property P, given that d_i refers to o_i, for all $1 \leq i \leq n$. This naturally gives rise to the following.

Let's assume that the truth conditions of a fragment of a natural language without talk about properties is correctly modeled with a certain formal language L. Adding talk about properties to that language should give us an infinitary expansion of L, according to the internalist. Now, to accommodate demonstratives, do the following. Add infinitely many new variables to L, which model the demonstratives. Build up formulas as usual, but don't allow ordinary quantifiers to bind these new variables. To accommodate talk about properties, we represent it as an infinite disjunction or conjunction as before, with one difference. Whenever we form an infinite disjunction or conjunction we also existentially or universally (respectively) bind all these new variables. Thus, now we do not simply represent "there is a property such that Φ" as the infinite disjunction over all the instances "$\Phi(P)$", i.e. as "$\bigvee_P \Phi(P)$". Now we take this disjunction and add existential quantification on the outside binding all the new variables. So, we now represent "there is a property such that Φ" as "$\exists v_1, v_2, \ldots \bigvee_P \Phi(P)$".[3]

The new free variables play the role of the demonstratives in this account, and the quantifier binding them plays the role of the arbitrary contexts that we allow in loose speaker expressibility. For example, the disjunction that spells out the truth conditions of (182) will contain a disjunct corresponding to

(183) Fred's interest in beer is that it tastes better than v_i.

Now there will be an existential quantifier that binds v_i from the outside. Since it will range over Diet Pepsis this disjunct will be true, and thus the disjunction will be true. And this will be so independently of there being a referring expression that refers to Diet Pepsi in the language in question

[3] It is a further, substantial, issue whether or not this quantifier has to be understood a-temporally, in other words as "there is, was or will be $v_1, v_2, \ldots$", and a-modally, in other words as "there is or could be $v_1, v_2, \ldots$. These are issues that have to be addressed (in the relevant form) by everyone in the debate, namely whether or not what properties there are depends on time or modality.

However, there is no upper bound on how many of these new variables will occur in these disjunctions. Since we allow, and have to allow, every predicate to occur in the disjunction, we can't give an upper bound on how many demonstratives may occur in these predicates. So, in the infinite disjunction there will be infinitely many variables that have to be bound, all at once, from the outside.

But this can be done. We just have to go to a higher infinitary logic. Not only do we need infinite disjunctions and conjunctions, we need quantification over infinitely many variables. Before we only used a small fragment of $L_{\omega_1,\omega}$, now we use a small fragment of L_{ω_1,ω_1}. This latter logic also allows for quantification over countably many variables.[4]

Given this new model of talk about properties we have the following:

- Properties are shadows of predicates, but not shadows of eternal predicates.
- Talk about properties gives rise to an infinitary extension of the original language, but not just to a small fragment of $L_{\omega_1,\omega}$, but to a small fragment of L_{ω_1,ω_1}.

So, the property of tasting better than Diet Pepsi is not expressible in Ancient Greek in the sense of language expressible nor in the sense of factually speaker expressible. It is however, expressible in Ancient Greek in the sense of loosely speaker expressible. And by the inductive argument we get that we have reason to believe that there are properties that are not expressible in English, but we get that only when expressible is understood in the sense of either language expressible or factually speaker expressible. Quantification over properties has to be understood as being based on what is loosely speaker expressible with predicates. Therefore it will be true that

(184) There are properties that are not expressible in English.

if expressible is understood as being language expressible or factually speaker expressible, but false if it is understood as being loosely speaker expressible.

[4] $L_{\omega_1,\omega}$ is an infinitary logic that allows conjunctions and disjunctions over countable sets of formulas, but only quantification over finite sets of variables (as in regular first or higher order logic). L_{ω_1,ω_1} allows for both conjunctions and disjunctions over countable sets of formulas, plus quantification over countable sets of variables. The basic language is usually the one of first order logic, but one can define infinitary expansions of other languages just as well. See (Keisler 1971) and (Barwise 1975) for much more on this. In our case here we only use very small fragments of these logics. All these fragments will be finitely representable, for example, and smaller than the smallest fragments studied in (Barwise 1975) or (Keisler 1971). I will elaborate on this in appendix A.

The internalist view is not threatened by there being inexpressible properties. That, of course, is not an argument in it's favor. It does, however, show that this simple refutation of the internalist view is too quick.[5]

So, the inexpressibility worries have shown that it was too simplistic to give the truth conditions of quantified statements as merely an infinitary expansion of the original language in the sense in which we only allow for new conjunctions and disjunctions. We will also allow for new, infinitary, quantifiers. This wasn't of any special concern in chapter 2 because we only dealt with a very simple situation. But now it is of importance. In fact, below we will consider another complication to the formal picture due to the threat of semantic paradoxes that comes with quantification over properties and propositions in general. But before we get into this, let's see if we answered the expressibility worries sufficiently.

4.4.5 More trouble?

So far we have only looked at the inductive argument for believing that there are inexpressible properties and propositions. But are there not other arguments? And are there not other versions of the inductive argument to the conclusion that there are properties not even loosely speaker expressible in English?

I haven't shown, of course, that there are no other arguments for there being properties that are not even loosely speaker expressible, nor that there are no properties that aren't loosely speaker expressible. But I think that once it is clear that the simple inductive argument can be endorsed by an internalist, the internalist is in a pretty strong position. After all, simply pointing to the intuition that there are properties that aren't loosely speaker expressible doesn't do much good. Being loosely speaker expressible means being expressible by some predicate (however complex) in some context. Who can claim to have any clear intuitions about what's expressible with these resources?

So, we have to see what reason one might have for believing that there are properties that aren't loosely speaker expressible. And, since this again can't be motivated by giving an

[5] Similar considerations apply to inexpressible propositions, and in particular to issues about the increased expressive power that a truth predicate (in combination with talk about propositions) gives rise to. For example, both (Horwich 1990) and (Field 1994a) reject a simple account of the function of the truth predicate as giving rise to an infinitary expansion of the language without it because of expressibility worries. With them out of the way a much more plausible minimalist theory of truth can be formulated. We will get back to this in chapter 6, section 6.5.

example of such a property, the only way to go will be a version of the inductive argument, but this time an inductive argument for there being properties that are not loosely speaker expressible. And to start such an argument we have to point to a property that is (loosely speaker) expressible in English, but not loosely speaker expressible in, say, Ancient Greek. What could that be? My experience is that those who are skeptical about the conclusions I would like to draw take recourse to properties from science on this occasion. For example,

(185) being a quark

One might think that it isn't even loosely speaker expressible in Ancient Greek. Whatever one's prima facie intuitions about this are, we should note that since this property is (language) expressible in English there has to be some relevant difference between English and Ancient Greek that allows for this property being language expressible in English. So, how did we come to be able to express it? That certainly is a tricky question, relating to some serious issues in the philosophy of science. Two possibilities come to mind, though, namely:

- "being a quark" is a theoretical predicate of physics. It is at least in part implicitly defined by the physical theory that uses it. Thus we can express it because we have the theory.

- We can express the property of being a quark because we have been in contact with observable phenomena that are caused by quarks, like traces they leave on some measuring instrument.

I would like to point out that if either one of them is correct then there is not problem for the internalist. The reason is simply the following. If the first is correct then the problem of expressing the property of being a quark reduces to expressing the defining parts of the theory that implicitly defines "being a quark", plus making the implicit definition explicit. Thus the problem is pushed back and doesn't really have anything essentially to do with "being a quark". If the second possibility is the right one then the increased expressive power does come from being in contact with more objects. Such contact does not have to be only with mid-size object, or direct contact.

Thus if we think about possible counterexamples to the claim that every property is loosely speaker expressible we run into a bit of a dilemma. If we reflect on how we came to be able to express a certain property that allegedly isn't loosely speaker expressible in Ancient Greek we see that and how it is loosely speaker expressible after all.

I conclude that there are no known counterexamples to the claim that every property is loosely speaker expressible.

4.4.6 The expressibility hypothesis

The above view about expressibility that is implicit in this discussion is, I think, our best shot at understanding how expressive power arises, where natural languages differ in expressive power and where they don't differ. I will thus conclude the section dealing with the alleged problem that inexpressible properties are for internalism with offering a bit of an account of expressibility in general.

I think the following expressibility hypothesis is plausible and even though it seems prima facie problematic, it is our best bet on how expressibility arises and how it differs from languages to languages. In addition, there are no known counterexamples to it. Here it is:

(EH) **The Expressibility Hypothesis**: For every property and every natural language, there is a predicate in that language that would express that property if uttered in the right context (similarly for propositions).

This hypothesis has to be distinguished from a weaker one that has been endorsed by Searle.[6] Searle's hypothesis is that

H1 For every proposition p, if you can think that p then you can say that p.

Or in other words, the content of any thought can be articulated in language. The present expressibility hypothesis is stronger because it states that

H2 For every proposition p, you can say that p if you are in the right context.

In other words, you can articulate any proposition whatsoever, in the right context.

How can that be right? Maybe we do not have any counterexamples to this claim, but is there any positive reason to believe it?

[6] See his (Searle 1969).

I think there is, because it fits nicely into a picture about where expressive power of languages comes from that is as good as any I have seen. According to this picture expressibility arises on the one hand from innate conceptual resources that are shared among humans and that are expressed in all human languages.[7] Besides these the basic source of expressive power is the particular situations that speakers have available to make statements about. So, their having access to certain contexts allows them to express certain propositions in these contexts. A third source of expressibility is now to separate what can be expressed only in certain contexts from these contexts. To do this the language community will introduce terms that allow the speakers of that language to express in any context what they originally only could express in certain contexts. As an example of this, consider names. With demonstratives and Fred right in front of me I can express the proposition that Fred is tall. But without having a name for Fred in my language I can only express this proposition when I have a certain relation to Fred, when he is within my visual field, say. But that is a real restriction that I might want to overcome, in particular if I have a need to talk about Fred behind his back. Thus I can introduce a name for Fred, originally taking recourse to the demonstrative and the context where Fred is in front of me, and then I have the ability to say what before I could only say when I had a certain relation to Fred now without having that relation. In a sense I have thus increased my expressive power. I have increased it in the sense that I can now say more in certain contexts, like the ones where Fred isn't around. I have increased my expressive power in the sense that I have made myself expressively more independent from being in a certain situation. But I haven't increased it in principle, in the sense that I can now say things that I couldn't have said before. Having the word 'Fred' in my language doesn't do that. I could have said what I can say with it without it, if I am in a certain context.

So, expressive power has three sources:

- Innate concepts.
- Context.
- Separating expressive resources from the requirements of a particular context.

[7] Of course, very very primitive languages might not have this feature. I will restrict myself to all mature human languages here. Those are the ones that articulate all innate concepts. I take Ancient Greek and English to be among them.

And this, generalized, gives rise to the expressibility hypothesis, and is what is underlying the internalists accommodation of inexpressible properties. According to this view the difference in expressive power of natural languages is one of what these languages can express in separation of particular contexts. And this will arise from what different objects the speakers of these languages have access to (Fred, Diet Pepsi, etc.) and which ones they take to be important enough to introduce a context insensitive word for them. But the difference of expressive power in this way arises only from different contexts the members of the language community have access to. In arbitrary contexts they will be able to express the same. The apparent difference in expressive power is real, but it is of a kind that isn't of much philosophical relevance for our concerns about the relation between predicates and properties. Different natural languages differ in what is language expressible with them, but not in what speakers can express with them (in the sense of being loosely speaker expressible). Internalism combined with an account of the role that demonstratives have, as spelled out above, nicely fits in with this view of expressibility.

To be sure, I have not shown that this view is correct. I have neither shown that there are innate concepts, nor that they are the same for all people etc.. But even though I haven't shown that this is the right view, it seems puzzling to me what else could be right. Unless one wants to endorse a kind of a holistic view of content, which I find implausible, I think there is basically no other way. At least I think I can claim that much: this is a plausible account of the source of expressive power of natural languages, and there are no known counterexamples to this view.

4.5 An account of the function of talk about properties and propositions

Internalists about (talk about) properties believe that quantification over them in ordinary, everyday situations of communication is internal quantification, and that in such uses property nominalizations are not used referentially. Nonetheless, according to internalism, talk about properties and propositions will often be literally and objectively true. This view seems to be easily refuted in a number of ways, but we have seen that such simple refutations are based on misunderstandings, either a too simplistic picture of the workings of natural languages, or on philosophical mistakes. I of course didn't get you through all this if I would have more in mind with internalism than just propose it as a possible view that isn't easily refuted.

In this section I would like to propose and defend that internalism is true. Internalism gives rise to the best and most plausible account of what we are doing in everyday life when we talk about properties and propositions. It gives rise to the best account I know of why we talk about them in the first place. To be sure, the account I will offer will have some gaps. Talk about properties and propositions is so diverse that there is no simple explanation why we do it. Such talk contains a variety of different cases, some of them are very tricky and relate to large scale issues in other disciplines, like the philosophy of mind. I, of course, won't have a sweeping account of all of them. I think the account I will propose will be helpful in these areas, too, but only in ways to be determined later and on another occasion.

4.5.1 Why we do it

The main idea of internalism is that the function of talk about properties and propositions is to serve the need for increased expressive power that we have on certain occasions when we want to make statements expressing complex facts about ordinary things. However, these complex facts about ordinary things don't involve talk about some other entities over and above the ordinary ones. Several of the above considerations, in particular about quantifiers and the function certain quantifier free expressions, suggest that the way in which we achieve this is based on general features of natural languages, features that haven't been very prominent in philosophical debates. When we look at what we are trying to do when we utter certain sentences that are about properties and propositions we will see that they are used very much like certain expressions that we encountered in chapters 2 and 3. I will also argue below that internalism better accounts for how we achieve certain increased expressive power than externalist attempts that try to do the same.

Talk about properties and propositions has several different functions, all of which nicely fit into an internalist picture. Among the functions of such talk are the following:

- **Having expressive power through quantification.** Quantification over properties and propositions gives rise to increased expressive power, as we saw above. We have to take recourse to such quantification in case of communicating in ignorance and in case we want to make certain general claims. I will defend in a moment that such quantification is better understood as internal, rather than external, quantification.

- **Expressing generalities related to default reasoning.** Talk about properties is also used to express certain general claims without quantifiers. On such occasions this is very much like what can be said using generics. Consider:

  (186) Being a philosopher is fun.

  (187) Finishing a dissertation is stressful.

  It would be a mistake to think that on these occasions I am using the noun phrase to refer to some entity and claim of it that it is fun or stressful. I am rather expressing a ceteris paribus generalization about individual people. These are ceteris paribus because the truth of either one of them allows for exceptions, and it is not clear for which ones or how many. Still, they figure in valid (in the right sense of the word) default inferences, as

  (187) Finishing a dissertation is stressful.

  (188) Fred is finishing his dissertation.

  (189) Thus: Fred is stressed out.

  Again, this does not strictly speaking follow, since not everyone has to be stressed out when finishing their dissertation for it to be true that finishing a dissertation is stressful. But by default, without overriding additional information, it can be inferred that Fred is stressed out when he is finishing his dissertation.

- **Articulating information with a certain focus structure.** As we have seen in chapter 3, property nominalizations and that-clauses on occasion have the function to communicate certain information, that can also be communicated without them, in a certain way, with a certain stress. This is an important feature they have and not to be ignored.

- **Ascribing content.** That-clauses have the function of ascribing content to mental states, speech acts, and other actions. This is a common focus of the philosophical debate, but it is the one I have the least to say about. Undoubtedly, this is what we do with that-clauses, and below I will argue that our doing this is better compatible with internalism than with externalism. The tricky question, however, is why we do this in the first place. I wanted to offer an account of not only what we do when we talk about properties and propositions, but also why we do it. It is not clear to me why we ascribe content to mental states, speech acts and other actions. The answer: "because they do have content" is not satisfactory, because lots of things have lots of properties that we never bother to ascribe to them. The answer to this question, why we do it, won't be easy, and I will have nothing to contribute to it here.

4.5.2 How we do it

This is what we want talk about properties and propositions for. We still have to ask ourselves how we manage to do with such talk what we want to do with it. In particular, we will have to see whether we achieve this through talking about some domain of entities, or through some other means. The internalist account that I am putting forward here accounts for how we achieve these individual functions as follows:

- We gain expressive power in exactly the way we want it through using internal quantifiers that express language internal generalities. Such generalities are independent of what entities are in a particular domain of discourse, which is, as we will see, exactly what we want when we use quantification over properties and propositions. Internal and external quantifiers are part of our language independently of any issues surrounding properties and propositions. But with them ranging over properties and propositions, too, we gain even more expressive power, and on occasion we have to take recourse to it.

- We gain the expressive power that allows us to express generalizations from default reasoning in the same way in which we gain it in the case of generics. Not because we make a precise claim about some extra entity, but because we represent certain information that encodes a inference rule in default reasoning in a simple subject predicate form.

- We achieve the focus effects we want by putting certain parts of the information we communicate into a noun phrase position. Through that syntactically distinguished position it gains special attention.

- We ascribe content through specifying the content in the content sentence of a that-clause. Unsatisfactory as this may seem, it points, I think, to an advantage that an internalist has over an externalist when it comes to the ascription of content. We will get to this below.

I take it that it is clear by now that all this can be done without the quantifiers being external, or the noun phrases being referential. The question now has to be: why should we think that this is true of our uses of property talk? Why not think that even though this can be done as the internalist would like it, it in fact is done differently?

In the next section I will discuss a number of arguments against an externalist view, in particular against the view that that-clauses and property nominalizations are referential expressions. In this section I'd like to mention only one more general issue in favor of an internalist view over an externalist one. It is the apparent lack of ontological presuppositions that we ordinarily recognize for the literal and objective truth of our ordinary claims about properties and propositions. This will get us back to the second puzzle we started out with. Let me explain.

Talk about properties and propositions is surprisingly removed from ontological concerns. It never, in everyday life, is a concern or even a legitimate worry, whether or not there exists a particular property, or proposition. This is quite different from when we undoubtedly talk about certain entities, as when we use proper names. There it always makes sense to say

(190) To be sure, if Fred doesn't exist then it isn't true that Fred won the championship.

This doesn't seem to have an analogue with talk about propositions. It never is an issue whether or not the proposition that Fred is tall exists. It never is a consideration to wonder:

(191) To be sure, if that Fred is tall doesn't exist then it isn't true that Joe believes that Fred is tall.

Why is that so? One might think it is so because we all implicitly assume that propositions and properties exist, and that every that-clause and property nominalization stands for one. But that would suggest that questions about the existence of properties and propositions are just trivial, not silly. But in fact they are considered very, very awkward by ordinary non-philosophers. When you ask an ordinary speaker what kind of thing being more than 5 feet tall is, whether or not it is located in space, whether or not it can be at more than one place at the same time, etc., then this will be taken to be a quite strange and will not meet much sympathy for the alleged seriousness of this question. These are truly philosopher's questions. The account for this awkwardness isn't that properties and propositions exist so obviously that any questions about them are silly, but that we never in the first place talk about any entities when we talk about properties and propositions. Sure enough, given a simple picture of language it seems like that talk about properties and propositions is just like talk about cars. But we have seen that independent of any philosophical issues there is good reason to believe that this simple picture is false, and that quantification and noun phrases have much more complicated functions than merely to externally range over a domain of entities, or to refer. Internalism makes perfect sense. We have seen that the arguments that seem to refute it directly are mistaken, that internalism shows how we can do with talk about properties and propositions what we want to do with such talk, and that internalism accounts for our attitude towards ontological questions about properties and propositions.

To be sure, the truth of internalism doesn't show that properties and propositions don't exist. It only shows that if properties are nothing more than what 'being a dog' and the like stand for in ordinary communication, and if propositions are nothing more than what that-clauses stand for in ordinary communication, then properties and propositions don't exist, since these expressions have a quite different function in ordinary communication than to stand for some entity. But one might think that properties and propositions should be understood differently. We will return to this shortly. Before that, let's look at some arguments against externalism.

4.6 Arguments against externalism

So far we have only seen that the arguments against internalism aren't any good. That made internalism a possible position, and I think once one sees that internalism is possible things plausibly fall into place, and our talk about properties and propositions then makes perfect sense. However, there are also a few arguments against externalism. They rely on that talk about properties and propositions is in important respects different from uncontroversial cases of talk about entities. These arguments give additional plausibility to internalism.

4.6.1 Some quick arguments

Let me mention some aspects in which property nominalizations and that-clauses are different from uncontroversially referring expressions. I will not get into the details of these arguments and into the possible replies that an externalist could come up with. The issues mentioned below quickly turn into lengthy debates. I only mention them to point to apparent problems for externalism.

Category mistakes

When I say

(192) Joe is more fun than Fred.

then I am referring to two things and say of the first that it is more fun than the second. However, when I say

(193) Being a philosopher is more fun than being an accountant.

then I don't do this. It would be a category mistake to think that the property of being a philosopher is more fun than some other property. A property by itself isn't fun. What's fun is to have it. So, when I say (193) then I am saying something like that all things being equal, it's more fun to be a philosopher than to be an accountant.

That-clauses and their content sentences

That-clauses consist of the expression 'that' followed by a sentence that specifies what the content of the proposition is that the that-clause is about. That is an important difference between that-clauses and referring expressions. The that-clause does not simply stand for something that has a certain content, or truth conditions, or the like. The that-clause explicitly states with the content sentence what the content is. That-clauses explicitly say that the content is. Referring expressions don't per se do this even if they were to refer to entities that have content or are contents. That-clauses seem more plausibly to be understood as explicitly stating the content in a noun phrase, or clause, rather than to refer to something that is the content.

4.6.2 Anaphora

We have encountered two views of nominalizations. One takes them to be referring expressions, the other takes them to be non-referring noun phrases. In this section I would like to point to some evidence that such nominalizations should not be understood as referential expressions. To do this I will in this section also consider event nominalizations, that is noun phrases that are about events, like

(194) a. John kissing Mary

b. John's kissing of Mary

c. the kissing of Mary by John

These have been ignored in most of this dissertation, mainly because they give rise to extra complexity that can not be dealt with here in sufficient detail. Even though we focused on talk about properties and propositions here, as opposed to talk about events, facts and the like, we can learn from how the former relate and interact to the latter even given our more limited goals.

In this section we will look at anaphora that have nominalized noun phrases as their antecedent, like

(195) That Fred will win the nobel prize is unlikely, but it would be deserved.

(196) John kissing Mary wasn't supposed to happen, but it happened.

In (196) "it" is a pronoun that has the nominalized noun phrase "John kissing Mary" as its antecedent, which is to say that its semantic value or semantic function is dependent on this phrase. How this dependence works will be the issue of this section.

There is a standard view of the semantic interpretation of anaphora that have a referential phrase as their antecedent, like

(197) John is tall, but he can't swim.

According to this view the pronoun "he" is also a referential phrase that picks up its referent from its antecedent, in this case "John". Simply put, if event nominalizations are referential, then the same story should work for them, and if that-clause are referential, then it should work for them, too. And it seems that there is no problem with this, at first. (196) seems to be fine with that proposal. However, there is a whole class of cases where this can't work. These show that the above standard approach to anaphora with referential antecedents can't be carried over to anaphora with nominalized antecedents (at least not in general). The reason for this is that many pronouns can be bound by the same antecedent formula, and that this can happen across categories. Let me explain.

One and the same antecedent can bind many anaphora. This is no problem, and occurs with clearly referential ones as in

(198) John is tall, but he can't swim, and he can't skate. In fact, he would like to be even taller.

all the occurrences of "he" in (198) are bound by the same antecedent, namely "John". Similarly, one might think it works in

(199) That philosophy is not an important discipline is a sad fact. We believe it, we know it, but we wish it wasn't true.

With (199) there is no problem with the standard approach, but what refutes that standard approach for nominalizations is that these anaphora can occur in different categories, even though the are bound by the same antecedent. Consider

(200) John kissing Mary was such a surprise to Fred that even though he saw it, he could not believe it.

The first occurrence of "it" in (200) is the argument of "saw", which takes either event or propositional arguments. Thus, it is fine to say

(201) Fred saw John kissing Mary.

but also to say

(202) Fred saw that John kissed Mary.

This is not to say that they mean the same, of course. They don't. (201) is ambiguous in a way in that generally occurs with perception verbs and event descriptions. It can either be taken to be read cognitively loaded, or cognitively neutral. In the first case Fred would not only observe the event, but also recognize what was going on. In the second he would observe the event, but might not realize what was going on. So, in the second sense it can be that Fred saw John kissing Mary, but he neither recognized John, nor Mary, nor what they were doing.[8] We will get back to this in a moment.

The second occurrence of "it" in (200) is the argument of "believes". It only takes propositional arguments. Thus it is fine to say

(203) Fred believes that John kissed Mary.

but not

(204) *Fred believes John kissing Mary.

But now, if nominalizations were referential expressions, and if the standard view of anaphora that have referential expressions as their antecedent were correct then we would have an expression that is the argument of "believes" in (200) and that refers to an event. The second "it" in (200), the one that is the argument of "believes", has an event nominalization as its antecedent. It refers to an event if it inherits its reference from its antecedent. Of course, only if "John kissing Mary" is a referential expression.[9]

It seems to me that the standard view of the semantics of anaphora that have referential antecedents is correct (I will support this in a second). So, I take the above examples as evidence against the view that nominalizations are referential expressions. But since I use anaphora as an argument against the referential view of nominalizations, what better account is there of their semantic function? It is one thing to say that a certain account has a problem, quite another to offer an alternative account.

How can we account for such examples? I think it is actually quite straight forward. Such anaphora are not referring expressions, but pronouns of a kind that depend on the

[8]See (Barwise and Perry 1983) for more on this ambiguity, and an account to deal with it in a semantic theory.

[9]The point could already have been made by considering examples like

(205) John kissing Mary was a real surprise, and Fred just couldn't believe it.

However, I think the longer ones with two pronouns are better. There it is clearer what is meant. That's why we had to go through a little bit of extra trouble. A further problem, which I skip, is that it seem wrong to say that seeing is a relation to propositions. It seems to be a relation to facts. This gives rise to further evidence for the point I would like to make here.

syntactic material present in their antecedent. They are substitution operations of a kind, that is, they should be understood as having the function of taking their antecedent and putting it in their place. But it can't be that simple, since in (200) the antecedent can't just be put in place of the second "it", the one that is the argument of "believes". To start with a simple example, we could represent the semantics of

(206) That John kissed Mary is a fact, and I believe it.

as

(207) [John kissed Mary]$_{i,prop}$ is a fact, and I believe [it]$_{i,prop}$

[John kissed Mary]$_{i,prop}$ hereby is supposed to be understood as the result of nominalizing the sentence "John kissed Mary" into a propositional nominalization, a that clause. The "i" is a marker, which indicates what the antecedent of the anaphora. [it]$_{i,prop}$ is supposed to be. It is the result of transforming the antecedent that is market with "i" into a propositional nominalization. This technique is general enough to work across categories. The problematic sentence (200) can now be understood as

(208) [John kisses Mary]$_{i,event}$ was such a surprise to Fred that even though he saw [it]$_{i,event/fact}$, he could not believe [it]$_{i,prop}$.

The only difference to (207) is that the kind of nominalization that the antecedent is put in differs from one anaphor to the next. And which kind it is will have to be understood as being determined by the position it occurs in.

So, there is an alternative account of the working of propositional anaphora. And if the standard account of anaphora with referential antecedents is right then this is good evidence that nominalizations are not referential noun phrases. But is the standard account right? Isn't the account just proposed such that the standard account falls out as a special case of it? Consider, again:

(197) John is tall, but he can't swim.

Certainly the "he" in it could be understood as a pronoun of laziness, one of a much simpler kind than the ones we just dealt with. But if anaphora that have referential expressions as their antecedent can be understood as pronouns dependent on the syntactic form of their antecedents, and therefore as basically being of the same kind than propositional anaphora, then it seems that the above arguments are not reasons to believe that nominalizations are not referential expressions. So, are pronouns with clearly referential antecedents such pronouns? I don't think so. The suggestion that they are works fine with cases like (197), but it does not work in all cases. However, the suggestion that they are referential expressions that inherit their referent from their antecedent works in all cases. To see this, consider

(209) Superman is extremely powerful, but when he is dressed as a reporter Lois Lane does not even believe that he could fly.

In (209) "he" can't be a pronoun of laziness since Lois Lane does not believe that Superman can't fly. She believes of a certain guy who happens to be Superman that he can't fly. That is an important difference.

So, there is an important difference between anaphora that have clearly referential expressions as their antecedent and those that have propositional or event nominalizations as their antecedent. I take this to be evidence that these nominalizations are not referential expressions.

To be sure, examples of this kind raise a number of very tricky and delicate issues. There are so many more things that can be said about these cases, and I don't expect anyone to take these examples as a refutation of the externalist position. They do, however, give rise to a problem for externalism, a problem that doesn't arise for internalism, and that should be taken as evidence against externalism. Before we leave this section, I'd briefly like to discuss some of the extra complexity that these examples give rise to, but that I won't be able to deal with in all necessary detail.

One of the issues that has to be looked at is the ambiguity that arises from the use of perception verbs and event descriptions, as pointed out above. A sentence like

(210) Fred saw John kissing Mary.

can be understood in two ways. First, in the sense in which Fred saw that John kissed Mary. This is cognitively loaded seeing, in which Fred recognized the people and what they are doing. Secondly, in the sense in which Fred is only required to be visually related to the event, but where he isn't required to recognize the people or what they are doing. Thus is the cognitively neutral sense of seeing. It might seem that (210) is ambiguous because in it "John kissing Mary" can be understood as standing either for an event, or for a more fact or proposition like thing. And because of this we get the ambiguity between the cognitively neutral and the cognitively loaded reading. And it might seem that this is most relevant to the above examples, since if event nominalizations can stand for proposition like things then the above examples don't show what they seem to show.

These are important and tricky issues, but I think the above argument still goes through when we keep them in mind, simply because we can force an event nominalization to stand for an event by modifying it in a certain way. Then we can still have an anaphoric chain with such a modified event nominalization as the antecedent of the chain and one of the anaphors in the chain has to stand for a proposition. The first can be done by talking about the temporal extension of the event. Consider

(211) John kissing Mary was a real surprise to everyone. It lasted only for a few seconds, and even though Fred Saw it, he couldn't believe it.

This sentence surely is ambiguous, too, but one of its reading is that Fred saw John kissing Mary, even though the kiss lasted only a few seconds, and he couldn't believe that John kissed Mary. And this reading of (211) seems to be incompatible with the anaphoric chain having a referential antecedent, assuming that the standard way of dealing with anaphoric chains having a referential antecedent carries over.

A further complication arises from the alleged possibility to deal with these examples as cases of ellipsis.[10] For example, the pronoun "it" that is in the scope of "believes" can be understood as being short for "that it happened". With this we can accommodate that the pronoun is referential, and that what is in the scope of "believes" is a propositional object. This allows us to deal with cases like

(212) Even though Fred saw John kissing Mary, he couldn't believe [that] it [happened].

Thus one might think that in anaphoric chains that go across categories, where some members of the chain stand for propositions, and others for events, the occurrences of "it that stand for propositions could be understood as being elliptical for "it happened". With this trick we could understand "it" to consistently refer to an event, but the occurrence

[10]Thanks to Jeff King for pointing this out to me.

of "it" within the scope of "believes" still would be fine, since it is really a part of the that-clause "that it happened". This is a legitimate reply to the above example, however, it is restricted to examples as the above ones and doesn't work in general. The reason for this is that we can turn an event nominalization into a that-clause with this trick, as that-EVENT-happened, but we can't do it the other way round. A that-clause can't be turned into an event nominalization by a similar procedure. The above example (200) was such that we started the anaphoric chain with a event nominalization and ended with what has to be a that-clause. But we can also do it the other way round. We can start with a that-clause and later require "it" to stand for an event, as in

(213) That John kissed Mary was a shock to everyone. Fred was so shocked that he couldn't believe his eyes when he say it happening.

"it" as it occurs in (213) can't be understood along the elliptical lines, but the same argument that I gave above goes through.

4.7 Towards solving the second puzzle

The second puzzle about ontology that we started out with in chapter 1 is the puzzle that in everyday and philosophical practice ontological questions about properties and propositions play no noticeable role. It seems that for our practice these questions are of no great importance. However, it also seems, or seemed in chapter 1, that basically everything say depends on there being properties and propositions. Without them nothing we say will be literally and objectively true.

It should be clear how internalism provides an answer to this puzzle. According to internalism our talk about properties and propositions has no ontological presuppositions for its literal and objective truth. Thus the existence of any special entities has nothing to do with the literal and objective truth of our talk about properties and propositions. Thus the common lack of interest in ontological worries about properties and propositions is quite justified, and ontological worries about them are no threat to the literal and objective truth of our talk about properties and propositions. This, in outline, is the solution to the second puzzle.

4.8 The expressive power of talk about properties and propositions

Talk about properties and propositions gives rise to increased expressive power, as we have seen. And internal quantification also gives rise to increased expressive power. Internalism claims that the increased expressive power that arises from quantification over properties and propositions comes from internal quantification over them. In chapter 2 we have seen how we can understand the increased expressive power of internal quantification, but we have only looked at this in the simplest case. By now we have seen a number of further complications, and these are relevant in saying more precisely what the expressive power is that we get from quantification over properties and propositions.

In chapter 2 we only looked at what the truth conditions should be taken to be of statements that contain internal quantifiers ranging over regular objects. Now we have to look at what the truth conditions should be taken to be when the quantifiers range over properties and propositions. In chapter 2 we saw that internal quantifiers have to have truth conditions such that having these truth conditions gives them a certain inferential role. The optimal way for these expressions to have the inferential role for which we want them is to make a contribution to the truth conditions such that the resulting sentence is equivalent to a certain infinite disjunction or conjunction. In chapter 2 we simplified and only considered names and other simple terms to figure in these infinitary statements. However, this simplification doesn't even apply in the case of quantification over properties and propositions any more. Internal quantification over properties and propositions will still give rise to the expressive power that we get with a certain fragment of infinitary logic, but this time it will be more complicated, for a number of reasons.

Consider a simple case of talk about properties, like

(214) There is a property Joe and Jill have in common.

If the quantifier is understood internally then this sentence will have it as its inferential role that any sentence of the form

(215) Joe is F and Jill is F.

will imply it. Thus the truth conditions of (214) can be seen as being the infinite disjunction of all the sentences that imply it, i.e. the disjunction

(216) $\bigvee_F$ (Joe is F and Jill is F).

But what are the Fs on this occasion? In chapter 2 we simply took all the names and simple terms as the class of terms that build up the infinite disjunctions and conjunctions. Now it has to be the class of all predicates of the language. This will give rise to three complications:

- Because of the considerations involving inexpressible properties we will have to allow for a special set of variables that model demonstratives to occur in the predicates of the infinite disjunctions and conjunctions, and we have to allow infinite blocks of quantifiers that bind these special variables.

- It isn't so obvious what the class of all predicates of the language is here. Having the internal quantifiers available allows for the formulation of new predicates, ones that are build up, in part, from the internal quantifiers. Thus we can't just take the class of all predicates of the original, internal quantifier free, language to model the truth conditions correctly, at least not if we don't want to simplify too much.

- Talk about properties and propositions is powerful enough to give rise to paradoxes, unless certain precautions are taken. If we allow properties to apply to themselves, and if we allow sentences to contain quantifiers that range over the very propositions that are expressed with them then we can give rise to paradoxes. Famous example are

  (217) The proposition expressed by an utterance of sentence (217) is not true.

  (218) Does the property of being a property that doesn't apply to itself apply to itself?

We will have to try to give a precise account of how to understand the expressive power of internal quantification over properties and propositions. And as it turns out, this can be done. Or at least, it can be done as well as anyone knows how to do this, internalism or not. There is no really satisfactory way to deal with paradoxes, but the best ways to do this can be carried over to the internalist account of the function of talk about properties and propositions. Also, such an account will accommodate all three of the above points, and will associate certain fragments of infinitary logic (built up from a base language, which corresponds to the language without talk about properties and propositions). We will look into more of the technical details of this in the appendix to this dissertation.

4.9 Properties and propositions in philosophy

So far we have focused on the function of talk about properties and propositions in everyday life. And we have seen that there talk about them is nothing but a device to make complicated statements about simple things. But it might also be that talk about properties and propositions on other occasions has a different function. In particular, so far we have implicitly assumed that if properties exist then they are whatever expressions like "being a dog" stand for, and if propositions exist then they are whatever that-clauses stand for. But maybe there are also other conceptions of what properties and propositions are. Take, for example, the role that properties and propositions play in philosophy. Some philosophers think of properties as language independent entities that define the basic features of reality, and they think of propositions as making up a domain of mind and language independent entities that we relate to when we have intentional states. They are the contents of our mental states, and they are there independently of anyone having mental states. Doesn't this way of looking at properties and propositions play an important role in philosophy? And even if internalism is the correct view of ordinary non-philosophical talk about properties and propositions, might not a different view be the correct one when it comes to talk about properties and propositions within philosophy?

4.9.1 Their role outside of metaphysics

Nobody talks so much about properties and propositions as philosophers. Why is that? One might that it is because philosophers have a special interest in certain philosophically special entities. But I don't think it is right.The reason why philosophers talk so much about properties and propositions is directly related to the reason why we talk about them in ordinary life. However, philosophers have even more need for such talk than we do in ordinary life. Talk about properties and propositions allows us to express generalities, and it allows us to do this despite our ignorance of the details. With it I can make claims like

(219) Fred believes everything he is told.

without knowing what he is told. And this is exactly why philosophers talk about properties and propositions all the time. It allows them to make general claims without having to have any knowledge about the details of what they are talking about. Consider, for example, the use that such talk has in the philosophy of mind, where we constantly talk about properties and propositions:

(220) Every mental property supervenes on a physical property.

(221) For every p and q, if you believe that not p unless q, and you desire p then all things being equal it is rational for you to bring it about that q.

Consider (220). In uses of property talk like this one we want to make general claims about all mental states, and about their physical realizations. We talk about properties for reasons of generality, not because we shift the domain of things we talk about away from the mind and onto some other things, the mental properties. Rather we use talk about properties to make general claims about the mind. In particular, we make general claims about the mind in ignorance. We don't know what the physical basis of any single one of our mental states is. But the expressive power that we get from internal quantification allows us to do just that. When we say that there is a physical property such that being in pain supervenes on it, we have no idea what this property is. And that's why we talk about properties to make the claims about the physical realization of metal states that we want to make. It is because of the desire to make general claims in ignorance about the details. Not because of a shift in what we talk about from the mind to some other entities.

Talk about properties and propositions in philosophical disciplines like the philosophy of mind is so prominent because there we have even more need for the expressive power that such talk brings with it than in ordinary life. That's why we talk about them there even more often than in regular, non-philosophical, communication. These uses of talk about properties and propositions fit very nicely into the internalist picture of what their function is.

4.9.2 The platonist's talk

However, not all uses of talk about properties and propositions in philosophy is like the above. When we talk about them in ontology, or metaphysics, it might seem to be quite different. For example, what about the persuaded platonist, who explicitly believes that properties are entities just like cars, and that talk about them is on no important way different? Is a persuaded platonists talk incompatible with internalism?

Well, yes and no. A persuaded platonist is free to use property talk the way they want to. And they can make statements about properties that do have ontological presuppositions for their literal and objective truth. A persuaded platonist can explicitly quantify over properties externally. Of course, that doesn't mean that properties exist. Only that the persuaded platonists talk has ontological presuppositions for its literal and objective truth. And as it turns out, their talk will be false, assuming that properties are still understood as being whatever property nominalizations stand for (in their ordinary use). So, an internalist account of this particular person's talk about properties will be false. But that doesn't mean that internalism is true in general. Internalism only says that the function of certain talk is a certain thing, not that every single person's use of this talk fits what it is for. Compare this to, say, generics. Generics have an important role for us to express information closely related to valid inference rules in default reasoning. This is not threatened by Fred's insistence that whenever he says

(222) The tiger is fierce.

that this has to be understood as an exceptionless generality about tigers, and thus is true only if each and every tiger is fierce.

Persuaded platonists are persuaded of the wrong view of what talk about properties and propositions is all about. But nonetheless, they can make statements that are only true if they were right, and if talk about properties and propositions would be talk about some domain of entities. The expressive resources of our language are rich enough to make statements that are only true if platonism is correct. However, our ordinary everyday talk about properties and propositions, and much of our talk about them in philosophy, does not have such presuppositions.

4.9.3 Realism

One might think that talk about properties and propositions plays another important role in metaphysics. Just like what has been claimed about a truth predicate, one might think that such talk plays an important role in articulating certain realist views of the world, and a internalist conception of such talk, just like a deflationary conception of truth, doesn't all us to express a certain kind of realism that we seem to be able to conceive and to articulate. The basic idea behind it is that the world exists independently of us, and that things have properties independently of us and independently of our ascribing predicates or using language in general. Internalism doesn't account for that, one might think, because it ties properties to predicates and thus fails to be able to accommodate such a form of realism.

I think this is a mistake. Internalism has no problem with a reasonable form of realism. For example, internalism doesn't entail that Fido's having the property of being a dog is dependent on our ascribing that property to him, or on our language use. In fact, the following is perfectly true according to internalism:

(223) Fido has properties even if no humans exist.

since it is equivalent to

(224) Fido is a dog, even if no humans exist, or Fido is a cat, even if no humans exist, ...

That things have properties only because humans ascribe them to them in no way follows from internalism.

In addition, internalist talk about properties is perfectly well suited to articulate our realist intuition, like that there is more to reality than we will ever be able to figure out. For example, the following are true according to internalism:

(225) There are properties of stars that we will never discover.

(226) There are propositions that will never be the content of a thought by anyone.

since they will be truth conditionally equivalent to

(227) Stars are heavy and we will never discover that stars are heavy, or stars can travel faster than light and we will never discover that stars can travel faster than light or ...

(228) That Fido ate exactly 5123 bones during the first 23456 hours of his life will never be the content of a thought by anyone, or ...

Note here that for a thought to have that content of (226), it doesn't mean that it has each one of the disjunctions of (228) as its content. (228) is an analysis, or making explicit, of the truth conditions of (226). It is not required to play any role in the metal representation of (226).

This is quite analogous to the discussion whether or not more than a deflationary conception of truth is required to (truly) articulate our realist intuitions and persuasion. The answer is that it isn't, unless one goes overboard with ones realist intuitions.[11]

Of course, some people have realist views that are incompatible with internalism, and that won't be true if internalism is true, and the other way round. If you believe that

(229) There are propositions that can never be articulated by anyone in any circumstances.

or

(230) There are properties of stars that are completely elusive to our expressive resources, and that can never be articulated by anyone in any circumstances.

then what you believe won't be true according to internalism. But why would you believe this. It is one thing to believe in realism, in the above sense, that the world is not dependent on us in a relevant sense. But it is quite another to believe that there are completely inexpressible propositions. To be sure, there are metaphysical views that are incompatible with internalism, and there are uses that some people want to make of property talk that come out false according to internalism. However, we have no reason to be worried by this, since we have no reason to believe that these metaphysical persuasions are true.

Internalism accounts for why talk about properties and propositions plays such a prominent role in philosophy, and it doesn't do anything like claim that platonists can't say what they think they say. It also fits in perfectly well with our use of property talk to express realist intuitions.

4.9.4 Other metaphysical uses

Of course, nothing I have said so far shows that properties and propositions might not play some other role in metaphysics. In particular, might there not be uses of properties

[11]See (Field 1986) for a discussion of this in relation to deflationary truth.

and propositions in metaphysics where they are not simply understood as being whatever property nominalizations and that-clauses stand for? Might properties ad propositions not be understood as the theoretical entities implicitly specified by some metaphysical theory?

I have said nothing that would rule this out in principle, and I doubt that this can be ruled out in principle. Maybe there will be some metaphysical theory that will specify something that might for some reason be called 'property' and 'proposition'. If so, these theoretical entities of this metaphysical theory won't be what our ordinary expressions 'being F' and that-clauses stand for. But besides that, we have to ask ourselves what such a metaphysical theory might possibly be attempting to do. Is it supposed to play some philosophically explanatory role? Well, one might think that such a theory will be required maybe to give an account of some metaphysically important notion, like causation, or law of nature, or the like. In the debates about causation and laws of nature people talk about properties all the time. Maybe here properties will play the role of being theoretical entities implicitly defined by some metaphysical theory.

I cannot rule this out, but I doubt it is true. The reason we talk about properties in these philosophical debates is because we want the increased expressive power, not because we implicitly subscribe to some substantial metaphysical theory that defines some theoretical entities. To see this we would have to look at the individual cases of what we do when we take recourse to properties and propositions in certain metaphysical enterprises. I won't be able to get into this here, but I conjecture that the account of the role of properties in philosophy outside of metaphysics will carry over to these cases, too.

One further issue we have to look at is the role that talk about properties plays in accounts of objectivity, and in particular in deliminating domains of discourse where we have objectivity from the ones where we don't have it. This was addressed already in chapter 1 as one of the general connections that ontology seems to have to objectivity. I will postpone this for now, and get back to it in chapter 6. There we will talk about it in a general discussion of objectivity.

4.10 Conclusion

In this chapter we have focused on the ontological presuppositions that our ordinary non-theoretical talk about properties and propositions has for its literal and objective truth. I have proposed internalism as the best account of what we do when we talk about them, why we do it, and how we do it. We have also had a look at the use of property and proposition talk in philosophy, and we have seen that we talk about them so much on philosophy for the same reason that we talk about them in the first place: the need for expressive power to make general statements. The role that talk about properties and propositions has in everyday life nicely covers the the role it has within philosophy. But of course, there are also other theoretical enterprises than philosophy that talk about properties. Most importantly, there is natural language semantics were we assign properties and propositions as semantic values to certain expressions. This use of talk about properties and propositions fits into a more general picture. So far we have focused on what we called *natural ontology* in chapter 1. Natural ontology is what is implied to exist by our shared ordinary everyday beliefs. We also saw that there are at least two other uses of the word 'ontology', namely *theoretical* and *philosophical ontology*. We will look at these broader issues about ontology in detail in chapter 6, where we will deal with he role that talk about properties and propositions has in semantics. For now, let's sum up what we have seen so far.

In chapter 1 we started out with two puzzles about ontology, one having to do that ontological questions are apparently more easily answer then they should be, and the other having to do that we don't seem to take ontological questions as serious as we apparently should. In chapters 2 and 3 we have seen that the first puzzle can be resolved by noticing that the trivial arguments that seem to solve the ontological questions too easily did not really have any ontological results, even though the arguments are valid. This required us to look at some general issues about quantification and noun phrases. At the end of chapter 3 we had seen that some talk about properties and propositions, namely the one that occurs in the trivial arguments and similar situation, doesn't have any ontological presuppositions for its literal and objective truth. In this chapter we have extended this to a view of what role talk about properties and propositions plays in everyday life in general. We have distinguished two basic approaches to this question, internalism and externalism, and I have argued that internalism not only is a view that isn't easily refuted, it is a view that makes sense of our talk about properties and propositions. I have proposed and defended internalism as an account of why we talk about properties and propositions, and how we achieve with such talk what we want to achieve with it.

The view about talk about properties and propositions defended so far has some, but only some, implications for ontology proper. At the end of chapter 3 we had only seen that inferences to the existence of properties and propositions can't be had that easily. That was consistent with the existence of properties and propositions. By the end of chapter 3 we had only seen that some talk about properties and propositions isn't talk about some domain of entities. By the end of this chapter now we have seen more. We have seen that our ordinary talk about properties and propositions never in the first place was talk about some domain of entities. This has limited ontological implications, as mentioned above in section 4.2. It has ontological implications given a certain view of what properties and propositions are. If they are nothing but whatever property nominalizations and that-clauses stand for, or refer to, then it follows from internalism that there are no such things, simply because if internalism is true then they do not even attempt to refer, thus don't refer at all. Whether or not this assumption is justified, and whether or not there is a use of talk about properties and propositions where this assumption doesn't apply is something we have to get back to in chapter 6, when we look at ontology in general. Before we do this, though, we should see if something like the internalist view of properties and propositions also applies to natural numbers. This is the topic of the next chapter.

Chapter 5

Numbers

5.1 Introduction

One of the classic problems about the relation between objects and objectivity is the problem about mathematical objectivity. How can there be objectivity in mathematics, unless it comes, so to speak, from some domain of mathematical objects? Many of the things we have looked at so far are very helpful in understanding this problem. But I won't be able to give an account of all of mathematical objectivity and how it relates to mathematical objects. That would not only be much too ambitious, but also be quite misguided. I think it is a mistake to try to give one single philosophical account of all of mathematics, as is common practice among both platonists and their enemies. Different parts of mathematics deserve quite different philosophical treatments. In this chapter we will only focus on arithmetic, the theory of natural numbers. Arithmetic is, I think, a special part of mathematics. It is the part, besides possibly geometry, that is most closely related to everyday non-mathematical life, and that can most easily be shown to arise from it. Contrary to e or π, we talk about 2 all the time in everyday life. But how our ordinary, everyday discourse about natural numbers relates to our mathematical discourse is not at all clear. It is particularly complicated because "two" occurs in apparently quite different ways in English. Understanding how they relate to each other seems to me to be the key to understanding the basis of arithmetic truth.

Roughly, there are three uses of "two" (or "2"), namely as in:

(231) Two men entered the bar.

(232) The number of men who entered the bar is two.

(233) 2+2=4

It has been one major tradition in the philosophy of mathematics ever since Frege to take it that the uses of "two" as in (232) show that "two" at least in these uses is a singular term that stands for objects, and thus that numbers are entities which arithmetic is about. It always has remained a bit unsatisfactory what proponents of these views had to say about uses of "two" as in (231). These uses often were taken as not very central to the philosophy of mathematics partly because there statements in first order logic with identity that are truth conditionally equivalent to them. But what such complicated first order logic sentences have to do with natural language sentences like (231), other than truth conditional equivalence, and how uses of "two" in such sentences relates to the other uses of "two" isn't usually dealt with in very much detail.

In this chapter I pursue an alternative route of looking at things. I will argue that uses of "two" as in (232) are not really all that central to arithmetic, and that rather uses of "two" as in (231) are the important ones. It seems to me that the particular way of looking at things that is presented in this paper has been ignored partly because of the quick transition from number quantifiers like "two men" to complex, but truth conditionally equivalent, quantifier blocks in first order logic with identity. To start out with, we will have a closer look at uses of "two" as in (231) in natural language. But before that, let me make some more general remarks.

5.2 Mathematical objects and mathematical objectivity

Whether or not numbers are objects, or more generally entities, seems to have important consequences for questions about mathematical and in particular arithmetic objectivity. If arithmetic statements are about a domain of entities then it seems that there is a clear sense in which we have arithmetic objectivity. An arithmetic statement is objectively true if these entities are as the statement says they are, and objectively false otherwise. Besides that, it seems that we also have what one might call **objective completeness**, namely: any arithmetic statement is either objectively true, or objectively false. Alternative views of arithmetic don't seem to have such nice features. For example, if one thinks that arithmetic truth should be more closely tied to axiomatizations of arithmetic, or that arithmetic truth is more closely tied to what follows from what we commonly believe about numbers, then there only seems to be a thinner sense of arithmetic objectivity, and a lack of objective completeness. If one thinks that truth in arithmetic is closely tied to what follows from what we commonly believe about numbers then there will be objectivity in arithmetic in a sense, assuming that there is objectivity in what follows from what. But this sense seems a bit thinner than the above sense, and in particular, there doesn't seem to be any guarantee of objective completeness any more.[1] There might be statements such that neither they nor their negation follow from what we commonly believe. In general, arithmetic objectivity seems to get tied to there being arithmetic objects, i.e. numbers being entities of some kind, and in particular objective completeness in arithmetic seems to be tied to arithmetic objects. Similar considerations seem to hold for other parts of mathematics as well.

The view I defend in this chapter will have it differently. Neither objectivity nor objective completeness will be tied to there being certain objects. This should not be understood that commitments to such objects can be avoided, or that arithmetic can be translated or reinterpreted in a certain way to achieve this. To the contrary, I think that if we look closely at how arithmetic arises out of ordinary talk, how we introduce the arithmetic symbolism and at a number of other things about what we do when we talk about numbers, we will see that talk about numbers never in the first place was talk about any entities. Sure enough, this seems to be false. There seem to be a number of good reasons to think that if there is objective truth in arithmetic then numbers have to exist. After all, arithmetic says that there are infinitely many numbers, or that 2+1=3, which seems to say that certain entities are identical, or the like. And if numbers exist then objectivity will, so to speak, come from them. Not so, I think. Let me explain.

5.3 Number determiners

The most common use of "two" in ordinary English is undoubtedly the one as in

(231) Two men entered the bar.

(234) I had two eggs for breakfast.

[1] This, of course, depends whether or not we are dealing with a semantic or deductive consequence relation, and whether or not we are dealing with a second order formulation of the axioms of arithmetic. Second order arithmetic is complete using a semantic notion of consequence. For a discussion of these issues, which I will skip for now, but return to a bit in section 5.4.4, see (Shapiro 1991).

(235) Two apples, please!

Besides these, ”two” is sometimes used where it occurs in the phrase ”the number of X is two”. Like:

(236) The number of men who entered the bar is two.

(237) The number of eggs I had for breakfast is two.

And there are, more rarely, uses that express more or less trivial mathematics, like

(238) Two is less than three.

(239) The number two is a prime number.

It seems that within natural language semantics it is nowadays uncontroversial[2] what the semantic function of ”two” in uses like (231) - (235) is. In these uses it is a determiner and the semantic function of them is accounted for in generalized quantifier theory. In this theory one gives a unified semantics of noun phrases like

(240) two men

(241) some men

(242) most men

(243) some drunk men

and many more. According to it, a determiner has as its semantic value a function from properties to sets of properties. It together with a noun, like “men”, forms a quantifier, which has as its semantic value a a set of properties. Which set of properties the quantifier has as its semantic value depends uniformly on the semantic value of the determiner and the semantic value of the noun. In such uses, “two” is part of a quantified noun phrase, like “two men”, which has a very similar semantic function as “some men”, another quantified noun phrase. In general, noun phrases can be distinguished into at least two exclusive kinds,

[2]I will ignore issues about plural for now (which are controversial), since the point I will make below has nothing to do with these debates. See, for example, (van der Does 1995) for the controversies about plural.

quantified noun phrases and referential noun phrases. These have quite different semantic functions. Since in the above uses "two" is a determiner and thus part of a quantified noun phrase, it is not a referring expression that has the semantic function of standing for some object, like "Fred".[3]

As a determiner, "two" takes a noun argument to make a full noun phrase. However, such determiners can occur in sentences without their argument being made explicit. Consider

(244) Three men entered the bar, and two stayed till sunrise.

(245) Every man entered the bar, but only some stayed till sunrise.

(246) Three eggs for breakfast is too much, two is about right.

In all these examples, the second occurrence of a determiner within these sentences is without an argument, at least without it showing up explicitly in the sentence. This is not too surprising and seems to be some form of ellipsis. The argument was just mentioned half a sentence ago and isn't repeated again. There are, however, a number of examples were the determiners also occurs without the argument being explicit, but where it does not seem to be a case of ellipsis. Consider:

(247) Some are more than none.

(248) Many are called upon, but few are chosen.

(249) None is not very many.

Such examples (except maybe (248)) are often used as expressing generalizations rather than being used elliptically. For example, last year a friend of mine only got one interview at the APA meeting. To comfort him I might say

(250) One is more than none.

This could be used elliptically, in the sense that what one says is that one interview is more interviews than no interview. But it could also be used to express a generalization, i.e. that one X is more X than no X. We often express generalizations in such context, usually much more obvious ones. A friend who was surprisingly left by his girlfriend might be comforted with

[3] See, for example, (Gamut 1991) for more on generalized quantifiers.

(251) Women are unpredictable.

which might be more comforting than

(252) Sue is unpredictable.

even thought the latter is implied, or at least strongly suggested, by the former.

This phenomenon occurs particularly often with number determiners. They are often used to express general comparisons of size rather than particular statements. Consider:

(253) Five are more than three.

(254) Two are at least some.

Let's call (an occurrence of) a determiner without explicit argument a *bare determiner*. As we saw, we have to distinguish between two kinds of bare determiners, the ones that are like the elliptical ones, and the ones that are used to express generalizations. Let's call the elliptical ones *syntactically bare determiners*, and the other ones *semantically bare determiners*. The distinction is supposed to be exclusive. An example of the first kind are also utterances of

(255) After dinner I will have one, too.

What thing it is the the speaker said he will have one of will be fixed by the context of the utterance. So, the context will fix a certain kind of thing, X, such that what the speaker said is true just in case the speaker has one X after dinner. So, "one" is a syntactically bare determiner here.[4] An example of semantically bare determiners, again, is an ordinary use of

(256) Two are more than none.

[4] In this example we are not really dealing with ellipsis (and that's why I called the case the one where the determiners are like the elliptical ones). In it the first argument of the determiner will be a particular one, even though it is not made explicit in the sentence what it is. Ellipsis is one case where this happens, and more generally the argument being contextually determined is another.

where the speaker is not intending to talk about any particular kind of object. In this case the utterance is true just in case for any X, two X are more (X) than no X, which is, of course, always true.

Another thing that we will have to look at first before we can go to issues more directly related to the philosophy of mathematics is operations on determiners and on quantifiers. Natural language allows us to build more complex determiners out of more simple ones. The best examples of this are Boolean combinations of determiners as in

(257) Two or three men entered the bar.

(258) Some but not all women smoked Havanas.

And we can build more complex quantifier phrases out of more simple ones. Boolean combinations are again a simple example:

(259) Some men and every women smoked Havanas.

(260) Two men and three children saw the movie.

Such operations can also occur on bare determiners:

(261) Two or three is a lot better than none.

(262) Few or many, I don't care, as long as there are some.

It is at first not clear if these operations in the above examples are operations on determiners, or on quantifiers. In any case, these are cases that get us closer to the philosophy of mathematics.

Consider now

(263) Two apples and two bananas make a real meal.

(264) Two apples and two bananas are only a few pieces of fruit.

(265) Two apples and two bananas are four pieces of fruit.

Here “Two apples and two bananas” form a complex quantified NP that consists of two quantified NPs that are joined together by “and”. The semantic function of “and” in these examples is a little bit more tricky than simply being a Boolean “and”. The reason is that we are dealing with plural NPs here and operations on plural NPs usually allow both so-called distributive and collective readings, just like plural NPs in general. To see this consider

(266) Three men entered the bar.

This could either be read collectively, where three men at once, as a group, enter the bar, or distributively, where three men entered the bar after each other, each one by himself. Similarly

(267) Three men and two women entered the bar.

could be read as that one group of five people entered, or that two groups, one of three men and in addition one of two women, entered the bar. The first would correspond to a collective reading of “and” and the second to a distributive one. Besides that, there are, of course, still the collective and distributive readings of each of “three men” and “two women”, so overall (267) has a lot of readings. In any case, this distinction is important when it comes to disambiguating cases like

(268) Pizza and cheap beer make me sick.

which could or couldn’t imply

(269) Pizza makes me sick and cheap beer makes me sick.

depending on whether or not “and” is read collectively.

Analogously with the case where we had sentences with only bare determiners, like

(270) Few are more than none.

we have sentences with only bare determiners, but where more than one of them are combined to form a complex quantified NP. Consider:

(271) Two and two are four.

Here we have three bare determiners at once in the sentence, the first two combined by "and" (in its collective reading). In most uses, these determiners will be semantically bare. So, (271) will be used to express a generalization. It will be true just in case, for any X, two X and two (more) X are four X, which is, of course, always true. The qualification "more" does not, I think, really have to be made explicit. "and" mostly should be read as "and in addition", not just in the above examples, but almost always. Consider:

(272) She only had an apple and dessert.

A usual utterance of this wouldn't be true if she just had an apple, even though fruit is perfectly fine dessert.

We saw a number of things which will be useful to understand basic arithmetic. On the one hand, we saw that in ordinary discourse we often utter sentences that contain bare determiners. In such sentences the determiner can be either syntactically or semantically bare. On the other hand we encountered a number of sentences where all the NPs in the sentence consist of bare determiners or are complex NPs consisting of bare determiners and operations on bare determiners. In all these sentences the determiners are part of a quantified NP. Quantified NPs have to be distinguished from referential NPs, which have a quite different semantic function. "two men" are in the same semantic category as "some men", and "two" (in these uses) is of the same semantic category as "some". Neither one of them picks out any particular objects, as referential NPs like "Fred" do.

It would be a mistake to think that a sentence like (271) is about determiners, at least not in the sense of about in which

(273) Fred is hungry.

is about Fred. If there is a sense of about in which (271) is about determiners then it has to be a sense in which

(247) Some are more than none.

is about determiners, too. True enough, functions from properties to sets of properties are the semantic values of "some" or "two" (in certain semantic theories). But having something as a semantic value and having something as a referent is quite different. Every phrase has a semantic value in certain semantic theories, and in different semantic theories they have different semantic values. Only few phrases have referents, and what their referents are does not depend on semantic theories.

There also is no mystery about truth when it comes to sentences with bare determiners. (247) is true, not simply true according to a fiction or the like, but literally and objectively true. It is also necessarily true. The same holds for (271). It is objectively and necessarily true.

(271) is true, and we can spell it out a bit more and be more explicit about what it says and why it is true. In doing this we will naturally use statements as

(274) If you take two, of any kind, and you add two more, of that kind, then you will get four, of that kind.

We use statements like these to spell out the truth conditions of (271), but that doesn't mean that (274) is more basic than (271). It is just used to elaborate on what the truth conditions of (271) are. Also, in (274) the conditional shouldn't be taken to be a mere material conditional. (274) isn't true whenever the antecedent is false. It rather means something along the lines that if you would do A then B would result. In particular, conditionals like (274) are not vacuously true if we talk about sizes of collections larger than the number of things that actually exist. Suppose the world only contains n objects. The the material conditional

(275) If you take n things, and add 3 more then you will get 2 things.

will always be true, since the antecedent is false. If there are only n objects in the world then you can't add 3 more. But of course, that is not how either (271) nor (274) should be understood.

These considerations certainly seem relevant to understand what we are doing when we do arithmetic. However, it would be quite premature to try to conclude anything from what we have seen so far. Arithmetic, after all is concerned with numbers, not number determiners. Still, we are closer to arithmetic already than one might think.

5.4 Numbers and arithmetic

5.4.1 Numbers and basic arithmetic

Introducing the formalism

Basic, quantifier free arithmetic is a trivial part of mathematics, but it is philosophically important for at least two reasons. First, understanding it is an important step towards understanding arithmetic. Secondly, it is in many ways the first and most basic part of mathematics. It is first when a child learns mathematics, and it is most closely connected to ordinary pre-mathematical discourse. Besides that, it is there where mathematical symbols are first introduced. To understand mathematics, it seems to me, one first has to understand the basic arithmetical truths. Among these I take truth like

(233) 2+2=4

I intentionally use the mathematical notation here, since it is itself an important issue what these really mean. What does (233) mean? How and why does this notation get introduced? Why are we using it? Well, to answer the first question, one simply has to read (233) out loud. But, as it turns out, different people say different things (at different times) when asked to do so. I did a little experiment here (and how often does a philosophy student have such an opportunity?) I asked various people to read "2+2=4" out loud. The results I got where

(271) Two and two are four.

(276) Two and two is four.

(277) Two and two equal four.

(278) Two and two equals four.

and all the above with "plus" instead of "and".[5] It seems that there are basically two ways to read "2+2=4", in the plural and in the singular. And on reflection, these two ways of looking at it seem to be quite different. It seems that if we speak in the plural and say

[5]There are of course a number of further possibilities that also seem to be right, like

(279) Two and two make four.

(271) Two and two are four.

then we are not talking about any particular objects. But on the other hand, when we speak in the singular and say

(276) Two and two is four.

then it seems that we are saying something about particular objects. Why the first is so we just saw. Why the second? Well, there are a number of reasons.

1. It seems that the "is" in (276) is the "is" of identity. What is said in (276) is pretty much that if one performs a certain operation on certain objects then the result is identical to another object.

2. It seems that "two" and "four" in (276) are singular terms that have to refer to objects in order for the sentence to be true.

3. It seems that we have the following, apparently valid, inference:

   (a) Two and two is four.

   (b) So, there is something such that it and two is four, namely two.

   (c) So, two exists.

It seems to be the standard persuasion that in (276) both "two" and "four" are referring expressions. There is, of course, a further issue whether or not they succeed in referring, and there certainly is disagreement. And there is disagreement among those who believe that "two" and "four" in (276) succeeds in referring about what they refer to. This is a disagreement about what objects numbers are.

I think it is prima facie very plausible to assume that when we speak in the singular and utter (276) then we are referring to objects. Maybe even more plausible, again prima facie, than that we are not referring to objects when we speak in the plural and utter (271). However, I think we have the best reasons to believe that when we are speaking in the plural we are not referring to objects. These reasons are just as good than the reasons we have for believing that "some men" is not a referring expression.

All this should be a bit surprising. There doesn't seem to be such a big difference between

(271) Two and two are four.

(276) Two and two is four.

Sure, one is plural, the other singular. But could it be that in the first we are not talking about any objects, but express a generalization using bare determiners, whereas in the second we talk about particular objects and operations on them? It seems that if your math teacher asks you to say a truth of arithmetic you would get away with either one. It doesn't quite seem right that only one of them is real mathematics whereas the other one is only folk talk. So, whatever one's view about arithmetic is, whether one believes it is about objects or not, one has to have an account of what the relation is between (271) and (276), and how either one of them relates to "2+2=4". And to have such an account is to have an account of how the use of "two" in ordinary discourse relates to the use of "two" in mathematics. It is to say what the relation is between "two" in

(280) Two apples are on the table.

and "2" in

(281) 3+2=5

It seems to me that the correct account of this as been overlooked, partly because philosophers of mathematics usually ignore how and for what purpose the basic mathematical formalism gets first introduced and how we first come to encounter arithmetic. It seems to me that looking at how the child learns arithmetic, which is, after all, how we all learned it, gives us key hints about what the relation is between the plural and the singular talk in the above examples. This will get us quite a step closer to see whether or not arithmetic is about objects.

Learning basic arithmetic

How a child learns arithmetic is of considerable interest from many points of view. On the one hand it is of interest to developmental psychologists, who take the child's mastering of counting, arithmetic operations and the number concepts as an important case study for large scale issues in developmental psychology. On the other hand it is of interest educators who would like to understand better how the child learns in order to be better in teaching children. For these reasons there is a substantial literature on how the child learns mathematics, in particular arithmetic, and how this can be improved by using certain teaching methods. It seems to me that this literature (or at least the small part I know of) is full of hints that are most useful for the philosophy of mathematics. Amongst others, they concern what the relation is between symbols like “2” and ordinary determiners like “two”.

In learning basic arithmetic children have to learn a number of different, but connected things. To name a few important ones:[6]

1. They have to learn how to count, that is how to continue the sequence: *one, two, three, four....* (the ordinal aspect)

2. They have to be able to determine size, or answer “How many?” questions. (the cardinal aspect)

3. They have to master change of size, adding things to a collection or taking them away.

4. They have to master the mathematical formalism, like ‘2”, “+” etc.

5. They have to learn how to solve arithmetical problems purely within the formalism, i.e. give the right answer to 2789+9867-34=?

The order of this list is also how things are ordered temporally. Of course one does not have to reach perfection at one stage to go to the next. Kids don’t first learn how to count all the way to 10^{10} and then learn how to add. First kids have to learn how to count to, say, ten or twenty. Then they have to learn to give the right answers to questions like

(282) Do you see these cats in the picture, Jonny? How many are those?

This will be done by simply counting the ones that one sees, at least if there are more then just very few. After that they will master judgments about changes of size, both in actual collections as well as in imagined ones. There are a variety of exercises to do this, and a standard text book on teaching mathematics to first graders should contain a collection of a good mixture of them.[7] Examples are

[6] Others are to understand the use of numbers in ordinary life, i.e. to understand the meaning of a “20” on a bus, a cake, a house wall, a dollar bill etc.

[7] See (Verschaffel and Corte 1996, 112f.) for a classification of such exercises, and also (Becker and Selter 1996).

(283) Here we have three marbles on the floor. Now I put two other ones there. How many marbles are now on the floor?

(284) Suppose Jonny has two marbles, and Susie has three more than Jonny, how many does Susie have?

During this learning process, which takes quite some time, the teacher will introduce the mathematical symbols. The students will learn the decimal system, that "2" is read "two", they will learn to count in symbols, i.e. continue the sequence *1, 2, 3, 4, . . .* . And the student will learn to represent what was learned in exercises like (283) and (284) in symbols. After doing exercise (283) a classroom situation might continue:

(285) That's right, Susie, three marbles and two more make one, two, three, four, five marbles. (*the teacher will write on the blackboard:*) 3+2=5. (*and say out loud any of the following:*) Three and/plus two is/are/make five.

After such exercises are mastered to a reasonable degree the education will continue in teaching the children to add, subtract and multiply using the symbols alone. At this stage the children will learn tricks for adding that are based on the use of the decimal system, like carrying over ones, or multiplying with the tens first, and the ones later and the like. The child is then supposed to solve simple arithmetical problems abstractly, without imagining a collection of marbles that gets increased or diminished. The child is supposed to be able to solve problems like

(286) $26789 - 789 + (2 \times 23) =?$

In this last case, for example, it is supposed to see more or less directly and without much calculation that "789" are the last three digits of "26789" and thus subtracting it from the latter gives "26000".

Simple as all this may sound, the learning process that I just described takes quite some time, in fact several years. Which years of the life of a child these are varies. In any case, it is a substantial and difficult task, even though it now seems trivial to us.

An account of basic arithmetic

So, why does any of this matter for the philosophy of mathematics. Well, I don't think it proves anything, but it is full of hints that are most relevant to understand what arithmetic is all about. In fact, I think it naturally gives rise to both

1. an account of what the relation is between "2" as used in mathematics and "two" when used as a determiner
2. an account of what the relation is between plural and singular talk about numbers

Both of these accounts, I think, are true. Here is what they are:

To get arithmetic of the ground on first has to use changes of sizes of collections. One adds and takes away marbles from a bowl, or the like. After doing this with different kinds of things one abstracts from the particular kind of thing one uses. So, one learns that for whatever things I put in the bowl, two of those and three more make five of them. This insight is then written down in a nice notation as "2+3=5".

This notation is not inessential. Soon one will stop thinking about collections and will use the notation to to perform much more complicated arithmetic operations, ones that one could not easily perform if one would still be thinking about collections. This will simply occur when arithmetic operations on large[8] numbers have to be performed. Such operations will be split down into operations on small numbers (like the ones, the tens, the hundreds etc.) using features of the decimal system. Now, it is important to note that there is a difference between the original learning of basic arithmetic truth involving only small numbers and the learning of arithmetic truth through calculations using the symbolism. Originally arithmetic truths were learned by abstracting from particular objects used in collections, that is going from a number of examples like

(287) Two marbles and three more marbles make five marbles.

to

(288) Two and three are five, for whatever one takes.

Later arithmetic truths like 987+123=1110 will be learned by manipulating symbols in a certain way. That does not mean that there is anything wrong in doing this. And it doesn't mean that arithmetic is about symbols or anything like that. But it shows that there is an interesting turn here. We start out by considering collections that are described using number determiners and learn certain truth about operations on these determiners. That's were we use the plural, as in (288). We introduce a symbolism to express such truths and

[8]Large here means something like more than ten or maybe thirty.

learn to calculate more such truths using manipulations on the symbols. As I said, this is a substantial learning task for the child, and, I think, it is were a certain shift occurs. Whereas in the simplest cases we treat arithmetic operations as operations on (bare) determiners, and use the plural to express them, when it comes to more complex cases we start using the singular. Certainly, we can say

(289) 987+123 are how many?

but usually we just say

(290) 987+123 is what?

Now, I think there is good reason why we do this, and it is important to see why we do it to understand what we are doing here.

Operations on determiners are rather unusual things. There are only a few which really occur in ordinary speech, as in "some but not all", and in these cases they are rather simple. Complex operations on complex determiners, as in

(291) Twenty four times forty six and eight but minus three are eleven hundred and nine, of whatever you take.

are very unusual, hard to understand and simply complicated in all kinds of respects. However, if we forget about that they are determiners and forget about the plural they require and rather talk in the singular and using the formal notation things become a lot simpler.

(292) 24 times 46 plus 8 minus 3 is 1109.[9]

simply is a lot easier to understand and to verify. Talking in the singular and using symbols makes things easier for us. That's why we do it. Consider it from another angle. Operations on determiners are logically much more complicated than operations on objects. Determiners are of a rather high type. They belong to type ((e,t),((e,t),t)).[10] If we talk in

[9] I use "plus" etc. instead of "+" etc. to make clear how we would say it out loud.

[10] At least according to one way of looking at them. There is no objective fact of the matter which type they belong to. That depends on part how one sets up ones semantics in general. However, there is an objective fact that they belong to a high type. The type notation I borrow from the semantic tradition that uses type theoretic hierarchies in semantics. (e,t) can be read as the type of functions from objects to truth values, and so on for higher types. For more of this way to talk about types of expressions, see, for example, (Gamut 1991).

the singular we talk about them as if they were regular objects, that is as if they were of type (e). Operations on this simple type are much more common and much easier for us to think about. That is why we naturally start talking in the singular when it comes to learning how to perform more complicated operations on number determiners.

I have to remind you that to learn how to add and subtract is a very difficult thing for a child to learn. It takes a long time and requires repeated exercises over that time. A child will usually be quite old (relatively speaking) before it will be able to solve an apparently not so hard arithmetic problem like (292). By that time they will be able to read, have conversations, take cookies out of a jar even though they know they are not supposed to, make up stories and other complicated things. Any stance that makes these difficult arithmetic tasks easier will be gladly adopted. Speaking in the singular, lowering the type of the things one operates on and dealing with them as if they were of lowest type, the type of objects, is, I think, among them.

None of this should be surprising once we remind ourselves where we came from, cognitively speaking. Our mind evolved with the need to think about objects and their properties and relations. What is most important to think about are regular midsize everyday objects. That's what our mind is really good at. And it isn't just as good at many other things. Many things that we think about nowadays, are cognitive luxuries, are really not the kinds of things that our mind was built to think about (so to speak). It should thus not be surprising that when we think about a certain problem that our mind isn't really made to think about, like arithmetical problems that go beyond trivialities, we naturally try to force the problem into the domain of thinking that the mind is made to think about. So, not surprisingly, once arithmetic gets hard for the child, it is much better off to employ its well developed ability to think about objects, even though arithmetic isn't really about objects. And to do so is, of course, not a bad thing. After all, we do manage to learn to perform complex arithmetical operations. The only downside is all the philosophical confusion this gives rise to.

Let me wrap up. It seems to me that if we look at how children learn to add and multiply we will see that even though arithmetic has sizes of (arbitrary) collections as it's subject matter, nonetheless, we will naturally violate the requirements this brings with it, namely speaking in the plural and operating on logically complex things when we are adding and subtracting. We will adopt speaking in the singular to simplify the process of learning how to solve more complex arithmetical problems. However, this is just a way of speaking and, above all, thinking, that we adopt to simplify our task. The subject matter of arithmetic is not affected by this. So the difference between the singular and the plural reading of “2+2=4” is insubstantial when it comes to truth conditions. The singular is a way of speaking here that is insubstantial just like using the singular in certain dialects of English, for example

(293) Your parents is nice people.

Basic arithmetic truth has the same status as truth expressed with sentences containing bare determiners. They are objectively and necessarily true, even though there is no reference to any objects. Furthermore, every one of such basic arithmetic statements is either objectively true or objectively false. For basic arithmetic we have objective completeness.

Speaking in the singular is speaking as if we would not operate on determiners and quantifiers, but on objects. This is a tool we adopt to simplify arithmetical calculations. But to this this is not to stipulate that there are such objects, nor to assume that there are such objects, nor does it show that there are such objects.

There are two more issues about quantifier free arithmetic that have to be addressed here, none of them, I think, bring in anything substantially new. First there is the issue about how “plus” relates to “and” and how either one of them relates to “+”. I think there is no real issue about “plus”. It is simply a Latin word that got adopted into the English language and means basically the same as “and” in its collective reading. As we saw above, “and” has different semantic functions on different occasions. Sometimes it is simply a Boolean connective, sometimes it it unifies collections, etc. “plus” is certainly fine instead of “and” in these latter uses. Consider

(294) Two apples plus two bananas give you all the vitamins for the day.

So,

(295) Two plus two are four.

means the same (modulo subtleties) as

(296) Two and two are four.

The the story for the singular just carries over.

The second issue is multiplication.[11] Multiplication is interestingly different from addition in many respects (besides interesting technical differences between arithmetic with and without multiplication). However, for our discussion now there is no substantial difference between addition and multiplication. Consider uses that the phrase "four times" has in natural language, like

(297) He came to the restaurant four times in the last year.

(298) Four times he had two apples without paying.

(299) Four times two apples makes eight apples he hasn't paid for.

(300) Four times two are eight.

Semantically "four times" seems to be more difficult than "four men" since in (297) "four times" seems to modify the VP. But anyways, the more general facts that I claimed about addition seem to carry over for multiplication, too. Sentences like (300) are true, and objectively true, without being about any particular objects.

5.4.2 Nominalizations, again

Arithmetic arises out of talking and thinking about sizes of (finite) collections of ordinary objects. And looking at it that way does not make it seem that it is about any objects. However, there are also rather plausible arguments that numbers exist. And if numbers exist then it seems that arithmetic is about them and therefore is about objects after all. One of the most influential arguments for this is the argument from nominalizations (or so I will call it). This section will deal with it. As you can imagine, the general view on the function of certain nominalizations that was developed in chapter 3 will be most relevant here.

The argument from nominalization goes back to Frege and plays an important role in his (Frege 1988) and has recently been revived by Crispin Wright in his (Wright 1983). The argument comes in several different forms, with different degrees of plausibility.[12] I will not

[11] Subtraction and other operations on the natural numbers that the natural numbers are not closed under will be discussed below.

[12] Consider, for example, the different presentations it gets in (Rosen 1993) and (Field 1989a)

get into the differences between these different presentations now and simply present a very simple form of it. It is also, I think, the most plausible version of this argument. It simply goes as follows:

1. Without question, everyone accepts that some sentences like, say,

   (301) Five men entered the bar.

   are objectively true.

2. (301) is truth conditionally equivalent to

   (302) The number of men who entered the bar is five.

   It is impossible that these two differ in truth value.

3. Therefore (302) is objectively true as well.

4. But (302) is an identity statement. It claims that what is denoted by two singular terms is identical. It claims that the number of men who entered the bar is identical to five.

5. But for a sentence of the form "s=t" to be objectively true the two objects denoted by "s" and by "t" have to exist.

6. Therefore, the number five exists.

And if numbers exist, then arithmetic most certainly is about them.

This is a very persuasive argument, and it has persuaded many philosopher. However, it is based on a mistake about what the relation is between

(301) Five men entered the bar.

and

(302) The number of men who entered the bar is five.

in these uses. As we have seen in chapter 3, the difference between them isn't one of reference to different objects. That wouldn't allow us to account for why we make such inferences this easily. It is rather a difference in how the same information is articulated.

The second one brings with it a certain focus, the first one can be uttered neutrally. Both can have additional focus due to intonation. In addition, it is a mistake to think that simply because the words "is the same as" or even "is identical to/with" occur between two noun phrase that these two noun phrases have to be referential. These expressions are used also with apparently non-referential noun phrases to express much more complex and elusive facts than the sameness of some entity, as in

(303) How John wants to live his life is identical to how Plato lived his life.

The standard route to the conclusion that objects are presupposed for objective truth in arithmetic is thus mistaken.

5.4.3 Quantified arithmetic

We have seen how talk about the sizes of collections gives rise to quantifier free arithmetic, and how this is not talk about any objects. Still, though, these statements, both formal and informal, are objectively true or false. Quantifier free arithmetic, however, is only trivial mathematics. To understand arithmetic we will have to look at quantified arithmetic. There are two issues we will have to discuss here. First, we will have to see whether or not adding quantifiers to the fragment of arithmetic we have looked at so far is in any conflict with the account of arithmetic as not being about any particular objects. Secondly, we will have to see whether or not the story about objective truth and objective completeness that we saw for quantifier free arithmetic carries over to quantified arithmetic.

The first question, of course, concerns whether or not ordinary quantification over numbers should be understood internally or externally. If it is understood internally then it will be clear from what we have seen so far that they literal and objective truth of such quantified statements would indeed have no ontological presupposition. This is because we have seen that the quantifier free statements of basic arithmetic don't have any ontological presuppositions for their literal and objective truth. And as you can imagine, I think they should be understood internally. After all, what are we trying to say when we say something like

(304) Every number has a successor.

Are we trying to say that every one of the numbers that exists has a successor? Shouldn't there then be at least in principle be an issue which numbers exist? It rather seems that when we say this we want to say something that is equivalent to

(305) 0 has a successor, and 1 has a successor, etc..

and is independent from some domain of elusive entities for its precise expressive strength. Whether or not numbers are objects is not of importance for the truth of such ordinary statements.

This also answers the second issue we have to address in this section, namely whether or not the objective status of quantifier free arithmetic carries over to quantified arithmetic. The answer is that it does, since the truth conditions of statements of quantified arithmetic are correctly modeled by infinite conjunctions and disjunctions of statements of quantifier free arithmetic. And disjunctions and conjunctions of objectively true or false statements are themselves objectively true or false. In addition, since every arithmetic statement will be truth conditionally equivalent to a Boolean combination of statements of basic arithmetic, we preserve objective completeness for quantified arithmetic. Thus not only quantifier free arithmetic is objectively complete, but quantified arithmetic is as well.

5.4.4 Formal Arithmetic

Real mathematics gets carried out with the use of the mathematical symbolism. And this includes arithmetic. In particular, when logicians look at arithmetic, they usually study certain first or second order theories. There the language of arithmetic is a formal language. It is not uncommon for philosophers of mathematics to take their starting point from there and to look at how we should understand these formal languages. In particular, since all usual formalizations of arithmetic take expressions like "two" to be formalized as constants, it is often assumed that "two" is a referring expression. But that seems a bit fast. Clearly the natural language talk is more basic than the use of formal languages. The formal languages are just written simple representations of mathematical talk. When we speak we still say everything in ordinary English. Even though the use of formal languages in mathematics is derivative on natural language, they are important tools in the study of arithmetic, and many other disciplines. Without their use certain important results, like the incompleteness theorems, would never have been found. But still, their role is often overemphasized, in particular in the philosophy of mathematics. I have heard more than once an argument of the following kind:

(306) It can be shown that everything one can do in theory T, formulated in a language L, can be proved in theory T', in language L'. Furthermore L' is nominalistically acceptable, contrary to L, so a certain branch of mathematics is nominalistically acceptable.

But prima facie it seems that the fact that one can give a formal representation of mathematical talk in one or the other formal language does not have much philosophical importance. After all, it is the informal talk that has to be understood, not the formal representations of it. Formal languages are simple representations of ordinary mathematical talk. Their main feature is that they are very useful in understanding our mathematical reasoning. And their being simple is part of their being so useful. As simple representations they do not have to correctly represent every aspect of ordinary mathematical talk. For example, it would certainly be a mistake to criticize first order formulations of arithmetic, like Peano Arithmetic (PA), that they do not give the right representation of natural language quantification. Certainly it's not quite right to represent "every number is Φ" as "$\forall x(N(x) \to \phi)$". Amongst others, this misrepresents the NP/VP structure of the original sentence. This does not have to be so. For example, in Montague's type theoretic framework, this can be avoided. But Montague's type theoretic framework is much more complicated than simple first order logic. Using it we would lose a lot of the nice things that we get in the first order representation. For the purpose of giving a formal model of mathematical reasoning, we do not care about representing the NP/VP structure of sentences. We care about completely different things. In natural language semantics, we do care about representing the structure of the sentence correctly. There first order logic has to go. But in formal models of mathematics other things are important. There we want to model our informal mathematical reasoning, have nice representations of the most general principles or statements that we take recourse to over and over, and have a representation of how we proceed deductively. So, it is at least a necessary condition for a formal model of mathematical language to be correct that it correctly mirrors valid deductions or proofs. In such a model we take a fragment of a natural language and some formal language, map sentences of the first onto sentences of the second, and do all this in such a way that an informal argument or proof in the first is valid iff the image into the second is a formally correct deduction. More precisely:

> **Definition** Let **F** be a fragment of a natural language that is used in mathematical discourse. A formal model **M** of **F** is a pair $\mathbf{M} = < \mathbf{L}, \varphi >$, whereby **L** is a formal language (with a semantics and syntactic calculus that specifies a notion of a proof) and φ is a recursively defined function that maps expressions of **F** onto formulas of **L** such that sentences of **F** are mapped onto sentences of **L**.
>
> A formal model **M** of a fragment **F** of a natural language is called an **Deductively Faithful Formal Model (DFFM)** iff
>
> $\sigma_1, \ldots \sigma_n$ are the assumptions in a valid argument or informal proof of σ_{n+1} iff $\varphi(\sigma_{n+1})$ can be deduced or proved in the calculus from $\varphi(\sigma_1), \ldots \varphi(\sigma_n)$
>
> The image under φ of all the assumptions that we implicitly make and that we accept without further proof will make a good set of axioms of our informal reasoning. Even better will be a simple set of sentences from which all the above can be deduced.

In modeling mathematical language we are more interested in modeling proofs than in preserving certain semantic properties of the natural language that is used in proving or arguing. That's OK, of course, formal models in mathematics don't have to be the same as the formal models of semantics, even if the model the same language. There are quite different interests in modeling.

This should make it clear why it is perfectly acceptable to use first order Peano Arithmetic as a formal model of arithmetical discourse, even though it neither models the NP/VP structure of the sentences correctly, nor does it model "two" correctly from a semantic point of view. Taking "two" to be a name or a constant is OK, if this helps in having a simple model of mathematical arguments. And, as we know, doing this gives us a very powerful and successful formal model.

5.4.5 Conclusion about arithmetic

According to the above account of arithmetic, arithmetic is not about objects, but every sentence in the language of arithmetic is either objectively true or objectively false. Furthermore, arithmetic truth is independent of any axiomatizations of arithmetic. The basis of arithmetic truth are changes of finite collections with respect to size. But, nonetheless, arithmetic is not about collections in the sense that we refer to collections in any of our ordinary statements like

(307) Three apples are on the table.

(307) does not imply that besides three individual apples there is a further entity which is a collection of apples.[13]

According to usual classifications this view is nominalistic because arithmetic is not about any abstract objects. However, it is also platonistic in the sense that every statement of arithmetic is either true or false, and objectively so. Is it also a logicist view? This, I think, is a more tricky question, because it is not clear what really counts as logicism. A number of things are true according to the present account that are often associated with logicism:

1. Arithmetic truths are necessary truths that do not imply the existence of any objects.

2. Arithmetical notions are invariant under permutations of the domain of objects.[14]

3. Knowledge of arithmetic truth does not seem any more (but also no less) mysterious than knowledge of other truths that are commonly considered as logical truths.[15]

In any case, I don't think that the question whether or not the account is a form of logicism is really of much importance. What is important is whether or not arithmetic truth is literal and objective truth, and whether or not every statement in arithmetic is either true or false independently of any axiomatization or description of arithmetic. Fortunately everything turns out to be just fine.

5.5 Beyond arithmetic

This section as well as this dissertation only aims at giving an account of our talk about natural numbers, both in everyday life and in arithmetic. But after what I have said so far the question naturally arises how it relates to other parts of mathematics. Is this somewhat logicist story supposed to carry over to all of mathematics? If yes, how is this supposed to go? If no, what about other number systems? What about rational numbers, and real numbers? What about set theory and geometry?

[13] See (Boolos 1984) for more on the relation between plural and talk about collections.

[14] See (Tarski 1986), (van Benthem 1986), (van Benthem 1989), and (McGee 1996) for accounts of the demarkation of logic in this spirit.

[15] I do believe that we know arithmetic truths a priori in a reasonable sense of the word. I hope that the above account has the potential to be extended in such a way to include an account of arithmetic knowledge. However, this will raise a number of tricky issues that should better be left aside for now.

What I have said so far applies to talk about natural numbers only. It doesn't carry over to any other part of mathematics, with some exceptions to be mentioned below. It is a mistake to think that one philosophical story will apply to all of mathematics, and that one account, platonism, fictionalism, structuralism, or what have you, can be one wholesale account of all of mathematics. Different parts of mathematics are different in many important respects. In particular, number theory is special in certain respects. In number theory we have both objectivity and objective completeness. The first is to say that truth in number theory is literal and objective truth, not given anything, like the number fiction, or what have you. But this might not apply to other parts of mathematics. In other parts we might not have objectivity in this strict sense, but only in the weaker sense spelled out in chapter 1. In addition, other parts of mathematics might not have objective completeness. It might be that other parts of mathematics give rise to questions that have no objective answer. As we have seen, objectivity and objective completeness are easy to conflate, but they might come apart. In particular, we might have objectivity without objective completeness. We will see a case where this might apply below.

In the remained of this chapter I would like to briefly outline some things about other number systems, and in particular how they relate to arithmetic. This will be brief and sketchy, since these further issues are beyond the scope of this dissertation. However, the contrast with arithmetic, and the relation of other parts of mathematics to arithmetic might be of interest nonetheless.

5.5.1 (Positive) rational numbers

There are a number of close similarities between the natural and the positive rational numbers. Both "2" and "1/2" function as number determiners, and both of them are used as bare determiners. For rational number determiners we get very similar examples as for natural number determiners:

(308) One cake is enough, but one and a half would be better.

(309) 90% of the men in the bar smoked Havanas.

(310) Two and a half are more than two and a quarter.

To be brief, it seems to me that everything I said about arithmetic carries over to the positive rationals. Sure, the operations on the rationals, or on rational number determiners, are more difficult to learn, and maybe the natural numbers are conceptually more basic and prior to the positive rationals.[16] But nonetheless, the basic story about objectivity just carries over, or so it seems to me. In this respect there is no difference between the positive rational and the natural numbers. In particular, I don't think that the positive rationals should be viewed as being defined from the natural numbers.

5.5.2 Negative numbers

Even though the positive rational numbers don't really bring in anything new, the negative numbers, both whole numbers and rational numbers, bring in something essentially different from a philosophical point of view. Negative number determiners don't make much sense and do not occur in English. Sentences like

(311) Minus three men entered the bar.

are nonsensical. Operations on negative numbers do not arise in the way operations on positive numbers arise, namely from talk about changes of the size of arbitrary collections. And, how could they? A collection with minus three members doesn't make sense. Negative numbers and operations on them are also, and because of this, especially hard to learn for the child. None of the tools for teaching that are most useful in the case of arithmetic are available to the teacher here. Asking the child to consider how many things of a certain kind are on the table when such and such a change of the number of things happens doesn't work with negative numbers. The child will rightly say that when you try to take 5 marbles away from the 3 marbles on the table then you'll reach the point where there aren't any marbles to take away. That there are minus two marbles left certainly won't be the answer. In teaching children about negative numbers teachers usually take recurse to cases were we use the negative numbers in real life. In particular, teachers will use the example of money owed. If I have only five dollars and spend ten then I owe someone five. So, I have minus five dollars. That is a helpful teaching tool, but it is rather metaphorical. I don't really have minus five dollars, but I owe someone five dollars. And that is to say that if I get n dollars then I will have $n - 5$ dollars, given that I pay back my debt promptly.

To understand the negative numbers we will, again, have to look how and for what purpose talk about the negative numbers and their symbolism were introduced, how we

[16]That is not so clear, though. Certainly the concept of one is more basic than the concept of one half. But it is not clear whether or not the concept two is more basic than the concept of one half.

learn about them and the like. Without getting into the details of this, let me just say that it seems to me that the negative numbers get introduced for the following reason. When doing arithmetic one does not only use addition and multiplication, but also subtraction. And subtraction makes perfectly sense within the way of looking at arithmetic that I endorsed above. Subtraction comes from taking away objects from collections. However, in arithmetic there is always one restriction to be obeyed. One can't take away more than there are. So, subtraction in arithmetic is only defined if what one subtracts isn't larger than what one subtracts from. $n - m$ is only defined if $m \leq n$. However, there are often cases where a problem seems to have a perfectly fine solution within arithmetic, like

(312) $(n - 7) + 5 = 3$

$n = 5$ seems to be the solution to this problem, but $(5 - 7)$ is undefined. These kind of problems can simply be avoided by doing the following. Assume that subtraction is always defined and assume that the natural numbers have inverses. So, extend the natural numbers in such a way that subtraction is always defined and certain obvious principles about subtraction hold. This, I think, is originally done to simplify calculations on the natural numbers.

I would like to take this at face value from a philosophical point of view. Negative numbers are introduced to simplify calculations on the natural numbers. Thus negative numbers are what is usually called ideal elements. Things that we assume, or introduce, to simplify matters with other things that we already take as given. So, in a sense, the negative numbers are derivative on the natural numbers.

Assuming ideal elements is a form of instrumentalism. Since we will encounter different kinds of instrumentalism,[17] let's label this kind *ideal element instrumentalism*. A case of ideal element instrumentalism consists in the following:

1. There is some domain **D** of discourse that makes sense prior to the introduction of the ideal elements.
2. Ideal elements get introduced for the purpose of simplifying matters with talk about **D**.

[17]We will also discuss instrumentalism more generally in chapter 6.

3. The ideal elements are, by stipulation, the same kind of thing than the things in **D**. So, **D** is stipulated to be enlarged to some larger domain $\mathbf{D}'$.

4. Truth about $\mathbf{D}'$ is truth about **D** plus the stipulations that introduced the ideal elements, or truth as-if the stipulations were true.

I would like to point out that this does not presuppose that the domain **D** is a domain of existing objects. I makes perfect sense to say that when we say

(271) Two and two is/are four.

then we are talking about numbers, even though we do not talk about objects, in the sense that we do not attempt to refer to any entities.

So, the truth of statements involving the negative numbers is truth given the stipulations about them, or truth relative to these stipulations. This is contrary to truth about the natural numbers, which is truth simpliciter, independent of any stipulations. This is quite analogous to any talk involving ideal elements. However, there is something special about the negative numbers, both the negative whole and negative rational numbers. The truths about the natural (or positive rational) numbers plus the stipulations that introduce the negative numbers determine an answer to every question about the negative numbers. As we saw above, every question about the natural numbers has an objective answer.[18] All the facts about the natural numbers, plus the stipulations about the negative numbers determine every reasonable question about the negative numbers. These stipulative facts are something like:

1. Every number n has an inverse, $-n$.

2. $n + (-n) = 0$

3. Every negative number is the inverse of of a natural number.

In other words, the whole numbers are the group generated by the natural numbers. It should be clear that all quantifier free statements about the negative numbers, or the negative and the natural numbers, have a determinate truth value (simply calculate). Quantified

[18] Well, every reasonable question. The question "What's the smallest natural number n such that the continuum has the size of $\aleph_n$?" isn't a reasonable question about the natural numbers, even if there is such an n. It's a question about the continuum.

statements will again have truth conditions equivalent to infinite conjunctions and disjunctions of quantifier free statements. Therefore this carries over to the whole numbers, and the rational numbers. The reason why these generalizations should be understood as internal generalizations is simply because when we add the ideal elements we consider them to be of the same kind as the original elements, and since we have an canonical description of each one of them we simply express internal generalizations, as we always did with the natural numbers or the positive rationals.

5.5.3 Real numbers

Rational numbers were used for the longest time to model length and other spatial properties of regular objects. And, as it is contained in most mathematics textbooks, it was a shocking discovery for some people roughly 2000 years ago that this could not always be done. The length of the diagonal of a square with sides of length one is irrational. To overcome this, a system of numbers was sought that would not have these deficits and that would be more complete than the rationals. Besides the usual closure properties that the rationals have, these new numbers should at least also be closed under square roots of positive numbers, and the like. This is a little bit like the introduction of the negative numbers, where we wanted completion under subtraction, but this time the story is much more complicated and different in important respects. I can't begin to outline the complicated history that the real numbers had, and how this affects how we should conceive of them philosophically. I would simply like to outline how I think the real numbers differ in one important philosophical respect from the natural numbers and the rational numbers, and how they are similar in a certain other respect to the negative numbers. I will have to be rather sketchy and programmatic. To fully work this out will require a lot more work. But I would like to briefly discuss the real numbers anyways, to point out an important difference between them and the natural or rational numbers. I will develop this more fully on another occasion.

Besides ideal element instrumentalism there is another kind of instrumentalist stance.[19] According to this stance, which we could call *stipulative instrumentalism*, one does not simply add some elements to an already well understood discourse, but one simply stipulates that there are objects of a certain kind. In doing this one will start out with a description of the entities that one wants to stipulate to exist. On the basis of such a description one then

[19] Again, we will have a closer look at instrumentalism in chapter 6.

stipulates that there are entities satisfying the description. This is more or less a classic instrumentalist stance. For example, according to certain views of talk about theoretical entities in science, a scientific theory implicitly provides such a description of certain entities (i.e. it specifies the causal role they are supposed to have and the like). One way of having an instrumentalist stance towards them is simply to stipulate that there are entities that have the causal role the theory assigns to them. Now, I do not have much sympathy with instrumentalism in the philosophy of science. So, this is just an example to illustrate a certain kind of instrumentalism. A kind of instrumentalism like this I find very plausible when it comes to the use of certain entities in modeling, for example, features of natural language (see section 6.3.1 for more on this). There is an important difference between the ideal element instrumentalism that we encountered above in the case of the addition of negative numbers to certain number systems and stipulative instrumentalism. Quantified statements in the former case expressed internal generalizations, and this was plausible because a) quantified statements about the original number system (the natural or positive rational numbers) expressed internal generalizations, and b) the new ideal elements all had canonical names, so an extension of internal generalizations to them was easy. The infinite disjunctions and conjunctions that model the truth conditions of these sentences arise in a simple way out of the ones that model the truth conditions of the sentences expressing generalizations over the natural or rational numbers. Quantified statements in the case of stipulative instrumentalism, however, will usually not express internal generalizations, but external generalizations over the stipulated domain. We will have no access to these objects other than through the descriptions that were used in their stipulation. And this will not necessarily give us a canonical, or any other, name for each one of them. In such a case, when we say "every Φ is Ψ" then we are not necessarily saying something equivalent to an infinite conjunction. We will say something like "Every Φ, whatever these may be, is Ψ".

As you might guess, it seems to me that the real numbers get introduced via stipulative instrumentalism. We have a need for a number system with certain properties. This is not ideal element idealism, since the real numbers are not introduced to simplify calculations on the rational numbers. They are introduced, amongst others, to model spatial and temporal properties. How exactly this works is a bit tricky, and not at least because the real numbers were around long before the precise definitions we have of them now (say, as Cauchy completion of the rationals). Because I won't be able to get into this, I will just make the following claim without much argument. The stipulative definitions or descriptions of the real numbers will not completely determine every question one can reasonably ask about them. This will, to some extend, arise from the fact that quantification about them will have to be understood as expressing external generalizations over a stipulated domain. The real numbers will differ from the natural or rational numbers in this important respect. For the former number systems every question was determined, either independently of any stipulations, or given certain stipulations. For the real number we will get question that simply have no objective answer, even given the stipulations. Since all the facts there are about the real numbers are the ones that are determined by the stipulations that introduced them, there will be statements about them for which there is no fact of the matter whether or not they are true.

As I said, this is very sketchy, and we have to be a bit careful. It all depends on what we take to be the stipulative descriptions of the real numbers, and what we take to be a reasonable question. For example, if we ask for nothing else than that the real numbers are a real closed field and we allow only questions in the language of real closed fields then every question will have a determinate answer.[20] But this isn't enough to be all we ask from the reals. We usually require Cauchy completeness. And to see whether this will determine all questions we will have to understand whether or not questions about arbitrary infinite sequences of rationals are determined. It seems to me that the same considerations will apply here, but this is an issue that I would like to declare all this as going beyond the scope of this dissertation.

5.6 Conclusion

In this chapter I have defended a certain logicist view of arithmetic, according to which arithmetic truth is objective truth, but not truth about any particular objects. In fact, arithmetic truth has no ontological presuppositions whatsoever. This view presupposed the general view about quantification and nominalizations defended in chapters 2 and 3. We have also seen how the natural numbers relate to other number systems.

So far we have looked at quantification and nominalizations in general, and at the function of our talk about properties, propositions and natural numbers in particular. However, one of our overall goals was to shed some light on ontology, and to sort out some of the

[20] A result due to Tarski. The language of real closed fields is first order logic with the following non-logical symbols $+, \times, -, <, 0, 1$

confusion issues that go on in that discipline. In particular, we wanted to answer the puzzles w started out with. We will do this in the next chapter.

Chapter 6

Ontology

6.1 Introduction

After having looked at what we do when we talk about properties, propositions and numbers, it is now time to get back to some more large scale metaphysical issues. In this chapter we will look at how what we have seen so far relates to some classic questions in ontology and metaphysics. In particular, we will look at different projects within ontology, and how they relate to each other, at nominalism, at the problem of universals, at minimalism about truth, and general issues about objectivity, and at some alternative approaches that try to deal with the same or similar issues as the ones we looked at in this dissertation. But first, let's solve the puzzles we started out with.

6.2 The solution to the puzzles

In chapter 1 we saw that ontology has some puzzling features to it, and we isolated two particularly puzzling aspects about ontology. One arose from there being two apparently equally correct, but incompatible, answers to the question 'How hard are ontological questions?'. The other arose from there also being two apparently equally correct, but incompatible, answers to the question 'How important are ontological questions?'. For both of these puzzles we have restricted ourselves to considering them for properties, propositions and natural numbers. Here is how we can now resolve these puzzles.

6.2.1 Puzzle 1

The puzzle arose because the following two seem to be true:

(P1.A) Ontological questions, including those about properties, propositions and numbers, are difficult questions that can't be resolved by trivial arguments. They involve serious theoretical or other work to be resolved, or might be beyond our epistemic capabilities.

(P1.B) The question whether or not there are properties, propositions and numbers has a trivial affirmative answer in each case. One way to show that this is so arises from the trivial arguments in the following. Any ordinary statement like

(313) Fred is tall

trivially implies that

(314) Fred has the property of being tall

which directly answers our ontological question because it implies

(315) There is a property Fred has, namely being tall. Thus there are properties.

With what we have seen in chapters 2 and 3 we can now resolve this puzzle as follows:

- We have not seen any evidence against (P1.A). It seems to be completely justified. Ontological questions are questions about what the building blocks of reality are, and such questions aren't easy.
- However, we now know a lot more about the trivial arguments on which (P1.B) is based. We have seen that they are perfectly valid, as they seem, but that they nonetheless do not answer the ontological question we wanted answered.
- The trivial arguments are valid because
  - The first step, from (313) to (314), is a transition from a sentence that communicates a certain information with no particular structure or emphasis to a sentence that communicates the same information with structure and emphasis. They communicate the same information with a different structure. In particular, they have the same truth conditions and imply each other.

– The second step in the argument, from (314) to (315), is also perfectly and trivially valid. However, it is only so in one of two ways in which the quantifier that occurs in (315) can be understood. We have seen that quantifiers in general have both an internal and an external reading. The inference is trivial if the quantifier is understood in its internal reading, since in the internal reading it has a certain inferential role, independently of whether or not the noun phrase in (314) is referential, or otherwise denotes some entity. The inference isn't valid if the quantifier is understood externally.

Thus the first puzzle is resolved by endorsing (P1.A), and rejecting (P1.B). The latter is rejected even though the argument that motivates it is perfectly valid. It just doesn't show what it seems to show.

6.2.2 Puzzle 2

The second puzzle arose because the following two statements seem to be true:

(P2.A) Most of the things we say have ontological presuppositions for their literal and objective truth. In particular, most of the things we say in ordinary life, philosophy and mathematics presuppose that there are properties, propositions and natural numbers. Thus the question whether or not there really are such entities as properties, propositions and natural numbers is of the greatest importance to us. If the answer, God forbid, were that they do not exist then basically everything we ordinarily say would be literally and objectively false, and most of our talk in everyday life, philosophy and mathematics would have been shown to have false presuppositions.

(P2.B) In practice we do not take ontological worries or concerns about properties, propositions and natural numbers very seriously. We do not first ask whether there are numbers before we engage in arithmetic, or whether there are propositions before we ascribe beliefs, or whether there are properties before we ask whether every mental property supervenes on a physical one. In practice we seem to postpone these question and leave them to the specialist. Thus they are not very important to us, at least in practice.

We can now resolve this puzzle as follows. (P2.A) would be fully justified if our talk about properties, propositions and natural numbers would indeed have ontological presuppositions for its literal and objective truth. For example, much of the religious talk of the pope has God as an ontological presupposition for its literal and objective truth, and thus for the pope and fellow believers the existence of God is an issue of the greatest importance. But we have seen by now that this is, appearances to the contrary, not so for our talk about properties, propositions and natural numbers. Our talk about them often is literally and objectively true, but it does not have ontological presuppositions for its literal and objective truth. Thus the uncertainties of ontology are no threat to our practice of talking about properties, propositions and natural numbers, nor are they a threat to how we should evaluate this practice. Thus (P2.A) has to be rejected since it is based on a mistaken view about what ontological presuppositions this talk has.

Because of this it turns out that the natural attitude that many philosophers and non-philosophers have towards ontological questions are (in this case) completely justified. Ontological questions about properties, propositions and natural numbers are indeed not very important. And they indeed don't have to be addressed first before we use such talk in everyday life, philosophy and mathematics. In this case the natural attitude wins over Quine-inspired philosophical sophistication. But that, of course, doesn't mean that ontological questions aren't important on other occasions.

This solves the puzzles, in the sense in which it shows which one of the two apparently equally good but incompatible answers to these questions is correct and which one is mistaken. It doesn't solve the puzzle with respect to giving an account of why the puzzles where puzzling in the first place. We shall look at this next.

6.2.3 Why the puzzles are puzzling

A puzzle or paradox consists of apparently incompatible or contradictory statements that each by itself seems to be well justified. Such puzzles, antinomies and paradoxes come in two kinds. On the one hand there are those that can be resolved. In this case one can show which one of the parts that gave rise to the puzzle was really based on a mistake, and where one can show that once this mistake is noticed and taken into account the puzzle dissolves. Such puzzles have a *happy-face solution*.[1] On the other hand there might be puzzles or paradoxes that do not allow for such a solution. They are based on a real, not only apparent, contradiction. They might still be resolvable in a looser sense of the word.

[1] I borrow this terminology from (Schiffer 1996a).

For example, one might be able to show why it is not as bad as one might think that there is a real contradiction. This can be called an *unhappy-face solution* to a puzzle or paradox.

Any happy-face solution to any puzzle has to have two parts. On the one hand, it has to show which one of the apparently well justified statements that gave rise to the puzzle in the first place is really false. And on the other hand it has to show why we got fooled into thinking it is true. We can call these parts of solving a puzzle *resolving the puzzle* and *explaining the puzzle*. Any philosophical puzzle of interest will not be based on a simple obvious mistake, but rather on a mistake that is for some reason or other tricky to avoid. How did we get fooled into thinking that the answers to our questions that gave rise to the puzzle were were both good? We have to address this now to completely solve the above puzzles. Why did we get fooled into thinking that our talk about properties, propositions and numbers has ontological presuppositions for its literal and objective truth?

I think this is easily accounted for. There are several different aspects to why this false belief is often had. Let me list a few:

- Quantifiers certainly have a close connection to ontology in some of their uses, in particular in the formulation of ontological questions. The fact that they don't have such a connection on all of their uses can easily be overlooked because of the subtle role that context plays in determining the content of an utterance with a quantifier in it. Such cases of contextual contributions to content are hidden from the speakers. A competent speaker doesn't have to know that they are. In addition, these cases of contextual contributions are often tricky to discover, and it is often tricky to determine how many different readings a certain phrase can have. Consider, for example, reciprocals. How many different reading does the following sentence have:

  (316) The French and the Germans don't like each other.

  Even though I am aware of the semantic underspecification of reciprocals, it is not obvious to me how many reading this sentence has.

- A further reason is that non-referential, non-quantificational noun phrases are semantically more elusive than referential ones. The latter simply stand for some object. But the former have a much more complicated and tricky function. In particular, not all the expressions in that class have the same function.

- Last, but not least, the philosophical tradition has been very much influenced by the success of logic in this century. But a good deal of the influence that logic had on philosophy is restricted to standard first order logic. It was for the longest time the paradigm of a formal language within philosophy. And in first order logic quantifiers always make the same contribution to the truth-conditions, and every noun phrase, i.e. term, is either a referring expression, i.e. constant or variable, or a quantifier, i.e. description. This picture, and first order logic, had a clear influence on Quine and many of his followers. In particular, the role of context is rather limited from this point of view. It is in essence the extended simple picture we encounter in chapter 2. According to this picture, what context does is restrict the quantifiers and fix the value of the demonstratives, i.e. free variables. And with it in the background it doesn't seem too surprising that the important aspects that lead to a resolution of the puzzles were overlooked.

6.3 Natural, philosophical, theoretical, and true ontology

We looked at the puzzles in part to arrive at a better understanding of what is going on in ontology in general. What is the philosophical discipline of ontology doing, and what should it be doing? We saw in chapter 1 that we have to distinguish a number of different uses of the word 'ontology'. To briefly remind you, they were:

- Natural ontology: what is implied to exist by our shared beliefs.
- True ontology: what in fact does exist.
- Theoretical ontology: the things that some theory or other talks about.
- Philosophical ontology: the philosophical theorizing about some entities or other.

We now have to look back at this issue again to see how these four notions relate to each other. This will give us a better understanding of the different projects that are pursued in philosophy these days under the name ontology. Intuitively one might think the picture of the relation of these different ontologies is as follows:

- Theoretical ontology contains parts of natural ontology (after all, theories are also about some ordinary things), but properly extends it (through talk about theoretical entities). Also, natural ontology contains things that are too specific or too unimportant for any theory to be about.

- Both natural and theoretical ontology have true ontology as their ideal. Both try to be as close as possible to true ontology. In particular, all our beliefs about what there is aim to be true, and our taking recourse to entities in theoretical enterprises aims at only taking recourse to those that are contained in true ontology, i.e. really exist.

- Philosophical ontology also aims at true ontology. It tries to come up with philosophical theories about what there really is. To be sure, in practice it has to use those that are in our natural and theoretical ontologies, but that is only because these provide our best access to true ontology.

There is no doubt that the relation between natural and theoretical ontology is correct as described. Natural ontology doesn't contain the things that complicated theories talk about, and no theories talks about some ordinary things that are in our natural ontology, like bread crumbs. To be sure, bread crumbs are theoretically covered by a number of theories, including those that deal with physical objects generally. But no theory will imply that there are bread crumbs, even though some will imply that there are dogs (under the name *canis familiaris*).

But the other two points are false. Neither philosophical nor theoretical ontology have true ontology as their ideal. Their connection is much looser. I will defend this point in this section, in particular by looking at properties and propositions again. This will be important to understand the overall picture that the view developed in this dissertation so far gives rise to. One might think, for example, that if one claims that properties and propositions, understood as whatever 'being F' and that-clauses stand for, don't exist then this is a theoretically restrictive position because one thereby believes that theories shouldn't take recourse to them and philosophers should stop making theories about them. I don't believe this at all. In fact, I think that theories are well advised to take recourse to them, and philosophical debates about propositions are most useful and desirable. To say this in no way contradicts anything I have said so far. To see how that can be, let's look at what use we have for properties and propositions in theories, and at philosophical theories about them. Properties and propositions have their main theoretical use in natural language semantics. What their use is in this discipline is what we shall look at next. This will help us understand what the relation is between theoretical ontology and true ontology. After that we shall look at how we should understand certain philosophical theories about properties and propositions, in particular the ones proposed under the heading of formal ontology. This will help us understand what the relation is that philosophical ontology has to all of natural, theoretical, and true ontology.

6.3.1 Properties and propositions in semantics

Natural language semantics

In natural language semantics properties and propositions play a prominent role. The role they play is that they are assigned as semantic values to certain phrases. Semantic values play an important role in semantics. They are used to represent the contribution that the phrase makes to the overall semantic value of the sentence. In particular, in a truth-theoretic semantics they are used to represent the contribution that the phrase makes to the truth conditions. Such semantic theories will try to assign to each phrase that occurs in a sentence some semantic value such that the truth conditions of an utterance of that sentence (in a particular context) can be recovered from the semantic values of its parts, the way they are combined, and certain information about the context of an utterance. In addition, the same semantic value for a particular phrase should be used in all sentences in which it can occur. This gives rise to a *compositional truth-theoretic semantics*. To specify such a semantic theory for a language one has to

1. Assign semantic values to the syntactically basic expression of the language.

2. Specify operations on semantic values and assign to each formula building syntactic operation one such operation on semantic values. Thus complex phrases get assigned semantic values on the basis of the semantic values that the primitive expressions that constitute them got assigned and on the basis of the operations on semantic values that the syntactic operations got assigned.

3. Recover the truth conditions of the sentences from the semantic values that they get assigned (or from the semantic values of their constitutive parts).

4. Do all this, so to speak, from a finite basis.

There is no question that properties and propositions play the role of semantic values in contemporary attempts to give a compositional truth-theoretic semantics for natural languages. What we will have to look at here is

a) whether or not this is compatible with an internalist view of talk about properties and propositions

b) whether or not this gives rise to evidence that properties and propositions exist

It seems that the answer to question a) should be that internalism and using properties and propositions as semantic values are incompatible, since it seems right to say that when we give a semantics for a natural language we think of properties and propositions as belonging to some domain of possible semantic values, and we assign to individual expressions members of that domain. And it seems that the answer to question b) should be that this is evidence for the existence of such entities, since after all, we have reason to believe that our best theories are true.

I think things are a bit more tricky. If we look at what use we make of the assignment of properties and propositions as semantic values we will see that a somewhat different view is right. And according to it the answer to questions a) and b) will be just the opposite.

Instrumentalism

The mere fact that we talk about certain entities in a certain theory doesn't mean that we should believe that these entities exist, even if the theory is our best theory. It also matters how and for what purpose we talk about these entities. The issue that matters here is whether or not an instrumentalist stance towards these entities is appropriate. To see how this question should be answered in the case of semantic values, we first have to look at instrumentalism in general, and then at semantic value instrumentalism.

Instrumentalism is mostly discussed in connection with theoretical entities of physics.[2] But instrumentalism is really a much broader view that may or may not apply to particular entities that particular theories take recourse to. As it turns out, the most widely discussed form of instrumentalism is rather implausible. Instrumentalism about unobservable entities in physics depends on what the goal of a physical theory is. Is a physical theory supposed

[2]See (van Fraassen 1980) for a defense of such a view.

to give a true description of the world, or an empirically adequate account of the observable phenomena. If one agrees with the first one will have to reject instrumentalism about unobservable entities. But if one agrees with the second then one should endorse it. After all, if all a physical theory tries to do is save the phenomena, then there can't be a correct answer to what auxiliary entities one should use. If saving the phenomena is all we are after then whether or not there really are electrons doesn't matter. But saving the phenomena isn't all we are after in physics. We sure want to get the phenomena right, but we also would like to have a true account of what lies behind them, what produces them, and what the world consists off at the smallest level.

Even though instrumentalism about unobservable entities in physics is implausible, that doesn't refute instrumentalism about all theories, or all entities to which some theory or other might take recourse. Whether or not an instrumentalist stance towards a certain kind of entity that a certain theory takes recourse to is legitimate will depend on what the goal of the theory is and what role the entities in question play in achieving that goal. For example, if it would be right that all that physics is trying to do is to have a simple and systematic account of the observable, and that all that electrons do in physics is to help us in having such a simple and systematic account then whether or not there really were any electrons wouldn't matter. Paradoxical as it may sound, in this case electrons could do all the work they are doing without having to exist. Of course, sentences of the theory that talk about electrons wouldn't be literally true. But that wouldn't matter in this case, because they aren't supposed to be literally true. They are just supposed to help in giving a simple and systematic account of the observable, and statements about the observable are supposed to be literally true.

Now, even though instrumentalism about the unobservable entities of physics is implausible, it isn't implausible because instrumentalism in principle is implausible. It is implausible because it gets the aim of physics wrong. To see whether or not an instrumentalist stance towards a certain kind of entity, and a certain part of a theory that talks about these entities, is appropriate, we have to see what we are trying to do with that theory, and how use these entities in doing this. In particular, we have to ask ourselves if it matters for the theory reaching the goal it is supposed to reach whether or not the entities in question exist. Let's look at this in the case of entities used as semantic values.

Semantic value instrumentalism

In discussing semantic value instrumentalism it is important to distinguish it from some other, clearly false, views. Even if one is an instrumentalist about semantic values that does in no way mean that one is an instrumentalist about semantics in general. For example, I take it to be beyond reasonable doubt that there is a fact of the matter what the truth conditions of an utterance are in a certain context.[3] Semantic value instrumentalism doesn't mean that the assignment of truth conditions is merely a useful tool serving some other purpose. There are correct and incorrect assignments of truth conditions, not just useful and unuseful ones. Furthermore, it is a fact of the matter whether or not an expression is a referential expression, and if so, what it refers to.[4] However, this, too, does not settle the question whether or not semantic value instrumentalism is correct. Even though it is a fact of the matter what the reference of a certain expression is, this does not determine whether it is also a fact of the matter what its semantic value is. As we have briefly discussed in chapter 3, referring expressions don't have to have their referents as their semantic value. It might for theoretical purposes be better to take something else than the referent as the semantic value of an expression. For example, instead of taking Fred to be the semantic value of 'Fred' one can take some other thing that is closely related to Fred, and from which the information can be recovered that Fred is contributed to the truth conditions by 'Fred'. This works if we take the set of all properties that Fred has to be the semantic value of 'Fred', and doing this is in no contradiction with claiming that 'Fred' refers to Fred, not the set of all properties of Fred. So, realism about reference and truth conditions doesn't settle the issue about instrumentalism about semantic values.

I think that if we look at what use we make of semantic values in a compositional truth theoretic semantics we can see that instrumentalism about semantic values is true. The use we make of them is such that the goals we pursue through their use can be successful independently of the existence of the semantic values. To see this, lets see what we want from a compositional semantics, and what doesn't matter for getting what we want from it.

When trying to give a compositional truth-theoretic semantics for a natural language

[3]This doesn't mean that they have to be completely specified. The might be underspecified, but that doesn't mean that there isn't a fact of the matter about what they are.

[4]This is not uncontroversial, but I firmly believe it. See (Davidson 1984a) for a contrary opinion.

we try to assign semantic values to the basic parts of the language, and operations on these semantic values to the syntactic ways of combining the parts such that the truth conditions of an utterance of a sentence of that language can be recovered from the information what the semantic values of the parts are, how they were put together, and information about the context of the utterance. What do we want from the semantic values in doing this? Well, that they allow us to do just that, that they are such that we can recover certain information about the truth conditions of the sentences from knowing that the semantic values of the parts are and how they were combined. Do we also, beyond this, try to assign the correct semantic value?

To be sure, there are correct and incorrect semantic values *within a given semantic theory*. To assign to a proper name its reference as its semantic value in Montague grammar is incorrect, since in Montague grammar all noun phrases have sets of properties as their semantic value. However, in other semantic theories they have their referent as their semantic value. Can there be a legitimate question about what the correct semantic value is, i.e. which one of these two theories assigns the correct semantic value, and which one assigns the incorrect one, not within a semantic theory, but simpliciter? I don't think so. Semantic values are used for a certain purpose. The are not themselves the domain which a semantic theory is about. A semantic theory tries to generate the truth conditions of sentences of a natural language in a compositional way. Semantic values are taken recourse to to allow us to do this.

Of course, for this to work the things that are supposed to play the semantic values have to satisfy some conditions. They have to be specified well enough so that they can perform this function. The have to allow for a recovery of what they contribute to the truth conditions of the sentence in which they occur. They have to be specified enough so that it will be clear which ones are different and which ones are the same. And that is pretty much it. Whether or not there really are such things as the ones we specify as semantic values doesn't matter. Just the same way as it doesn't matter for the instrumentalist about electrons whether or not there really are such things that satisfy the description that electrons are supposed to satisfy.

To illustrate this, consider the following hypothesis. Suppose you had the most reliable information that the type theoretic entities that are used as semantic values in Montague grammar don't exist. How would this new and firm belief of yours affect your attitude towards Montague grammar? Well, it would depend on why you think these type theoretic entities don't exist. If you believe they do not exist because the description you have of them is inconsistent, then you would rightly say that Montague grammar has been based on a mistake. It used things as semantic value that have an inconsistent definition. But if you believe that these things are perfectly coherent, they just do not exist, then this would have a quite different effect. Here we wouldn't say that Montague grammar has false presuppositions. We wouldn't be much concerned at all. So these things don't exist, but we can still do Montague grammar as we always did. No contradiction will follow. We can still assign type theoretic entities as semantic values, and we can still use them as we always did. What this thought experiment makes plausible is that we never in the first place assumed or presupposed that the things that we use as semantic values are real things. For the way in which we use them it is only required that they are well defined and clearly specified. Whether or not there really are such things doesn't matter.

So, using properties and propositions as semantic values doesn't require that one believes that they exist, nor does their successful use in semantics provide any evidence that they exist. In particular, to believe that properties and propositions don't exist does in no way mean that one is on conflict with contemporary semantic practice, or that one is restricting oneself for ontological reasons in ones theorizing in semantics.

Numbers, too, play a role in the theoretical ontology of many theories. However, the case of numbers is a lot more complicated than the case of properties and propositions in this respect. One the one hand, real numbers are much more popular in non-mathematical theories than natural numbers, and we have barely started to talk about them in chapter 5. In addition, the role that mathematics, and numbers in particular, play in scientific theories is too complicated to be dealt with here, and I will have to postpone the discussion of this to another occasion.

So, by considering instrumentalism, and in particular by noting that instrumentalism sometimes is the right stance towards a certain use of certain entities in certain theories, we can see that the ideal limit of theoretical ontology isn't true ontology. The theoretical ontologies of some theories are comfortably independent of true ontology, because they use these entities in an auxiliary way in which it doesn't matter for the overall success of the theory whether or not talk about them is literally true.

6.3.2 Formal ontology

Nowadays there is a whole subdiscipline in metaphysics called formal ontology. People who are active in it try to come up with axiomatic mathematical theories about entities of a certain kind. For example, they try to establish what closure conditions the class of events satisfies, or how propositions can be seen as set theoretic entities of a kind, and the like. How should we look at formal ontology? One way to look at it, and this is the way people who are engaged in formal ontology like to look at it, is that formal ontology tries to give a precise theory about certain entities that are part of realty. Formal ontology is a philosophical discipline that tries to give a precise mathematical theory about a particular aspect of reality that isn't dealt with much by the empirical sciences. Formal ontologist often state their mission as follows:

> We metaphysicians have discovered that the world is a lot richer than what one might naively believe. Our natural and theoretical ontology contains all kinds of objects, most of them abstract, some of them concrete. But our ordinary scientific theories, except mathematics, are only really concerned with the concrete ones. Besides these, however, there are large numbers of other kinds of entities. We need theories about them, too. And we formal ontologists do just that. We give precise mathematical theories about all these other objects that are ignored in the other parts of science. And that is the only essential difference: physics makes theories about unobservable concrete entities, we make theories about unobservable abstract entities. And just as in physics, we have reason to believe that our best theories are true.

I think this way of looking at formal ontology is much too narrow, and in fact quite misguided. The value of formal ontology is quite independent of whether or not the things that such a theory is about exist. And the success of a formal ontology provides no evidence that the things the theories tries to give a mathematical model of exist.

Consider our use of propositions in natural language semantics, and formal ontology approaches to what propositions are. Assuming the semantic value instrumentalism I defended in the last section, it will not matter for what semantics wants to do with propositions whether or not there really are such things. All that will matter is that they can play the role they are supposed to play as semantic values. But for them to play this role we have to have a precise specification of what they are like. Now, an account offered by a formal ontology that gives a precise mathematical model of propositions is exactly what we want here. We want to have precise conditions when propositions are the same or different, how propositions can be build up from the semantic values of the parts that occur in the content sentence of the propositions and the like. So, a formal model of propositions, taking them to be, say, structured entities that contain certain constituents, is of great value to a semantic value instrumentalist. As it turns out, semantic value instrumentalism is true, and thus the value of formal ontology isn't limited to giving a true mathematical model of certain entities that are part of reality.

In addition, the success of a formal ontology shouldn't be taken to suggest that the entities it gives a formal model of exist. Take the case of propositions again. Suppose we have a formal model of propositions that is such that it gets sameness and difference of propositions just right. To say this is to say

- the formal ontology assigns certain precisely characterized things as the values of that-clauses, such that
- the assignment of these things to that-clauses is such that two that-clauses get assigned the same thing iff they are replaceable in sentences without change of truth conditions. For example, if the assignment to 'that p' is the same as to 'that q' then x believes that p iff x believes that q.

Such conditions of formal ontologies, however, do not suggest that success in achieving them means that the things that the formal models are about are real. The can just as well be the things that an instrumentalist semantics takes recourse to.

Formal ontology is a valuable discipline. But one shouldn't be too quick to think that it is giving a mathematical model of real features of reality. On the other hand, what reality contains is no limit for what formal ontology should theorize about.

6.4 Nominalism

I have argued above that expressions that expressions like 'being F' and that-clauses are not referential expressions, and thus if properties and propositions are nothing over and above what such expressions refer to then properties and propositions do not exist. Reality does not contain such entities. This, of course, does not mean that reality does not contain any abstract objects. Everything I have said so far is perfectly consistent with there being lots and lots of abstract objects. However, nothing I have said requires there to be any abstract objects either. Even though what I have said so far is no defense of nominalism broadly construed, it is in certain respects related to nominalism. To see this more clearly we shall look at nominalism more closely in this section.

6.4.1 The varieties of nominalism

There are a variety of nominalistic positions. On the one hand, nominalistic positions differ in how broad and far-reaching their claims are. Most radically, there is what we can call **global nominalism**, which is the claim that everything there is is concrete, and nothing is abstract. Global nominalism has to rely on something like an argument that there is something wrong with abstract objects in general. Even though most anti-nominalists take nominalism to be this claim,[5] most nominalistically inclined people wouldn't want to defend anything this strong. After all, it doesn't seem clear what is wrong with being abstract in principle. What does seem fishy to nominalistically inclined people is that we would ever talk about any abstract objects. In particular, nominalists defend what we can call **local nominalism**, that is the claim that particular kinds of things are not abstract objects. Thus (local) nominalists about numbers defend the claim that numbers are not abstract objects, and (local) nominalists about properties defend the claim that properties are not abstract objects. There are a variety of ways in which such local nominalisms can be defended. We can distinguish at least the following (I will explain them in the case of properties only, the other cases carry over in an obvious way):

- **Eliminative nominalism** Talk about properties has to be given up, since it presupposes that there are certain things, properties, that there aren't. Since talk about properties has false presuppositions it has to be abandoned, at least when we are engaged in anything other than loose and popular talk. This view is the view of Quine in (Quine 1980).

- **Reductive nominalism** Talk about properties is not to be abandoned, but it has to be translated or reduced away into talk about concrete objects only. This can be done claiming that talk about properties can be translated away into talk about cars,

[5] See (Burgess and Rosen 1997).

trees, and other concrete entities. Such an approach will usually try to associate with each sentence that talks about properties one that might look quite different on the face of it, but is claimed to be truth conditionally equivalent to the former. Because of this the claim is that property talk is acceptable to the nominalist and should be understood as being a brief and useful way of really saying what its nominalistic equivalent says. For example,

(317) Being a philosopher is fun

might be claimed to be equivalent to

(318) Every philosopher has fun.

or something like this.

- **Non-reductive, non-eliminative nominalism** According to this approach talk about properties and propositions is not to be eliminated, nor does it have to or should it be translated away. Nonetheless, it is claimed, is it compatible with the claim that properties are not abstract objects. According to this view it is to be accepted that there is a property we have in common, and this is neither to be eliminated nor to be translated away. Nonetheless, this doesn't imply that any abstract entities that are properties exist. This view can be defended in several different ways:

  - By claiming that properties are concrete entities, like mereological sums of tropes, or the like. According to this strategy the grammar and semantics of a sentence that talks about properties will be left as it seems to be, but the domain that this talk is about will be claimed to consist only of concrete entities.
  - By claiming that when we say that there is something we have in common we are not making a literally true statement, even though we are making a statement that is to be accepted and not to be eliminated. This can be so because

    * the statement we make is to be understood in expressivist terms. When we talk about properties we are not making factual statements, but we are expressing something.
    * the statement we make is not literally true, but it is true given the property fiction. Thus given certain background assumptions about properties, there

is a property we have in common. However, speaking literally, there are no properties.

- By claiming that our talk about properties is literally true, but it isn't about any domain of entities at all. Because it isn't about any entities at all, it doesn't imply or presuppose that properties and propositions are abstract entities. This is, of course, the view defended in this dissertation.

All forms of nominalism except, surprise, the one defended in this dissertation strike me as implausible. However, I won't get into a discussion of them in any details except about the case of that talk about properties is only true within the property fiction, which will be discussed below, in section 6.9. I just wanted to point out how many different forms of nominalism there are. In particular, arguments against some of them won't apply to others. For example, it has often been argued that nominalism has to be false because some talk about properties can't be translated away.[6] But this of course only applies to reductive nominalism, which is just one of many kinds.

6.4.2 A world without properties?

Nominalism about properties is true. Properties are not abstract entities, because properties are no entities at all. Our talk about properties was never supposed to be about some domain of entities. We have seen evidence for this claim by looking at what we do when we talk about properties. But how does this look from the metaphysical side? What is the world made of if there aren't any properties? Are there just individuals, and that's it?

I, of course, have no worked out view what the world is made of. But the following view about the material world seems to make perfect sense to me: the world consists of individuals having properties and bearing relations to each other. In a sense there are just individuals, but this shouldn't be understood as that there are only individuals independent of the properties and relations they have. The world consists of red apples and fast cars, not apples, cars, the property of being red, the relation of instantiating a property, etc. etc..

It is a mistake to try to separate the things that have properties from the properties they have. One extreme way to make this mistake is to think that the world consists of properties and property-less bare particulars that have the properties. But this mistake is

[6] See (Jackson 1977) for an argument to this effect.

natural to make if one thinks that properties are entities. If a red apple really is built up from two things, the apple and in addition the property of being red that the apple has a certain relation to then it seems reasonable to ask what each one of the relata is, the apple and the property of being red. The apple by itself then gets separated from its being red, and all of its other properties by a similar reasoning. What we are left with is a bare particular and a number of properties that it relates to. There are other ways of making similar mistakes, but we have seen how to avoid them by now. A red apple isn't made up of two things, a property and a thing that has the property. There is only one thing, an apple, that has properties. But having a property isn't having a relation to some other entity. Things have properties. But from an ontological point of view only one kind of entity is involved here, the things that have the properties.

So, the present view in no way denies that reality contains things that have properties. But it denies that the move from that to the question what these properties are, how they relate to the things that have them, etc., is a legitimate move. This move is based on the mistake of thinking that when we say that a thing has a property then we are really talking about two entities, the thing that has the property, and the property it has.

6.4.3 Content without propositions?

If nominalism of the kind I defend about propositions is true then propositions do not exist, at least not if we assume that propositions are whatever that-clauses stand for, or refer to. But what does this say about content? Can one be a realist about content without believing that propositions are among the things that make up reality? Can beliefs have content without relating to some entity that is the content of the belief?

Realism about content is not at all threatened by the view defended here. In fact, I think content makes much more sense if it doesn't consist in having a relation to some abstract object. Even if you believe that if one has a belief with content that p then one has a relation to some entity which is that p, you won't believe that one has this relation to this thing simpliciter. That is, there will be other things about you that make it that you have this relation to this proposition rather than another (and that you have this relation, rather than another). What these other things about you are that make it that you relate to, say, the proposition that Fido is a dog is your physical relationship to a certain dog, your relationship to dogs in general, your being a member of a certain language community, and many other things.[7] And it will be because of all that you will have a belief with that content. Your being related to a proposition (we are still assuming for the moment that propositions are entities) will be secondary to that and, in a sense, a result of it. The basis of content is to be found in all these other things.

So, that propositions are entities isn't central to realism about content. It only follows from realism together with the claim that that-clauses are referring expressions. What is central, though, is that ascriptions of content are literally true. But with the view developed above we have no problem with that. That-clauses are non-referring expressions that specify content, but don't refer to something that is the content. So, that-clauses specify what the content of a belief is with the content sentence that is part of a that-clause. This isn't a substantial theory, to be sure, but it is on the face of it no worse than saying that that-clauses refer to things that are the content of the belief. This latter view is easily, but falsely, motivated by looking at quantifier inferences and a number of other considerations that we by now have seen to be based on confusions.

6.5 Deflationary truth

What has been said so far is quite relevant for several discussions about topics outside of ontology proper. One of them is the debate about the function of a truth predicate, in particular about minimalism about truth, or deflationary truth. I won't get into the details here, but I'd like to point out that the following two problems about truth are directly affected by what has been said the above:

- It seems to be a triviality, maybe even a conceptual truth, that

  That p is true iff p.

  But how can that be so obvious, since on the left hand side we seem to ascribe some property to some entity that we don't talk about on the right hand side?

- Deflationists about truth claim that the function of a truth predicate is metaphysically thin and merely to say what has been said before, or to get certain increased expressive power. However, both Horwich, in (Horwich 1990), and Field, in (Field 1994a), take recourse to rather complicated formulations of this to accommodate that propositions

[7] See (Field 1978) for an outline of a picture of how physical beings can relate to abstract propositions.

are language independent entities that are not all expressible in contemporary English. For example, Horwich thinks that the correct theory of truth wiIl have infinitely many axioms not all of which can be expressed in English. This gives rise to the question in what sense such a theory of truth can be called minimalist. Is this complexity really necessary?

These two problems can be seen in a different light because of what we have seen above:

- The truth equivalences are just as obvious as the ones in the cleft constructions, since we are not talking about any other entities on the left hand side, but rather take recourse to a different way of articulating the same information as with the right hand side.
- Inexpressible propositions can be accommodated within a certain simple picture of the expressive power that we get from quantification over propositions. The source of the expressive power that we get through a truth predicate, applied to propositions, can be seen as arising quantification over propositions. No taking recourse to language independent propositions is required to accommodate this.

Of course, this is no defense of a deflationary conception of truth, even when we restrict ourselves to ascriptions of truth to propositions. I am inclined to believe that deflationism isn't true across the board. It seems to me that a truth predicate is used in two ways, related to a distinction made in chapter 1:

1. That p is true iff that p is correct according to the standards of correctness that are the relevant ones on a certain occasion.
2. That p is true iff that p is correct according to certain special standards of correctness, ones that require some kind of direct relation between language and world .

In the first sense, it will be true that Sherlock Holmes is a detective, and in the second it will be false. It is quite plausible to assume that the primary function that a truth predicate has is to help with the increased expressive power that we get through quantification over properties. This is indeed metaphysically thin, given the present account of quantification over properties. A secondary function, one that the truth predicate somehow got recruited to perform, is to make the distinction between

- a statement living up to the standards of correctness that are the ones that it is supposed to live up to on the occasion of its utterance, whatever these standards may be on that occasion, and
- a statement living up to special, somewhat restrictive, standards of correctness.

That all this happens while we ascribe truth to propositions, rather than utterances, is no problem, by the way, since the context of the individual ascription will matter. In particular, the following can be uttered both truly and falsely:

(319) It's true that Sherlock is a detective.

This can only be the roughest sketch, of course, and I leave the details for another occasion.

6.6 The problem of universals

The classic problem in ontology is the problem of universals. It is probably the oldest clearly formulated and controversial question in ontology, and it is only one of a few that is widely recognized within philosophy as an important philosophical problem. Even though the problem of universals can be stated in just a sentence or two, and motivated with just a few sentences, it is a problem that is quite different in character than other philosophical problems. The most striking difference is that the problem of universals, contrary to almost all other major philosophical problems, is extremely hard to motivate to non-philosophers, even non-philosophers who otherwise are very much open to philosophical problems. Try to persuade your colleague in, say, mathematics or economics that it is an important problem in philosophy to find out if there is something we have in common, and if yes, what that thing is and how it relates to us. Is it fully present in both of us? Is it abstract or concrete? This problem often strikes non-philosophers as somehow wrong headed. Other philosophical problems don't have this feature. The problem about the existence of God, the nature of morality, the mind body problem, freedom of the will, etc., all can be well motivated to non-philosophers, and they commonly do strike non-philosophers as serious problems. Not so the problem of universals.

This is not only my personal opinion that I have because it fits my view of this issue. It is also held by one of the most prominent present philosophers working on the problem of universals, David Armstrong. Here is a quote from him:

> "The problem of universals has the interesting characteristic that it is almost impossible to explain to a non-philosopher what all the fuss is about. It is truly philosopher's philosophy. Perhaps that should make us suspicious of it. Yet I believe that Plato's instinct was correct when he treated it as the central question in metaphysics." (Armstrong 1984, 41)

I disagree with Armstrong about the second part of this quote, but agree about the first part. The problem of universals is not a central problem in philosophy, or ontology. It is based on a confusion. And I think the confusion is such that it nicely accounts why philosophers would take it to be a serious problem, but non-philosophers can't get themselves to see what all the fuss is about. The problem of universals is motivated with a very compelling argument. But the way in which it is compelling is such that it will only be compelling to philosophers, or a certain kind of philosophers. The standard move to motivate the problem of universals is based on the following reasoning:

1. Fido is a dog. Fifi is a dog.

2. So, there is something Fido and Fifi have in common.

3. So, there is some thing, or entity, that they have in common.

4. What is that thing, and how can they have it in common?

To be sure, usually there are a lot more bells and whistles in between, but this way of reasoning is the backbone of motivating the problem of universals as a philosophical problem. By now we have seen what is going on in this kind of reasoning. The first step is perfectly valid. There are lots of things we have in common, and so do Fido and Fifi. But the second step in this argument is a mistake. From the fact that there is something we have in common it doesn't follow that there is a thing or entity that we have in common. This would only be valid if the quantifier in 2. were an external quantifier. However, the inference from 1. to 2. is only valid if the quantifier is understood internally. In particular, the questions asked in 4. assume that the quantifier in 3. is external.

Even though main line of reasoning that motivates the problem of universals is based on a mistake, its the kind of mistake that is easy to make. It seems that there is a perfectly valid form of reasoning that starts with nothing controversial and ends up with the claim that when we have something in common we have some kind of relationship to some other thing. Philosophers are trained to follow where-ever reason takes them, and they are quite willing to to consider that the world contains some special entities that we relate to when we have something in common. These arguments lead us down the wrong path, but as trained philosophers we are often willing to follow where we think we have to go, even if the conclusions we end up with seem really quite strange or a bit absurd. Non-philosophers are often much less willing to follow where they don't want to go.

The problem of universals is based on a confusion. The confusion arises from overlooking that quantifiers are used in more than one way. With the view developed in this dissertation we can look at the problem of universals the following way. The problem of universals is the problem of answering the following question:

The Problem Is there something we have in common?

We now know that this question can be understood in more than one way. Given this and some other things that have been said above we can give the answer to the question.

The Solution Yes and No, depending on how you understand the question. Yes, trivially, if the quantifier in the question is understood internally. This answer to the question follows immediately from the fact that we are all human. No, if the quantifier in the question is understood externally. Here I assume that whatever it would be that we, externally, have in common is whatever expressions like 'being human' would stand for or refer to. Since we now know that these expressions never in the first place try to refer to anything we now know that there are no entities that they stand for and which we have in common.

Even though the problem of universals is based on a confusion and a mistake this should not be taken to mean that ontology, or metaphysics, should be dealt with by trying to find the mistakes that underlie these problems. Most metaphysical and ontological questions discussed in philosophy are not based on mistakes or a confusion. The problem about the existence of God, the existence of immaterial minds, the nature of morality, etc., are all perfectly legitimate serious questions that are in no way wiped of the table by pointing to confusions in the use of language. Most metaphysical problems aren't like that, but the problem of universals is based on a confusion about language. Just one out of many problems, not metaphysics in general.

6.7 Objectivity without objects

It is natural, and even suggested by the words, to closely tie objectivity to the existence of certain objects, or entities. For a domain of discourse to be objective it has to be about objects existing independently from us, or so one might think. In addition, the properties ascribed to these objects have to exist. As we saw in chapter 1, it is not uncommon to closely associate the objectivity of a certain domain of discourse with the existence of certain properties. Ethics is objective iff there are ethical properties. Aesthetics is objective iff there are aesthetic properties. Furthermore, there is a natural tendency to conflate objectivity with objective completeness. This way of looking at objectivity, by closely tying it to ontology, naturally gives rise to the view that if we have objectivity at all in a domain of discourse then we have to have objective completeness in it. If the objectivity of morality comes from which moral properties are instantiated, then it seems that every moral question should have an objective answer. It is true if the moral property ascribed to an action or person is in fact had by that action or person, and false otherwise. There seems to be no room for any other options. Similar considerations carry over to tying mathematical objectivity to mathematical objects, and the objective completeness that seems to follow from that. However, in ethics as in the philosophy of mathematics there seem to be cases were this isn't so. Some moral seem to admit of no objective answer, and some mathematical problems don't seem to have one either. This gives rise to a dilemma, since it seems that if one ties ontology to objectivity in the above way then one can't have both

- that there is objectivity at all in a domain of discourse, and
- that there is no objective completeness in that domain.

If objectivity comes from the existence of certain objects then it seems objective completeness follows. However, by now we have seen that this picture is too simple. We have by now seen that a number of parts of the above picture are mistaken. The include the following:

- We have encountered cases where we do have objectivity, in the strongest sense of the word, going beyond mere intersubjective agreement, but no objects. We have seen that we have objective truth in arithmetic, but that arithmetic has no ontological presuppositions for its objective truth. And we have seen that talk about properties and propositions is objectively true, but also doesn't have ontological presuppositions for its objective truth.

- We have seen that saying there is objectivity in a domain of discourse iff the properties ascribed in that discourse exist, is a mistake, since it is based on the wrong conception of what talk about properties is all about. Talk about properties never is talk about some domain of entities, and thus the existence or non-existence of certain entities that this talk is supposed to be about can't distinguish objective from non-objective domains of discourse. This way of looking at objectivity is based on an externalist conception of talk about properties, which we have seen to be a mistake.

- Separating ontology from objectivity gives rise to the possibility of there being objectivity without objective completeness. We have not really seen a case of this, but we have touched on it in the discussion about the real numbers. This way of looking at things only opens the door to a plausible account of the objective incompleteness in certain (objective) domains of discourse. We haven't seen a case of a domain of discourse where this holds, but the above shows that this is possible and conceivable. Whether it actually applies to a certain domain of discourse will depend on lots of details about that discourse.

6.8 Neo-Carnapianism about ontology

6.8.1 A kind of neo-Carnapianism

In chapter 1 I outlined how Carnap would have been able to solve the puzzles from chapter 1 if his position about ontology would have been defensible. But Carnap's position, as Carnap held it, is quite problematic. Carnap thought that ontological questions are without cognitive content, and based on misuse of language. They are not fully factual questions. Carnap based this view on his verificationist conception of language, which hardly anyone believes these days, and rightly so, it seems. But Carnap still was onto something important. Carnap thought that there are two ways to talk about what there is. According to one of them, where we ask internal questions (in his terminology), the answers will be trivially affirmative. If one internally asks whether or not there are Fs then the answer will trivially be 'yes'. The reason for Carnap for this was that these questions were internal to a linguistic frameworks, and in any framework where you can ask the question whether or not there any Fs the answer will be 'yes' since according to the framework there will be Fs. There would be no use for talk about Fs from within the framework if according to the framework there wouldn't be any Fs. But Carnap said that this can't be what philosophers are after when they want to do ontology. It can't be, in part, because the answers to these questions that the philosophers are after can't be trivial. What the philosophers are after, according to Carnap, is to try to find out whether or not the framework corresponds to reality. And this is asked with external questions. But these external questions go beyond what can be meaningfully said. According to Carnap external questions, and in particular the philosophical discipline of ontology, are meaningless and based on a mistake. A typical philosopher's attempt to try to say more than what can be said, just like all of metaphysics.

I agree with Carnap that there are two ways to ask what there is. But these two ways are not based on anything more tricky and substantial than that in ordinary communication quantifiers are used in two different ways. Since questions about what there is involve quantifiers, the fact that quantifiers can be used in two different ways means that questions about what there is can be used in two different ways. This gives rise to that there are internal and external questions about what there is. However, contrary to Carnap I don't think that there is any difference between the two when it comes to being factual, or meaningful. Both internal and external quantifiers (in my sense, not Carnap's) are completely factual, and completely meaningful. The only difference between the two is what they contribute to the truth conditions of the utterance in which they occur. So, ontological question, questions formulated with an external quantifier, are completely factual and meaningful. In particular, ontology, the philosophical discipline, is perfectly meaningful. It is not based on some mistake, nor is the rest of metaphysics. Carnap was mistaken in thinking that metaphysics is impossible, or devoid of factual content. But he was right in thinking that there are two quite different questions we might be asking when we ask what there is. One is (usually) trivially answered with 'yes', the other is not trivial, but also completely factual and meaningful. It is only a difficult question. It deals with what kinds of things the world is made up from, and this is just not that easy to answer. It might be beyond our epistemic capabilities to answer it in certain cases, but even that does in no way mean that these questions are meaningless.

6.8.2 The methodology principle

One of the main upshots that the kind of neo-Carnapianism defended in this dissertation has for ontology in general is simply that the most common quick argument for a particular ontology is false. In discussions about ontology I hear it all the time that one side of the debate says

(320) Hey, you used the word 'property', so you are, after all, committing yourself to properties in your ontology!

It is often perceived that if one believes that properties do not exist then one has to stop talking about properties and reject all theoretical enterprises and theories that take recourse to properties in some way or other. In particular, if one quantifies over properties then one most certainly commits oneself to properties. But this way to see what ones ontological commitments are is a mistake. Simply because one uses the word 'property', or any other word doesn't mean that any entities called 'properties' are presupposed for what one says to be literally and objectively true. The methodology to achieve results in ontology has to be different. That one talks about a kind of thing alone doesn't mean that there are any entities presupposed to exist. What one does with this talk, what its function is, is what matters. To stress this, I will put it into a principle:

The Methodology Principle

Talking about a certain kind of thing, K, either with quantification over Ks, or with noun phrases that are about Ks, or ones using Ks in theoretical enterprises, alone does no mean that one is ontologically committed to Ks. What ontologically commits one to Ks is what one does when one talks about them.

- Only if such talk is to be understood externally and if it at the same time is meant to be literally and objectively true does it have direct ontological relevance.
- If it is meant to be literally and objectively true, but it is to be understood internally, or if it isn't meant to be literally and objectively true, then we can't get any direct ontological implications from it.
- Determining whether or not talk about Ks is internal or external, literal or non-literal, will be a tricky matter that depends on the details of what we are doing or trying to do when we talk about Ks. It will depend on what the function of such talk is (on that occasion).

Several of the points developed in the preceding chapters are necessary to defend this methodology principle. I will list the most important ones:

- Distinguishing topical from referential aboutness. Otherwise any talk about Ks might trivially seem to have ontological relevance.
- Distinguishing internal from external quantifiers, in particular noting that we have a need in communication for quantifiers to play more than one role.

- Noting that there are non-referential, non-quantificational noun phrases, and that we use such noun phrase to communicate complicated information about ordinary objects in a simple subject-predicate sentence, as with generics, for example.

- Noting that some uses of talk about certain things in theoretical enterprises should be understood instrumentally, and that for the uses we make of them in these theoretical enterprises it is irrelevant whether or not there really are such things.

This way of doing ontology applies in general, not just to properties, propositions and natural numbers. In addition to defending this general view of ontology, I have argued that in the special case of properties, propositions and natural numbers, our ordinary talk about them doesn't have any ontological presuppositions for its literal and objective truth. Much of this talk is meant to be literally and objectively true, but some of our theoretical talk about them (like taking recourse to properties and propositions in semantics) isn't.

6.9 *Remnants of Meaning*, language created entities, and pretense

There are other approaches to deal with the problems we have dealt with here. One of them is in its goal quite similar to the present one, the other two are different in what they argue for, but they also try to give an account of the some puzzling features that ontology has. Two of these accounts are by Stephen Schiffer. The first one is proposed in his book *Remnants of Meaning*, the second one is one that he developed after he abandoned the first one. The third approach is due to Steve Yablo. I will discuss all three in this order.

6.9.1 Schiffer's *Remnants of Meaning*

Remnants of Meaning isn't mainly about ontology, but about the philosophy of mind and language. In it Schiffer defends the view that even though ontological physicalism, the view that there are no non-physical entities of any ontological category, is true, still, there are non-reducible, but true, sentences ascribing mental properties to people, including belief ascriptions. Schiffer argues that such belief ascriptions cannot be reduced to the level of the physical, and that they cannot be accommodated in a compositional truth theoretic semantics. Overall, Schiffer concludes that the questions that define most of contemporary philosophy of language have false presuppositions, because they try to give a theory of meaning that can't be given.

Even though Schiffer's book isn't mainly about ontology he makes a number of claims related to ontology that he takes to be required to defend his overall view of mind and language. These claims are very much in the spirit of what I have defended in this dissertation. In particular, they include:

- that-clauses and expressions of the form 'the property of being F' are not referential expressions
- nominalism about properties and propositions
- quantifiers in trivial quantifier inferences involving properties and propositions do not have any direct ontological relevance

By the time I first read this book, about 5 years after it appeared, Schiffer had already given up all three of these points, and started to develop the view he presently holds, *pleonastic Fregeanism*, which will be discussed shortly. It seems to me, though, that this is a mistake and that the older view is closer to the truth than the newer view. In particular, it seems to me that the apparent reasons for abandoning the older view, and the problems it had, can be overcome. I will momentarily discuss the reason that Schiffer mentions himself why his view about properties and propositions should be abandoned in favor of his more recent one. But first, let me mention some aspects of Schiffer's view in *Remnants of Meaning* that I find problematic. After that we will have a look how these points have been resolved by what has been said above.

- *Substitutional quantification.* Schiffer argues that quantification over properties and propositions does not bring any ontological commitment with it because it can be understood as substitutional quantification. He does not give much argument for this, but quotes (Gottlieb 1980) approvingly. But as we have seen above in section 2.3.7, it is not at all clear how substitutional quantification relates to ontology. Simply because a quantifier can be given substitutional truth conditions doesn't mean that it carries no ontological commitment.
- *Inexpressible properties.* Schiffer mentions the worry about inexpressible properties in a footnote, and suggests that he might be able to deal with it by giving quantification

over properties a substitutional semantics in an extended language. This is a mistake, I think, and I suspect that this one is to a good deal responsible why Schiffer thought that his view in *Remnants of Meaning* should be given up. I will discuss this shortly.

- *The role of compositionality.* Schiffer denies in *Remnants of Meaning* that natural languages have a compositional truth-theoretic semantics. Not only that, he thinks such a denial in necessary to defend nominalism about properties and propositions. Without such a denial one would have to assign a semantic value to expressions like 'humility'. But according to the nominalist, 'humility' refers to nothing, thus no semantic value can be given 'humility'. Thus nominalism about properties and propositions seems to be committed to denying the insight and value of semantic theories that take recourse to properties and propositions.

- *Quantifier free talk about properties and propositions.* Schiffer claims that expressions like 'being a dog' and 'humility' are non-referential, but no positive account of what their function is is given.

These points are problems for this view as it was stated by Schiffer, but none of them is essential to the view. In fact, with what we have seen so far we can answer all of them. Let me go through them again, and say how we can now deal with them:

- Substitutional quantification by itself does not help nominalism. What would help would be to argue that quantifiers in general, not just when they range over properties and propositions, have more than one role to play. What would help is to argue for all of the following:

  - quantifiers in general are used in more than one way, and have more than one function in communication
  - of of these functions can only be performed if the quantifiers are detached from direct ontological relevance
  - when we quantify over properties and propositions we use the quantifier in this way

  If all these were true then quantification over properties and propositions would indeed not have any direct ontological implications. I have argued for all of these points above.

- Dealing with quantification over inexpressible properties by taking recourse to an extension of our language, adding more names and predicates, is a mistake. Taking recourse to extensions of ones language will, in fact, make substitutional quantification analogous to objectual quantification. The difference between these two ways of looking at quantifiers will collapse and there won't be anything special about the substitutional quantifier. Not only is taking recourse to substitutional quantification a mistake, treating inexpressible properties this way is a mistake also. We have seen above, in section 4.4, that inexpressible properties can and should be accommodated differently.

  If one believes that quantifiers play two different roles, imposing domain conditions and having a certain inferential role, then this will have relevance for ontology, as argued above. But even if one believes this, and one believes that quantification over properties in ordinary communication is based on quantifiers used for their inferential role, then one should not accommodate inexpressible properties by taking recourse to extensions of our language. In doing this one would thereby acknowledge that there is more to quantification over properties than a use of the expressive power that we get from quantifiers in their inferential role. And if we would use more than this it wouldn't be clear any more why these uses of quantifiers are not directly relevant to ontology. By accommodating inexpressible properties the way it was done above we do only rely on the expressive power we get from the quantifier's having a certain inferential role. However, it is the inferential role in a context sensitive language.

- The issue of compositionality isn't directly related to the question of nominalism about properties and propositions. Expressions like 'humility' and that-clauses can be assigned semantic values whether or not they are referring expressions. Semantic values of referring expressions don't have to be their referents, and non-referring expressions can have semantic values, too. In addition, even the assignment of 'properties' and 'propositions' as semantic values of 'being F' and that-clauses is independent of the issue of nominalism, as we have seen in section 6.3.1.

- An account of the function of quantifier free talk about properties and propositions can be given, and much can be said about this even if nominalism about properties and propositions is true. We have looked at some cases of this in chapter 3.

6.9.2 Schiffer's present view

Schiffer gave up the view he held in *Remnants of Meaning* partly because "there are serious philosophical difficulties involved in trying to spell out the kind of non-objectual quantification involved." (Schiffer 1996b, 152). I understand how these might arise if one tries to accommodate inexpressible properties and propositions by taking recourse to extensions of ones language. This way of doing it makes non-objectual quantification very much like objectual quantification.[8] However, with the distinctions and views arrived at in chapter 2 we can say there is a perfectly fine sense in which quantifiers are used other than to range over some domain of entities. I don't see what other worries Schiffer might have in mind, and I think that there is no problem about quantification for a view like the one defended in this dissertation.

Schiffer's present view about properties and propositions is trying to find a middle way between nominalism and heavy duty platonism. In particular, he tries to give an account of how it can be that we start talking about properties and propositions so easily. The view he holds is trying to combine a form of conceptualism, the view that properties and propositions are language created entities, with realism, the view that they are language independent entities. His view is that properties and propositions exist, but that they don't have hidden and substantial natures. There is nothing substantial to be discovered about what things they are. All that can truly be said of them is what follows from, or is determined by, our talk about them. Schiffer takes this to be analogous to our talk about fictional entities. According to Schiffer, fictional characters are language created entities. When an author makes up a story and introduces a character they, in a sense, create the characters that occur in the story. For us to successfully refer to these characters nothing more is required than to step back and talk about the story. So, when Conan Doyle makes up the Sherlock stories he creates the fictional character Sherlock Holmes, in the sense that nothing over and above his doing this and our starting to talk about his story is required for us to succeed in referring to Sherlock Holmes. Something like this also happens when we talk about properties and propositions, according to Schiffer. All that is required for our referring to them to succeed is to start talking about them using the *something-from-nothing* transformations that allow us to introduce such talk without change of truth conditions. These are the schemas

[8] Bonevac spells out how apparently problematic cases for substitutional quantification can be accommodated using extensions of ones language. See (Bonevac 1984).

(321) a. p iff it's true that p.

b. x is F iff x has the property of being F.

According to Schiffer, properties and propositions are language created entities in the same sense in which the fictional character Sherlock Holmes is. But they are also language independent in the sense that they would be there even if no one would have started to talk about them through using the *something-from-nothing* transformations. However, starting to talk about them is all that is required to discover them.

I think there are a number of problems with this view. One of them has to do with the sense in which one can say that properties, propositions or fictional characters are created by our talking in a certain way. To be sure there undoubtedly is a sense in which Conan Doyle created the fictional character Sherlock Holmes. But it is in no way clear that when he did this he created some entity. We often speak of creating something in a much broader way. Sometimes one creates a thing, an entity, as when one bakes a cake. But often one makes some other change to the world that shouldn't be described as creating some entity, even though one does create something. Consider the following ways in which we speak of someone creating something:

(322) a. With his inappropriate remarks John created a lot of tension among the people in the room.

b. The rising gas prices created an uproar among Palo Alto's SUV owners.

It is in no way clear that Conan Doyle created a thing when he created the fictional character Sherlock Holmes. To be sure, Conan Doyle changed the world in a certain way when he wrote his novels. He made the world such that it was now possible to talk about Sherlock Holmes. But this can be understood in two ways in which one can talk about Sherlock Holmes, as topical aboutness or as referential aboutness (see section 4.3.1 above). Distinguishing these two ways in which one can talk about Sherlock, we can distinguish two ways in which one can create the fictional character Sherlock Holmes:

- **Topical creation** To create a fictional character in this sense is to change the world in such a way that talk about the character, in the sense of topical aboutness, becomes possible (for a group of people). To do this it will be necessary to introduce some term for the character, give some description of what he is supposed to be like, and communicate this to the group of people who will now be able to talk about him, in this sense of 'about'.

- **Ontological creation** To create a fictional character in this way is to change the world in such a way that talk about the character becomes possible, in the sense of referential aboutness, and to do this by creating the referent of what this talk is about.

Conan Doyle created Holmes in the sense of topical creation. There is no question about that. He changed the world in such a way that a certain discourse in our language community became possible, one where we can say things like

(323) Sherlock would have solved that problem in no time!

and our saying this will be meaningful to our fellow language community members. But did Conan Doyle do this by ontologically creating Sherlock? To be sure, this is controversial, but it seems unlikely to me that this is the right thing to say. One of the problems with saying this is that it is not clear how we can make sense of ontological creation like that. How can we create some entity simply by speaking, other than the sounds we utter? It is unclear, and I think implausible, to think of creation in fiction as ontological creation. To be sure, it is creation. The author produces a real change in the world, that is related to what other members of the language community can talk about. But all that can happen without ontological creation.

I mention all this to raise some doubt about Schiffer's taking recourse to language created entities in fiction to use this as a model of language created entities more generally. It is quite implausible in the case of fiction that authors ontologically create their characters. The reason why Schiffer takes recourse to this is to have talk about properties and propositions be referentially about properties and propositions, and to have quantification over them to be objectual (or external, with the terminology of this dissertation). As we saw, Schiffer though that taking recourse to non-objectual quantification will lead to philosophical trouble, and one of the reasons for this, I speculated, was to accommodate inexpressible properties and propositions in such a framework. I'd like to note here, though, that it is quite unclear how his present view can deal with inexpressible properties and propositions. If properties and propositions are language created entities, and are mere shadows of predicates, how can there be any inexpressible ones? It seems that inexpressible properties and propositions are a problem for this view just as much as for Schiffer's older view. If quantification over properties is objectual / external quantification over some domain of language created entities, how can there be true quantified statements where the quantifier has no witness for its truth? The solution to the problem of inexpressible properties and propositions give in chapter 4 might be available to Schiffer, too. But with it and the way of dealing with quantifiers given in chapter 2 we do not get any advantage from claiming that properties and propositions are language created entities. We only get the additional trouble of spelling out how our talking a certain way can create these entities.

A minimalism about properties and propositions, in the sense that there isn't anything substantial to be discovered about them, has in fact been defended before, but in a different part of the woods. Platonists who believe in some form of a plenitude of abstract objects sometimes believe that abstract objects are partial, in the sense that there are properties such that the abstract object neither has it nor its negation. In particular, platonists like Zalta[9] think that which one of the plenitude of abstract objects we talk about depends on what properties we ascribe to the thing. For example, what properties certain mathematical objects have, like the sets that ZF is about, is closely related to what properties these sets have according to ZF. Since ZF does not determine whether or not the continuum hypothesis holds the sets that ZF is about also do not determine them. There is no fact of the matter whether or not the continuum hypothesis holds about them.

This kind of platonism about abstract objects, taking them to be partial objects, fits well with some of the intuitions that motivate Schiffer's recent view. The platonist view, however, does not take recourse to language creation, and it seems hard to me to see why one might prefer a language creation version of this view, rather than a form of platonism based on partial abstract objects.

6.9.3 Yablo's pretense account

Yablo recently defended a neo-Carnapian approach to ontology.[10]. But his kind of neo-Carnapianism is in a quite different spirit than the one defended in this dissertation. His is, in a sense, much more Carnapian. Above I have argued that because quantifiers are semantically underspecified with respect to whether they impose domain conditions or occupy

[9]See his (Zalta 1983) or (Linsky and Zalta 1995).
[10]See (Yablo 1998).

a certain inferential role. And because of this one can utter a sentence of the form "Are there Fs?" with two different truth conditions. One of them will have a trivially affirmative answer, and the other will have a non-trivial answer. Thus I agreed with Carnap that we have internal and external questions about what there is, and the internal ones (usually) have trivial affirmative answers. But contrary to Carnap I held that both of these kinds of questions are perfectly meaningful and factual. This was contrary to Carnap who thought that external questions, the ones that make up the questions of ontology, are meaningless, i.e. without factual content. In (Yablo 1998) Yablo defends a form of neo-Carnapianism that agrees with Carnap about the meaningfulness of ontological questions. He distinguishes two attitudes about ontological questions: the curious, who try to find out what the answer to such questions are, and who take them to be completely meaningful and factual questions, and the quizzical, who think that there isn't anything to be found out about them. Yablo thinks that the quizzical approach to such questions is correct, and thereby he defends a view that, as he himself says, seems to be quite counterintuitive, namely that there is no fact of the matter about what there is. He thinks that the internal/external distinction should be drawn in connection with the literal/metaphorical distinction. In particular, we can ask what there is, taken literally, or we can ask what there is within a certain game of pretense. According to Yablo, much of our ordinary talk is not meant to be literal and much of our non-literal, make-believe, talk gives rise to statements about what there is. To have a discipline of ontology one would have to separate the literal from the non-literal, but Yablo doubts that this is possible, and he thus doubts, with Carnap, that there can be a philosophical discipline worth the name ontology.

It is unclear if the fact that there is no way to distinguish the metaphorical from the literal really has such radical consequences as Yablo suggests. It seems that at most such a distinction can't be drawn by us, but not that there isn't a fact of the matter which utterances are only metaphorically true and which ones are literally true. All that I see reason for in Yablo's paper is that it is really hard to make the distinction between the literal and the metaphorical, but not that there is no fact of the matter about this difference. And if that is all we get then it would only follow that it is hard to find out what is implied to exist by our beliefs, since it is hard to distinguish the literal from the metaphorical, but not that there is no fact of the matter about what there is. After all, how can there be no fact of the matter about what things the world is made up of?

In a different paper Yablo uses some of these ideas to deal with some of the same questions that this dissertation dealt with.[11] In particular, Yablo discusses the role that talk about certain entities has that gets introduced through paraphrases or equivalences. One of his main examples is

(324) An argument is valid iff it has no countermodels.

Yablo notes that the existence of countermodels does not seem to matter much for our taking recourse to them. He then offers a pretense account of our talk about them. Basically, we pretend that there are countermodels for certain purposes. Yablo not only deals with countermodels, but wants his account to apply quite generally to all kinds of entities talk about which we introduce through such a paraphrases, or explications. One of the ones he mentions is that

(325) x is F iff x has a certain relation to the property of being F.

Yablo's overall goal is in a very similar spirit as the goal of this dissertation, but the way it is supposed to reached is quite different. In the present terminology we can say that the problem that we both try to deal with is that it seems that certain talk about certain things does not seem to have ontological presuppositions for living up to the standards of correctness that it is supposed to live up to. Yablo concludes that if this is so then it is out of the question that the standards of correctness that govern talk about these things is literal and objective truth. It seems to be out of the question because it seems to imply directly that there are things of that kind, and thus that this talk would have ontological presuppositions after all. Yablo rather thinks that such talk should not be understood as aiming at literal truth, and he proposes to account for the function that such talk has within a pretense framework.[12] When we talk about countermodels, the average star or properties we are not speaking literally. That's why no entities are required for this talk to do what it is supposed to do.

As much as I like the overall spirit of Yablo's account, I don't think it can be right. At least not in the general way in which he seems to want it to be. To be sure, sometimes we talk about things for certain purposes, and such talk should not be understood as aiming to

[11]See (Yablo 1999).
[12]See (Walton 1990) for more on pretense.

be literally true. In fact, I have defended this myself above, when we looked at the function that talk about semantic values has. But I think this can't be right when it comes to our ordinary and everyday talk about properties, propositions and natural numbers, besides several other kinds of things. In particular, I think Yablo is mistaken to quickly conclude that if such talk is meant to be literally and objectively true then it clearly has ontological presuppositions. That this is a mistake was the main point of this dissertation. In addition, I think it is a mistake to think that talk about properties, propositions and natural numbers is not literally true. Is this at all plausible for case like the following?

(326) Being a philosopher is fun.

(327) That John is rich is unlikely.

(328) Three and two are five.

Yablo does not give any arguments why such talk should be understood non-literally, other than that this would account for the fact that such talk doesn't seem have ontological presuppositions to live up to the standards of correctness it is supposed to live up to. The main reason to claim that it isn't meant to be literally true, as far as I can see, is to account for the lack of ontological presuppositions that this talk has to do what it is supposed to do. But by now we have seen that this in no way is in conflict with a certain domain of discourse being aimed at literal and objective truth. Lack of ontological presuppositions for a domain of discourse to live up to the standards of correctness it is supposed to live up to is in no conflict, in principle, with this standard being literal and objective truth. In addition, we have seen that this is in fact the case with our talk about properties, propositions and natural numbers. Whether it carries over to other domains of discourse will depend on what function this discourse has, and this will depend on a number of details.

I do believe that something like the pretense account Yablo offers is correct for certain domains of discourse. In fact, such an account is very close to the form of instrumentalism defended about about semantic values. I think this is correct for certain domains of discourse, but it isn't correct for properties, propositions and natural numbers in everyday talk, in mathematics and much of philosophy.

6.10 Conclusion

Talk about properties, propositions and natural numbers never was meant to be about some domain of entities. But that doesn't mean that it isn't literally and objectively true, and it doesn't mean that we can't use properties, propositions and natural numbers in semantics and other theoretical enterprises. What it does mean, though, is that a certain set of metaphysical questions is based on a confusion. To ask whether or not the property of being a dog is fully present in both Fido and Fifi is based on a confusion, even though there is something that Fido and Fifi have in common, namely the property of being a dog. The results of this dissertation are bad news for people who like a certain approach to ontology and to characterizing objectivity. They are good news for everybody else, or so I'd like to think.

Appendix A

Quantification and Infinitary Logic

We have seen in chapter 2 that we need quantifiers for more than just imposing certain conditions on the domain of objects that make up reality. We also need them to occupy a certain inferential role, and to do so independently of issues in ontology. Quantifiers, we have seen, do in fact do both. On occasion they impose domain conditions, and on occasion they occupy a certain inferential role. But they can't do both with the same contribution to the truth conditions in a language like ours. So in a language like ours the expressive power we get from quantifiers is greater than merely to impose domain conditions. In chapter 2 I have outlined what the truth conditions of a certain quantified statement have to be for it to occupy the inferential role for which we want it. We have seen that the quantified statement has to be equivalent to certain infinite disjunctions or conjunctions. In this chapter we shall look at infinitary logics in some more detail, and in particular deliver on some of the promises that were made above. I chapter 4 I said that the expressive strength we get from internal quantification over properties and propositions is greater, and a bit more tricky to capture, than internal quantification over objects would give rise to. In addition, I said that with such quantification there is the threat of paradox, but that the standard approaches to dealing with paradoxes that arise from quantification over properties and propositions can be carried over to the present way of understanding such quantification. All of this involves looking more precisely at certain infinitary logics that will be used to model the expressive power of internal quantification in general, and over properties and propositions in particular. To do this, let's first look at infinitary logic in general. Infinitary logics are a powerful extension of ordinary, finitary logics. The most well known of these logics are extremely powerful, and in fact so powerful that it seems rather unlikely that they will have anything to do with natural languages. But besides these very strong infinitary logics we will also consider small fragments of infinitary logic that are much more controllable, and much less powerful. After that we will look at a more precise characterization of the expressive power of internal quantifiers in a simple case, then we'll look at quantification over properties and propositions, and finally at how standard accounts of dealing with the paradoxes carry over to the present approach.

A.1 Fragments of infinitary logic

A.1.1 $L_{\infty,\infty}$

In standard first order logic we have a restriction in the recursive definition of being a well-formed formula. Only finite boolean combinations of formulas will themselves be formulas. In particular, only the conjunction or disjunction of finitely many formulas will itself be a well-formed formula. But this restriction isn't essential to the formula building operations, nor to the standard way of giving a semantics for such formal languages. And similarly with quantification. In first order logic we only allow quantifiers to bind one variable at a time, and when combined with the finite boolean combinations we, in effect, only allow quantifiers to bind finite sets of variables. Again, this isn't essential to either giving a precise syntax or semantics. In infinitary logic one gives up these restrictions. The most radical way to do this is to simply allow conjunctions and disjunctions over *arbitrary* sets of formulas, and quantification over *arbitrary* sets of variables. So, suppose you start out with variables, constants, predicate symbols and function symbols (finitely or infinitely many), then you build the atomic formulas, as usual, and close the formulas under the following operations:

> if all members of a set Φ are formulas, then $\bigvee \Phi$ is a formula, and $\bigwedge \Phi$ is a formula

The semantics of these is a straightforward generalization of the finite case:

> $\bigwedge \Phi$ is satisfied by a sequence s iff all of the members of Φ are satisfied by s.

> $\bigvee \Phi$ is satisfied by a sequence s iff one of the members of Φ are satisfied by s.

And similarly for quantifiers:

> If ϕ is a formula then $\exists\vec{x}\phi$ is a formula, where $\vec{x}$ is any set of variables.

Analogously for universal quantifiers. The semantics of quantified formulas is an obvious generalization of the one for first order quantified formulas:

> $\exists\vec{x}\phi$ is satisfied by a sequence s iff there is a sequence $\check{s}$ that differs from s only in what it assigns to the variables in $\vec{x}$, and $\check{s}$ satisfies ϕ.

$L_{\infty,\infty}$ is extremely powerful. Basically everything can be described and characterized in this logic, assuming a big enough base vocabulary. So, $L_{\infty,\infty}$ is too strong to be of much interest.[1] But there are a number of ways of making $L_{\infty,\infty}$ weaker, and more interesting for us here.

A.1.2 Cardinality restrictions of $L_{\infty,\infty}$

The straightforward way of weakening $L_{\infty,\infty}$ is to impose restrictions on which sets of formulas are allowed in forming infinitary conjunctions and disjunctions, and which sets of variables infinitary quantifiers can bind. In the case of $L_{\infty,\infty}$ we allowed arbitrary sets, independently of their infinite size or anything else. One way to restrict this is to allow only conjunctions and disjunctions of sets of formulas that have a certain bounded cardinality, and to allow quantifiers to bind only sets of variables of another bounded cardinality. The resulting logics are called $L_{\kappa,\lambda}$. Here κ and λ are infinite cardinal numbers, and we allow only disjunctions and conjunctions of sets of formulas of cardinality less than κ, and only quantification over sets of variables of size less than λ.

Even though this is clearly a substantial restriction, the resulting languages still have extremely strong expressive strength. For one example, in the language $L_{\omega_1,\omega}$ built up from the language of arithmetic one can explicitly define every set of natural numbers, independently of its complexity. Since for every natural number n '$x = \overline{n}$' is a predicate of that language, the disjunction

$$\bigvee_{n\in S}(x = \overline{n})$$

is wellformed in $L_{\omega_1,\omega}$ for every set of numbers S, since there are only countably many natural numbers.

These large infinitary logics are much too complex to have anything to directly to do with natural languages. They are of interest in set theory and model theory, in particular,

[1] It make a rare appearance in philosophy in (McGee 1996).

they are very useful in studying large cardinals, or to characterize complicated models. But all that has little to do with natural languages.

A.1.3 Admissible fragments of $L_{\omega_1,\omega}$

However, it is possible to specify infinitary logics that are less expressive than the ones we have seen so far. One well known case of this are the so called admissible fragments of infinitary logic, in particular countable admissible fragments of $L_{\omega_1,\omega}$. Admissible sets are sets that are a model of a certain rather weak set theory, KP, or Kripke-Platek set theory.[2] The smallest models of KP are initial segments of Gödel's constructible universe, $\mathbf{L}$. In particular, the ordinals that are such that $\mathbf{L}(\alpha)$ is a model of KP are called *admissible ordinals*. $\mathbf{L}(\alpha)$ is the set of all constructible sets that can be constructed before stage α. The first α above ω that is such that $\mathbf{L}(\alpha)$ is a model of KP is ω_1^{ck}, the first non-recursive ordinal. Thus $\mathbf{L}(\omega_1^{ck})$ is the smallest admissible set that contains ω.

We can encode infinitary formulas in a certain straightforward way as sets, and via this coding we can speak of admissible fragments of $L_{\omega_1,\omega}$. For example, suppose we deal with a language that has countably many primitive symbols, i.e. predicate symbols, constants etc.. Identify each one of them in a 1-1 way with a hereditarily finite set. Then identify formulas build up from them with larger sets. For example, the rule of building infinite disjunctions can be seen as identifying a certain set with an infinitary formula, as follows:

> If $\phi_1, \phi_2, \phi_3, \ldots$ are formulas then $\{\bigvee, \{\phi_1, \phi_2, \ldots\}\}$ is also a formula, where $\bigvee$ is identified with some set.

We can now restrict ourselves to the infinitary logic that only contains infinitary formulas that are contained in $\mathbf{L}(\omega_1^{ck})$, which we will call $L_{\mathbf{L}(\omega_1^{ck})}$. This will only allow for the formation of disjunctions and conjunctions over sets of formulas of a certain complexity. After all, all these sets have to be definable (with parameters) at an earlier stage of the construction of the constructible universe. $L_{\mathbf{L}(\omega_1^{ck})}$ is substantially weaker than $L_{\omega_1,\omega}$. For one thing, the former is countable, the later uncountable. Also, the former does not allow for the explicit definition of all sets of natural numbers any more. Only those that are constructible before ω_1^{ck} are definable.

[2]See (Barwise 1975) for more on admissible sets, and admissible fragments of infinitary logic.

Admissible fragments of infinitary logic have a certain technical importance that we won't get into here (it is spelled out in (Barwise 1975)). For us here they are only an example of one way to define rather weak infinitary languages, like $L_{\mathbf{L}(\omega_1^{ck})}$. However, weak as these languages are in comparison with $L_{\omega_1,\omega}$, they are still rather complex. Much more so than we need for the purpose of this dissertation, as we will see below.

A.1.4 Even smaller fragments of $L_{\omega_1,\omega}$

We can restrict ourselves even further by allowing only even less complex conjunctions and disjunctions. There are a number of ways in which one can make such restrictions, and it will depend on for what purpose one wants these infinitary logics what restrictions will be reasonable. One might want to restrict the complexity of the sets of formulas from which conjunction can be build even further. For our purposes here, we will use infinitary languages that only allow for the formation of infinite conjunctions and disjunctions of a very simple kind. In the next section we will have a look at what these languages are, and how they relate to other infinitary languages.

A.2 The expressive strength of internal quantifiers

A.2.1 The base cases

What expressive strength we get through internal quantification can be nicely modeled by seeing what fragment of infinitary logic is needed to model these quantifiers. We have seen in chapter 2 that we can account for the inferential role that internal quantifiers give rise to by associating them with certain infinite disjunctions and conjunctions, quite independently of how the rest of the language is modeled. But what infinitary languages are needed to do so will differ depending on what we quantify over. For example, in chapter 2 we only looked at a very simple case where we assumed that there was a fixed class of terms of the language, and that the internal quantifiers gave rise to an infinitary expansion of the original language where disjunctions and conjunctions were expressible that involved a disjunct for each one of these terms. For example

(329) Something is F.

understood internally was modeled as

(330) $\bigvee_{t \in T}(F(t))$

where T is the class of terms of the original language without internal quantifiers. We did not deal with new terms that might arise in this expanded language with internal quantifiers. When we did the same for quantification over properties and propositions we also assumed that no new predicates or sentences where introduced in the class of expressions over which a disjunction was formed. But, of course, this was too simplistic. Once we have internal quantifiers we can form new terms, predicates and sentences. And it seems that they, too, should occur in disjunctions and conjunctions formed over all the terms, predicates, and sentences, of our language. We will look at this in the next section, and see if this is any particular trouble for the present view, and how this affects how we should look at the expressive strength of internal quantifiers.

An example of the base case for internal quantification over properties is given by

(331) There is a property that Bush and Nixon both have.

This has the truth conditions

(332) $\bigvee_{P \in \mathbf{P}}(P(b) \wedge P(n))$

where b is Bush, n is Nixon and $\mathbf{P}$ is the class of all predicates in the language in question that do not contain internal quantifiers. In addition, we have seen that to accommodate inexpressible properties we have to allow a special kind of variables modeling demonstratives, and infinitary quantifiers that bind them from the outside. I will largely ignore this aspect for now, though.

So far, as the base case, we get the expressive power of the following infinitary expansion of the original language by adding internal quantification over properties to it.

- Suppose L models the original language, and $\mathbf{P}_L$ is the class of (unary) predicates of L.
- Besides the operations that formulas of L are closed under, formulas of L^+, the language containing internal quantification over properties and propositions, is also closed under the following operation:

  - Add auxiliary second order variables F_i, and propositional variables P_i to L and define atomic (auxiliary) formulas of L^+ using these.
  - If '$\phi(F_i)$' is an (auxiliary) formula of L^+ with 'F' an auxiliary property variable, then

    $\bigvee_{P \in \mathbf{P}_L} \phi(P/F_i)$

    is a (auxiliary) formula of L^+. Similarly for propositional variables, and conjunction.

- The formulas of L^+ are the (auxiliary) formulas with no auxiliary variables.

The expressive strength in this simple case is the infinitary language we get by taking the language consisting of all the formulas of L^+. To be sure, how complex this language is will depend on how complex the set of predicates of L is.

A.2.2 The trickier cases

Once we have internal quantifiers in L^+ we can use these to form new predicates and new sentences, ones that couldn't be formed in L. Thus what infinitary formulas can be formulated in L^+ will be trickier. On top of this, it can give rise to paradoxes. If quantification over propositions ranges over propositions that themselves involve quantification over propositions then paradoxes can occur. And similarly with quantification over properties. This, however, is no special problem for internalism about talk about properties and propositions. Everybody, internalist, externalist, fictionalist or whatever, has to deal with the paradoxes. What I would like to point out in the following is that the standard tools to deal with the paradoxes, a Tarski-style hierarchical account, or a Kripke-style partial account, are available to the internalist just as well as to everybody else. Thus the internalist does not have any more problems dealing with the paradoxes than anybody else. In addition, when we look at the hierarchical accounts we can have a bit closer look at what fragments of infinitary logic one has to take recourse to in order to accommodate this. We will see that we will take recourse to fragments that are quite a bit smaller than the smallest admissible fragment, $L_{\mathbf{L}(\omega_1^{ck})}$.

A.2.3 Hierarchical approaches

A hierarchy of infinitary languages

One of the standard approaches to deal with semantic paradoxes is to involve hierarchies of stronger and stronger languages that each are consistent, and allow us to approximate the apparent circularity that we have in natural languages. In English we can formulate expressions that talk about themselves, as in

(333) What this sentence says is not true.

And this ability of sentences talking about their own truth or falsity is one of the sources that gives rise to paradox. Tarski proposed to consider a hierarchy of truth predicates, each applying only to sentences containing truth predicates lower in the hierarchy. This way consistency can be guaranteed and paradox can be avoided.

A similar approach carries over to an internalist conception of talk about properties and propositions (as well as an externalist conception of such talk). We have to acknowledge that once we have internal quantifiers in our language we can form new predicates with it, ones that we couldn't form before. And this gives rise to new infinitary conjunctions and disjunctions that we can express. Adopting a hierarchical approach is a straightforward way of endorsing this extra complexity while maintaining consistency. In the case of internalist quantification over properties, for example this would go as follows.

Suppose L is our base language, the one without internal quantification over properties. We then define a hierarchy of infinitary languages on the basis of L that corresponds to a hierarchy of internalist quantifiers over properties (i.e. quantifiers of level α range over properties of 'complexity type' less than α):

Let $\mathbf{P}$ be the class of predicates of L. Let L^1 be L extended by internalist quantifiers over properties that are expressible with predicates of L. The infinitary formulas of L^1 exactly like the ones of L^+, the base case from above:

- All formulas of L are formulas of L^1.
- In addition, we allow special predicates F_i to occur in the formula building recursions. Thus the formulas of L extended by these predicates are (auxiliary) formulas of L^1.

- If $\phi(F_i)$ is an (auxiliary) formula of L^1 then $\bigvee_{P\in\mathbf{P}} \phi(P/F_i)$ and $\bigwedge_{P\in\mathbf{P}} \phi(P/F_i)$ are (auxiliary) formulas of L^1. Here $\phi(P/F_i)$ is the result of substituting P for all occurrences of F_i in ϕ.
- The formulas of L^1 are all the (auxiliary) formulas that do not contain any auxiliary predicate symbols, i.e. F_is.

To extend this to higher levels of a hierarchy of infinitary languages we generalize this definition in a straightforward way. We simultaneously recursively define a hierarchy of sets of predicates $\mathbf{P}^\alpha$ and of languages L^α for all ordinals α:

- Let $\mathbf{P}^\alpha$ be the union of all predicates of the languages L^β, for $\beta < \alpha$.
- The formulas of L^α contain all the formulas of L. The auxiliary formulas are all those that can be built from L with extra auxiliary predicates F_i. The infinitary formulas are defined recursively as:

  If $\phi(F_i)$ is an (auxiliary) formula of L^α then $\bigvee_{P\in\mathbf{P}^\alpha} \phi(P/F_i)$ and $\bigwedge_{P\in\mathbf{P}^\alpha} \phi(P/F_i)$ are (auxiliary) formulas of L^α.

- Again, the formulas of L^α are all the auxiliary formulas with no auxiliary predicate variables.

So, the larger α, the larger the fragment of $L_{\omega_1,\omega}$ that L^α corresponds to. More and more infinite disjunctions and conjunctions become expressible.

Doing this for internal quantifiers over propositions is quite analogous. This way of setting up a hierarchy of languages has been used extensively in proof-theory, for example in the study of ramified analysis.[3] If the language of Peano Arithmetic is our base language then L^α will be equivalent to RA^α, ramified analysis at level α.

So, if one thinks that hierarchical approaches are the key to solving the paradoxes we get from quantification over properties and propositions, internalism is no worse off than externalism. Hierarchical accounts can be formulated for either one of them. In addition, looking at these hierarchical accounts we can see more about the expressive strength we get from internal quantifiers. The fragments of infinitary logic that are expressible with them

[3]See (Schütte 1960) and (Feferman 1968).

nicely measure the expressive strength that these languages have. The more infinitary formulas can be expressed the greater the expressive strength.

Comparing size

The fragments of infinitary logic we get here are on a much smaller scale than the ones encountered above, where we looked at cardinality restrictions and admissible fragments. Obviously, for all countable α, L^α will be countable. In addition, if α is recursive then L^α will be a proper fragment of the smallest admissible fragment of $L_{\omega_1,\omega}$. In fact, we can quite directly compare the smallest admissible fragment of $L_{\omega_1,\omega}$, which consists of all the infinitary formulas of $L_{\omega_1,\omega}$ that are contained in the constructible universe up to level ω_1^{ck}, with the fragments L^α. This can be done as follows.

We will only be concerned with countable α for now, and we will restrict ourselves to a countable base language L. We can then code all the finitary formulas of L in the hereditarily finite sets, $\mathbb{HF}$. One way of doing this is to associate with each of the basic symbols of the language a set in $\mathbb{HF}$, the code of the symbol, as follows:

- Predicate symbols $P_1, P_2, \ldots$ get associated with sets $< 0, 1 >, < 0, 2 >, \ldots$.
- Constants $c_1, c_2, \ldots$ get associated with sets $< 1, 1 >, < 1, 2 >, \ldots$. Similarly for function symbols.
- The logical symbols of L, $\exists, \neg$, etc., get associated with $< 4, 1 >, < 4, 2 >, \ldots$.

I will use $\overline{\phi}$ for the set associated with, or the code of, a formula ϕ. Complex formulas then get associated sets as follows:

- $\overline{\phi \wedge \psi}$ is $< \overline{\phi}, \overline{\wedge}, \overline{\psi} >$, and similarly for all other finite formula building operations.
- If Φ is a countable set of formulas then $\overline{\bigvee \Phi}$ is $< \overline{\bigvee}, \{\overline{\phi} : \phi \in \Phi\} >$. Similarly for conjunctions.

All codes of, or sets associated with, finitary formulas are in $\mathbb{HF}$. The complexity of the set of formulas over which infinitary disjunctions and conjunctions are formed will determine where in the constructible hierarchy the codes of these formulas will be. All the infinitary disjunctions and conjunctions that can be formed with our languages L^α are so simple that they will occur very low in the hierarchy of constructible sets. More precisely, they occur as follows.

Let λ_α enumerate the countable limit ordinals. That is $\lambda_0 = \omega, \lambda_{\alpha+1} = \lambda_\alpha + \omega$ and for limit ordinals γ, $\lambda_\gamma = \sup_{\beta<\gamma} \lambda_\beta$. In terms of ordinal arithmetic, $\lambda_\alpha = \omega \times (\alpha + 1)$.

Let $L_\alpha^{\Sigma_n}$ be the infinitary logic that consists of all formulas of $L_{\omega_1,\omega}$ that are definable by a Σ_n formula (of set theory with parameters) on the constructible hierarchy up to level α. In particular, $L_{\omega_1^{ck}}^{\Sigma_\infty}$ is the smallest admissible fragment of $L_{\omega_1,\omega}$, and for $n < \omega$ and $\alpha < \omega_1^{ck}$, $L_\alpha^{\Sigma_n}$ is a proper fragment of $L_{\omega_1^{ck}}^{\Sigma_\infty}$

The comparison of size of admissible fragments of $L_{\omega_1,\omega}$, and the languages L^α that we are dealing with here, is summed up in the following:

Theorem: L^α is a fragment of $L_{\lambda_{\alpha+1}}^{\Sigma_0}$.

Proof: By strong induction on α. For $\alpha = 0$, L^α is the finitary language $L_{\omega,\omega}$, all of whose formulas are coded by hereditarily finite sets, which are all definable with no unrestricted quantifiers before level ω of the constructible hierarchy, which is to say they are in $L_{\lambda_0}^{\Sigma_0}$, and thus in $L_{\lambda_1}^{\Sigma_0}$.

Assume we have shown that for all $\beta < \alpha$, L^β is contained in $L_{\lambda_{\beta+1}}^{\Sigma_0}$. We have to show that L^α is in $L_{\lambda_{\alpha+1}}^{\Sigma_0}$. All the formulas of L^α that are also in some L^β, for $\beta < \alpha$, are in $L_{\lambda_{\alpha+1}}^{\Sigma_0}$, since they are, by induction hypothesis, already in $L_{\lambda_{\beta+1}}^{\Sigma_0}$, which is contained in $L_{\lambda_{\alpha+1}}^{\Sigma_0}$. It is thus sufficient if we concentrate on the formulas of L^α that aren't in any L^β, for $\beta < \alpha$. These are the formulas that contain disjunctions (or conjunctions) formed over the set of predicates $\mathbf{P}^\alpha = \bigcup_{\beta<\alpha} \mathbf{P}^\beta$. These disjunctions are of the form

$$\bigvee_{P\in\mathbf{P}^\alpha} \phi(P/F_i)$$

Here ϕ is any formula of L^α, and can thus itself contain disjunctions over all predicates in $\mathbf{P}^\alpha$.

Let $\gamma = \sup_{\beta<\alpha} \lambda_{\beta+1}$. We will have to show that all new formulas of L^α can be defined with no unbounded quantifiers in the constructible hierarchy at level $\gamma + k$, for some finite k.

Since $\overline{\bigvee_{P\in\mathbf{P}^\alpha} \phi(P/F_i)} = \{\overline{\bigvee}, \{\overline{\phi(P/F_i)} : \phi \in \mathbf{P}^\alpha\}\}$, we have to see how we can define this latter set in $\mathbf{L}$. This involves defining a substitution relation for formulas of L^α, defining the set of predicates $\mathbf{P}^\alpha$, and considering that ϕ can itself contain infinite disjunctions over $\mathbf{P}^\alpha$.

$Sub(\phi, \psi, a, b)$ can be defined as there being a construction sequence s_1 for ϕ and a construction sequence s_2 for ψ such that they differ only in that wherever a occurs in s_1, b occurs in s_2. This can be done in the usual way (see (Devlin 1984, 39f.)) without unbounded quantifiers at the appropriate level.

Defining the set of predicates of a language is straightforward, whether or not the language is finitary (through defining the set of free variables along with the formula building recursion, and taking the (unary) predicates to be those with one free variable). If $\alpha = \beta + 1$ then this can always be done at only finitely many levels above λ_β. However, if α is a limit ordinal, and λ_α thus a limit of limit ordinals, we will only be able to define $\mathbf{P}^\alpha = \bigcup_{\beta<\alpha} \mathbf{P}^\beta$ at level $\lambda_\alpha + k$, for some finite k, and thus it will not be a member of $\mathbf{L}(\lambda_\alpha)$, but only be a member of $\mathbf{L}(\lambda_{\alpha+1})$.

All that remains to be seen is that for every formula of L^α there is some k such that this formulas is definable at level $\gamma + k$, without unbounded quantifiers. To do this we have to consider that infinite disjunctions over $\mathbf{P}^\alpha$ may contain infinite disjunctions over $\mathbf{P}^\alpha$ as their disjuncts. However, in the simple language we are dealing with here, we can assign to each infinite disjunction over the predicates in $\mathbf{P}^\alpha$ a rank as follows:

- $\bigvee_{P\in\mathbf{P}^\alpha} \phi(P/F_i)$ has rank 1 if ϕ does not contain any infinite disjunctions that do not already belong to some L^β, for $\beta < \alpha$.
- $\bigvee_{P\in\mathbf{P}^\alpha} \phi(P/F_i)$ has rank $n+1$ if the highest rank of a disjunction (or conjunction) that occurs in ϕ is n.

From the definition of the formula building operations we see that every disjunction has only a finite rank. Even though there can be infinitely many disjunctions within the scope of a disjunction, there can be only finite depth with respect to rank. We can now do a secondary induction with respect to this notion of rank. It is sufficient to show that for each new rank we only need to ascend finitely many steps in the constructible hierarchy to define all the necessary formulas with no unbounded quantifiers. But this is now straightforward, since all finite boolean combinations can only increase the level we need to ascend to a fixed finite amount, and (given what we have seen about the definability of $\mathbf{P}^\alpha$ and Sub) this is all that can happen in going from one rank to the next . $\Box$

Corollary: If α is recursive then L^{α} is a fragment of $L_{\mathbf{L}(\omega_1^{ck})}$.

Proof: If α is recursive then $\alpha < \omega_1^{ck}$. However, ω_1^{ck} is inaccessible through addition and multiplication from below. That is, if $\gamma_1, \gamma_2 < \omega_1^{ck}$ then $\gamma_1 + \gamma_2 < \omega_1^{ck}$ and $\gamma_1 \times \gamma_2 < \omega_1^{ck}$. Since $\lambda_{\alpha+1}$ is nothing other than $\omega \times (\alpha + 2)$ we get that $\lambda_{\alpha+1} < \omega_1^{ck}$.
By the theorem L^{α} is a fragment of $L^{\Sigma_0}_{\lambda_{\alpha+1}}$ and thus a fragment of $L^{\Sigma_\infty}_{\omega_1^{ck}} = L_{\mathbf{L}(\omega_1^{ck})}$. $\Box$

This shows that all the recursive levels of our hierarchy of infinitary languages give rise to proper fragments of the smallest admissible fragment of infinitary logic. So, even for very large recursive ordinals, we only need very little infinitary logic to accommodate the internal quantifiers in this framework. None of the extreme expressive strength of $L_{\infty,\infty}$ or $L_{\omega_1,\omega}$ is needed here.

Let me say as a final remark on hierarchies that all hierarchical approaches have to deal with the question which one of these many languages in the hierarchy should be taken to correspond to English (on a particular occasion). There is no clear and obvious answer to this, and in fact it might be seen as a good example of an underspecified feature of natural languages to a fan of a hierarchical approach. In any case, the internalist is in no worse position here than the externalist.

A.2.4 Partial approaches

Alternatively, one can deal with the paradoxes in a type free, but partial framework. In it not every sentence gets assigned a truth value, but the ones that do get assigned truth values do obey certain conditions about how their truth values have to relate to each other. A well known way of doing this is in (Kripke 1975). In it one specifies how an initial partial assignment of truth and falsity has to be extended in a certain well-behaved way. This can be iterated, and gives rise to an inductive definition of 'fixed-points' of partial assignments of truth.

This approach carries over to fragments of infinitary logic given by internal quantifiers over properties and propositions in two ways. On the one hand, internal quantification over propositions is closely related to quantification over sentences and a truth predicate. For example, the infinite disjunction $\bigvee_{P \in \mathbf{P}} P$ is equivalent to $\exists x(\tilde{P}(x) \wedge T(x))$. Here $\tilde{P}$ is a formula defining all the (codes of) formulas in $\mathbf{P}$. In other words, internal quantification over propositions can be directly translated into restricted quantification using a truth predicate.

In particular, the partial theory of self-applicable truth carries over directly to this fragment of a language with a truth predicate.[4] Associate with each internally quantified statement a restricted quantified statement in a language with a truth predicate, give a standard semantic treatment of the language with a truth predicate in a Kripke-style framework, and take the partial truth-value assignment to carry over to the language with the internal quantifier.

An alternative way of dealing with this is to directly associate the internal quantifiers with infinitary languages again, and then give a partial truth value assignment to the infinitary language (in a model). To do this, and to accommodate the fact that we are now in a type-free framework we will have to allow infinite disjunctions that contain themselves as disjuncts. In this framework a disjunction like $\bigvee_{P \in \mathbf{P}} P$ will be a disjunction over all sentences of the present language, including itself. Having infinite disjunctions that contain themselves as disjuncts is, however, not as contradictory as it sounds. Such disjunctions can be understood as being formed from non-wellfounded sets of sentences. Such sets can be specified in the framework of (Aczel 1988) or (Barwise and Moss 1996), and an inductive definition of a partial truth value assignment can then be given for them just as well as for languages containing their own truth predicate.

It is a little bit trickier in this approach to closely associate expressive strength with the size of the fragment of $L_{\omega_1,\omega}$ that a language gets associated with. But everything will be well-behaved, partly because the inductive definition closes off at ω_1^{ck}. One can make such associations by looking at what is definable in such a language, or by directly associating a Tarski style hierarchy with a Kripke style theory of truth.[5]

A.3 Conclusion

In this appendix we have looked at how we can understand the expressive strength of internal quantification over properties and propositions, and how we can deal with circularity and the threat of paradox. The points I wanted to stress were the following:

- The threat of paradoxes from quantification over properties and propositions is not unique to internalism, but applies to any other approach to such quantification as well.

[4] This, of course, can be done with a hierarchical theory of truth, too.
[5] See (Burgess 1986), (McGee 1991) and (Halbach 1997) for more on this.

- The standard ways of dealing with the paradoxes carry over to an internalist view of quantification over properties and propositions.

- The fragments of infinitary logic we get in an hierarchical account of the the strength of such quantification are fragments of the smallest admissible fragment (for recursive levels of the hierarchy).

- The size of the infinitary languages we take recourse to model the truth conditions of internal quantifiers are a nice measure of their expressive strength. The larger the fragment of infinitary logic we need to use, the more expressive the kind of internal quantifier we are dealing with.

References

Aczel, Peter. 1988. *Non-Wellfounded Sets*. CSLI Publications.

Alston, William. 1958. Ontological commitment. *Philosophical Studies* 9:8–17.

Armstrong, David. 1984. Self profile. In *D. M. Armstrong*, ed. Radu J. Bogdan. Reidel.

Armstrong, David M. 1997. *A World of States of Affairs*. Cambridge UP.

Asher, Nicholas. 1993. *Reference to Abstract Objects in Discourse*. Kluwer.

Asher, Nicholas, and Michael Morreau. 1995. What some generic sentences mean. In Carlson and Pelletier (Carlson and Pelletier 1995), 300–38.

Barwise, Jon. 1975. *Admissible Sets and Structures*. Springer Verlag.

Barwise, Jon, and Lawrence Moss. 1996. *Vicious Circles*. CSLI Publications.

Barwise, Jon, and John Perry. 1983. *Situations and Attitudes*. MIT Press.

Becker, Jerry P., and Christoph Selter. 1996. Elementary School Practices. In *International Handbook of Mathematics Education*. Kluwer.

Boghossian, Paul. 1990. The status of content. *Philosophical Review* 99:157–184.

Bonevac, Daniel. 1984. Systems of substitutional semantics. *Philosophy of Science* 51:631–656.

Boolos, George. 1984. To be is to be the value of a variable, (or some values of some variables). *Journal of Philosophy* 81:439–49.

Burgess, John. 1986. The truth is never simple. *Journal of Symbolic Logic* 51:663–681.

Burgess, John, and Gideon Rosen. 1997. *A Subject with no Object.* Oxford UP.

Carlson, Gregory N., and Francis Jeffrey Pelletier (ed.). 1995. *The Generic Book.* University of Chicago Press.

Carnap, Rudolf. 1928. *Der logische Aufbau der Welt.* Weltkreis Verlag.

Carnap, Rudolf. 1956. Empiricism, semantics and ontology. In *Meaning and Necessity.* University of Chocago Press.

Davidson, Donald. 1984a. Reality without reference. In *Inquiries into Truth and Interpretation.* Oxford UP.

Davidson, Donald. 1984b. Truth and meaning. In *Inquires into Truth and Interpretation.* Oxford UP.

Devlin, Keith. 1984. *Constructibility.* Springer Verlag.

Feferman, Solomon. 1968. Lectures on proof theory. In *Lecture Notes in Mathematics, vol 70.* Springer Verlag.

Field, Hartry. 1978. Mental representation. *Erkenntnis* 13:9–61.

Field, Hartry. 1980. *Science without Numbers.* Princeton UP.

Field, Hartry. 1986. The deflationary conception of truth. In *Fact, science and Value*, ed. G. MacDonald and C. Wright. 55–117. Blackwell.

Field, Hartry. 1989a. Platonism for cheap? Crispin Wright on Frege's context principle. In *Realism, Mathematics and Modality* (Field 1989b).

Field, Hartry. 1989b. *Realism, Mathematics and Modality.* Blackwell.

Field, Hartry. 1994a. Deflationist views of meaning and content. *MIND* 103:249–85.

Field, Hartry. 1994b. Disquotational truth and factually defective discourse. *Philosophical Review* 103:405–452.

Fodor, Jerry A. 1987. *Psychosemantics.* MIT Press.

Frege, Gottlob. 1988. *Die Grundlagen der Arithmetik.* Felix Meiner Verlag.

Gamut, L.T.F. 1991. *Logic, Language, and Meaning.* University of Chicago Press.

Gottlieb, Dale. 1980. *Ontological Economy: substitutional quantification and mathematics.* Oxford UP.

Halbach, Volker. 1997. Tarskian and Kripkean truth. *Journal of Philosophical Logic* 26:69–80.

Hofweber, Thomas, and Jeff Pelletier. n.d. Encuneral noun phrases. forthexisting.

Horwich, Paul. 1990. *Truth.* Blackwell.

Jackson, Frank. 1977. Statements about universals. *MIND* 86:427–9.

Keisler, H. Jerome. 1971. *Model Theory for Infinitary Logic.* North-Holland.

Kim, Yookyung. 1996. *A Situation Semantic Account of Existential Sentences.* Doctoral dissertation, Stanford University.

Kim, Yookyung, and Stanley Peters. 1995. Semantic and pragmatic context-dependence: The case of reciprocals. In *Proceedings of the Fifth Semantics and Linguistic Theory Conference: SALT V*, ed. T. Galloway and M. Simons. Cornell UP.

Kripke, Saul. 1975. Outline of a theory of truth. *Journal of Philosophy* 72:690–716.

Lambrecht, Knud. 1994. *Information Structure and Sentence Form.* Cambridge University Press.

Lees, Robert B. 1960. *The Grammar of English Nominalizations.* Indiana University Research Center.

Lewis, David. 1986. *On the Plurality of Worlds.* Basil Blackwell.

Linsky, Bernard, and Edward N. Zalta. 1995. Naturalized platonism and platonized naturalism. *Journal of Philosophy* 92:525–555.

Marcus, Ruth Barcan. 1993. Quantification and ontology. In *Modalities.* Oxford UP.

McGee, Vann. 1991. *Truth, Vagueness, and Paradox.* Hackett.

McGee, Vann. 1996. Logical operations. *Journal of Philosophical Logic* 25:567–80.

Parsons, Terence. 1980. *Nonexistent Objects*. Yale UP.

Pelletier, Francis Jeffrey, and Nicholas Asher. 1997. Generics and defaults. In *Handbook of Logic and Language*, ed. Johan van Benthem and Alice ter Meulen. Elsevier.

Perry, John. 1997a. Indexicals and demonstratives. In *A Companion to the Philosophy of Language*, ed. Bob Hale and Crispin Wright. Blackwell.

Perry, John. 1997b. Reflexivity, indexicality and names. In *Direct Reference, Indexicality and Proposition Attitudes*, ed. Martin Anduschus Wofgang Künne and Albert Newen. CSLI-Cambridge University Press.

Price, Huw. 1992. Metaphysical pluralism. *Journal of Philosophy* 89:387–409.

Quine, Willard V. O. 1966. On Carnap's views on ontology. In *The Ways of Paradox and Other Essays*. Harvard UP.

Quine, Willard Van Orman. 1980. On what there is. In *From a Logical Point of View*. Harvard UP.

Rochemont, Michael S. 1986. *Focus in Generative Grammar*. John Benjamins Publishing Company.

Rochemont, Michael S., and Peter W. Culicover. 1990. *English Focus Constructions and the Theory of Grammar*. Cambridge University Press.

Rosen, Gideon. 1993. The refutation of nominalism (?). *Philosophical Topics* 21:149–186.

Schiffer, Stephen. 1987. *Remnants of Meaning*. MIT Press.

Schiffer, Stephen. 1994. A paradox of meaning. *Nous* 28:279–324.

Schiffer, Stephen. 1996a. Contextualist solutions to scepticism. *Proceeding of the Aristotelian Society* 96:317–33.

Schiffer, Stephen. 1996b. Language created, language independent entities. *Philosophical Topics* 24:149–167.

Schilpp, Paul Arthur (ed.). 1963. *The Philosophy of Rudolf Carnap*. Open Court.

Schütte, Kurt. 1960. *Beweistheorie*. Springer Verlag.

Searle, John R. 1969. *Speach Acts*. Cambridge UP.

Shapiro, Stewart. 1991. *Foundationalism without Foundations: a case for second order logic*. Oxford UP.

Stevenson, Charles L. 1963. *Facts and Values*. Yale UP.

Tarski, Alfred. 1986. What are logical notions? *History and Philosophy of Logic* 7:143–54.

van Benthem, Johan. 1986. *Essays in Logical Semantics*. Reidel.

van Benthem, Johan. 1989. Logical constants across varying types. *Notre Dame Journal of Formal Logic* 30:315–42.

van der Does, Jaap. 1995. Sums and quantifiers. *Linguistics and Philosophy* 16:509–550.

van Fraassen, Bas. 1980. *The Scientific Image*. Oxford UP.

Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25:221–261.

Verschaffel, Lieven, and Erik De Corte. 1996. Number and Arithmetic. In *International Handbook of Mathematics Education*. Kluwer.

Walton, Kendall. 1990. *Mimesis as Make-Believe*. Harvard UP.

Wittgenstein, Ludwig. 1984. *Tractatus Logico-Philosophicus*. Surkamp Verlag.

Wright, Crispin. 1983. *Frege's Conception of Numbers as Objects*. Aberdeen University Press.

Yablo, Stephen. 1998. Does ontology rest on a mistake? *Proceedings of the Aristotelian Society* Supp. Vol. 72:229–261.

Yablo, Stephen. 1999. A paradox of existence. In *Empty Names, Fiction and the Puzzles of Non-Existence*, ed. Anthony Everett and Thomas Hofweber. CSLI Publications.

Zalta, Edward N. 1983. *Abstract Objects*. D. Reidel.

Zucchi, Alesandro. 1993. *The Language of Propositions and Events*. Kluwer.