STUDIES IN THEORETICAL PHILOSOPHY

vol. I

VITTORIO KLOSTERMANN

STEPHAN KRÄMER

On What There Is For Things To Be

Ontological Commitment and Second-Order Quantification

VITTORIO KLOSTERMANN

Bibliographische Information der Deutschen Nationalbibliothek
Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliographie; detaillierte bibliographische Daten sind im Internet über *http://dnb.dnb.de* abrufbar.

Gedruckt auf Alster Werkdruck der Firma Geese, Hamburg,
alterungsbeständig ∞ ISO 9706 und PEFC-zertifiziert.
Druck: betz-druck GmbH, Darmstadt
Printed in Germany
ISSN 2199-5214
ISBN 978-3-465-03868-9

To my family

Contents

Preface

This book is based on my PhD thesis, called *Second-Order Quantification and Ontological Commitment*, which was written between 2007 and 2011 at the University of Leeds. During that time, I was very lucky to work, at one stage or another, with three exceptional supervisors, excellent philosophers, and all round great guys in Joseph Melia, John Divers, and Robbie Williams. I am most grateful to them for their time, encouragement, and invaluable philosophical input. Very many thanks are due, too, to my examiners Ian Rumfitt and Jason Turner, for reading the whole thesis with so much care and, apparently, interest. I learned a great deal from their many insightful questions and criticisms.

The thesis was awarded a *Wolfgang-Stegmüller-Prize* in 2012; I thank everyone involved and especially Margret Stegmüller. Subsequent revision of the thesis was carried out while employed in the Hamburg-based part of a project on nominalizations funded by the DFG and ANR, whose support is gratefully acknowledged. I am very thankful for the continued and most valuable support and encouragement from Benjamin Schnieder, the lead of the Hamburg group. I could not have wished for a better boss.

Special thanks go to Tobias Rosefeldt and again Benjamin in their capacity as editors of this fantastic new series. It is a great honour to me that the series opens with my book.

Many others have given helpful comments on some of the material for this book. With apologies to those I have forgotten, thank you all to: Ross Cameron, Katharina Felka, Nick Haverkamp, Mirja Holst, Øystein Linnebo, Alex Oldemeier, Howard Peacock, Peter Simons, Robert Schwartzkopff, Moritz Schulz, Alex Steinberg, Duncan Watson, Nathan Wildman, and Timothy Williamson.

Chapter 2 is based on my article 'Implicit Commitment in Theory Choice', *Synthese* (2014) 191:2147–2165. Section 4 of chapter 8 is based on my paper 'Semantic Values in Higher-Order Semantics', *Philosophical Studies* (2014) 168:709–724. I thank the editors concerned for permission to reuse the material in this book.

My biggest debt, however, is to my family, and it is to them that I dedicate this book. My parents and siblings have provided excessive amounts of more than kind of support (moral, philosophical, financial, orthographic, . . .), for which I am very grateful. Most of all, I thank my wife Anka, among many other things for proofreading the entire manuscript and suggesting innumerable improvements to form and content. The book, and everything else, would be much worse without her.

S. K.

Introduction

I am sitting in a pub with a pint of ice-cold lager, and I say to myself: This beer is cold. What does this statement tell us about the world? For one thing, it tells us that there is such a thing as this beer. For another, it tells us that there is at least one thing which is cold – to wit, this beer. But the statement tells us even more. In addition to these two facts, which concern what sorts of things there are, it also tells us something about what there is for these things to be. In particular, it tells us that there is such a way for a thing to be as cold. And it also tells us that there is at least one way for a thing to be such that this beer is that way – to wit, cold. Or so I would say. You might say something slightly different instead. You might say that in addition to there being this beer, and there being something cold, what my statement tells us is that there is such a thing as the property of being cold. And it also tells us that there is at least one property that this beer instantiates – to wit, the property of being cold.

This book is about the relationship between these last two kinds of claims, i.e. claims about what there is for things to be, and claims about what properties (or similar entities) there are. More specifically, the book tries to answer these questions: Are statements about what there is for things to be just disguised versions of statements about what properties there are? Or is there perhaps some important difference between the two kinds of claims? Call this *the Disguise Question*. If there is a difference, do claims of the first sort nevertheless commit one to corresponding claims of the second sort? That is, if one endorses the claim that there is such a way for things to be as cold, does one also have to hold that there is such a thing as the property of being cold? Or could one perhaps consistently accept the former claim while rejecting the latter? Call this *the Commitment Question*. The view I shall defend is that there are indeed important differences between the two kinds of claims, and that claims of the first kind do not commit one to corresponding ones of the second kind.

How do these issues relate to the terms that make up the subtitle of this book, 'ontological commitment' and 'second-order quantification'? A claim concerning what things there are, such as the claim that there is something – namely, this beer – which is cold, is called a *first-order quantification*. A claim concerning what there is for things to be, such as the claim that there is some way for a thing to be – namely, cold – such that this beer is that way, is called a *second-order quantification*. So our Disguise Question can also be described as the question whether second-order quantification is just a special kind of first-order quantification.

Ontology is the discipline devoted to the question what sorts of things there are. A first-order quantification is therefore an ontological claim. It says that there is something of a certain sort – a property, say, or a cold beer – and thereby constitutes a partial answer to the ontological question. Now suppose that endorsing a second-order quantification *commits* one to also endorsing a corresponding first-order quantification. So if one accepts that there is a way for things to be such that my beer is that way, one also has to accept, on pain of inconsistency, that there is a property which my beer instantiates. Then, even if the second-order quantification is not itself an ontological claim, it carries the same *ontological commitment* as its first-order counterpart, which is an ontological claim. So our Commitment Question can also be described as the question whether second-order quantifications carry the same ontological commitments as their first-order counterparts.

Having described the questions that this book investigates, as well as the answers it defends, it seems appropriate to spend a few words to try and explain why one might reasonably care enough about the matter to make it worth reading (indeed, writing) a whole book to move from the questions to the answers. The next section is my attempt to do that. In the subsequent sections of this introduction, I explain the topic and main thesis of this book in more detail, and I will give a sketch of how the overall argument unfolds over the course of the book. I hope this will be of some use to the reader in helping him or her to maintain an idea of where I am going, and why.

0.1. Why Bother?

What does it matter whether second-order quantifications are just disguised first-order quantifications over properties? One way to answer this question would be to describe some important implications for connected, contentious topics in logic, semantics, and metaphysics. But this would take us rather far afield, and it would only provide somewhat indirect reasons to care about my topic.[1] I think that the question is also interesting in its own right, and so I shall focus on trying to explain why. It is worth emphasizing first of all that although terms like 'first-order quantifier' and 'second-order quantifier' describe certain linguistic, representational devices, the question is not, or not just, about language and the way we represent the world. It is also about the world's fundamental, general structure.

To see why, imagine that you set out, with the ambition typical of us metaphysicians, to give a complete account of the world. Among other things, such an account will have to say what sorts of things there are. You might start writing up a list of individual *things*, including my pint of lager, your neighbour's cat, the Eiffel Tower, and so on. Call this your *t-list* ('t' for *t*hings). You might also write up a separate list of the *sorts of things* to which the individuals on your list belong: cold things, furry things, monuments, and so

[1] Some of the connections the topic has to other debates will surface later on, though.

on. Call this your *st-list* ('st' for *s*orts of *t*hings). By means of these two lists, you can take account of facts like that there is such a thing as my pint of lager, and facts like that there is at least one object that is cold.

But what about the fact that there is such a way for things to be as cold, and the fact that there is at least one way for things to be such that my pint is that way? How do you take account of these? It depends. Suppose that second-order quantifications are just disguised first-order quantifications. Then these facts are simply the facts that there is such a thing as the property of being cold, and that there is at least one property that my pint instantiates. To take account of them, you simply write 'the property of being cold' as a further entry on your t-list, and you add 'properties instantiated by my pint' to your st-list.

But now suppose that second-order quantifications are not just disguised first-order quantifications, that they are a separate, sui generis kind of quantification. Then it seems that to take account of the fact that there is such a way for things to be as cold, you have to start a new list, on which you write *ways for things to be*: cold, hot, small, tall, yellow, red, and so on. Call this your *w-list* ('w' for *w*ays). And to take account of the fact that there is at least one way for a thing to be such that my pint is that way, you start a fourth list, on which you write the *sorts of ways for things to be* that are instantiated by the ways for things to be on your previous list: ways my pint is, ways the Eiffel Tower is, and so on. Call this your *sw-list*. If you proceed like this, you present the world as being different from how the previous account has it. In particular, you present it as having a richer, more complicated fundamental *structure*.[2] And to me, at any rate, whether the world has or lacks this kind of structure seems a worthwhile question.

Setting aside whether you should start new lists to do justice to your second-order facts, it might be suggested that even if you did that, you would also have to add something to your t-list and your st-list. That is, given that there is such a way for a thing to be as cold, whatever else you did, you would have to add the property of being cold to your t-list. And given that there is at least one way for a thing to be such that my pint is that way, whatever else you did, you would have to add the sort properties instantiated by my pint to your st-list. That suggestion is effectively the claim that second-order quantifications carry the same ontological commitments as their first-order counterparts. Why should you care about this issue, independently of whether second-order quantifications are just disguised first-order quantifications?

To ask what the ontological commitments of a statement are is to ask, roughly, what ontological claims are true according to that statement. There are at least two kinds of situations in which it is important to know what the ontological commitments of a given statement are. The first situation is one in which you know, or take yourself to have good reason to believe, that the statement in question is true. Such a reason might, for instance,

[2] All the more so, if you extend this view of second-order quantification to even higher orders of quantification. (I explain below what third- and higher-order quantifications look like.)

consist in the statement's being part of a theory generally accepted by scientists. Knowing what its ontological commitments are then allows you to draw conclusions about what exists, or what you may take yourself to have good reason to believe exists. The notion of ontological commitment accordingly has an important place in the *methodology of ontology*.

The second kind of situation in which it is useful to know what the ontological commitments of certain statements are is when we are unsure which of several competing theories to prefer, i.e. in situations of theory choice. According to the widely held methodological maxim called *Occam's Razor*, given two competing theories that are otherwise equally good, we ought to prefer the one which carries less substantive ontological commitments. If the theories are not equally good, the one that makes more substantive ontological commitments is at least in this respect at a comparative disadvantage, and we have to see whether the benefits it may offer in other respects outweigh the additional ontological cost it carries. In a slogan, the suggestion is that *ontological parsimony* is a theoretical virtue. The notion of ontological commitment accordingly has an important place in general *scientific methodology*.

These two observations provide some reason to be interested in the ontological commitments of statements in general. There are also reasons that pertain more specifically to the ontological commitments of second-order quantifications. These come out especially vividly in the context of a particular proposal for a methodology for ontology, namely that of Willard van Orman Quine. In idealized form, that methodology takes roughly the following shape: First, we take our best, total theory of the world. Second, we reformulate it in our language of regimentation: a mathematically precise formal or semi-formal language to which we can apply the tools of formal logic and thereby work out the deductive consequences of our regimented theory. Third, we determine what sorts of things the resulting formal theory is ontologically committed to. Our view of ontology should then be aligned with what our regimented best theory is committed to. If it is committed to Fs, we should accept that there are Fs, and if it is not so committed, then we need not, perhaps ought not, believe that there are Fs.

Quine is explicit about what kind of formal language he has in mind for the purpose of regimentation, namely an extensional first-order language. We shall later see exactly what such a language looks like. For now it suffices to note that such a language allows us to express claims concerning what individuals there are, and what sorts they instantiate, but nothing else. There are of course other formal languages that one might use to regiment total theory. For instance, there are languages that include modal operators expressing necessity or possibility, or tense operators allowing the expression of what once was and what will be. More to our point, there are second- and higher-order languages, allowing us to express claims, among other things, about what there is for things to be. So the question arises whether the results we obtain by applying the Quinean methodology remain the same, no matter whether we choose a first- or a second-order language of regimentation. If second-order quantifications carry the same ontological commitments as their first-order counterparts, the answer would appear to be positive. But if they do not, it may well be negative.

0.2. Ontological Commitment

The term 'ontological commitment' was coined by Quine in his famous paper 'On What There Is'. Already in Quine's writings, and even more so in the subsequent literature, the term and its derivatives are used with a kind of systematic ambiguity. For ontological commitments are ascribed to various kinds of things, including sentences, propositions, theories, discourses, beliefs, persons, and assertions. And what it is for, say, a person to be committed to things of some kind is of course not the same as what it is for a theory to be committed to things of that kind.[3]

We shall mainly be concerned with commitments of sentences, theories, and persons.[4] The things that may literally be said to have commitments, on any ordinary understanding of the word, are of course persons rather than sentences or theories. Theories and sentences might then be said to carry a certain commitment in the derivative sense that by endorsing the theory or sentence, a person would incur that commitment. It turns out to be theoretically more convenient, however, to start with a notion of what a theory or sentence is committed to. For what a person is ontologically committed to is, roughly, what she ought to believe there is, given her overall system of beliefs. If we try to identify, on this basis, what her commitments are, we have to look at which attitudes she takes towards what sentences. In particular, we have to determine what there is according to the sentences she accepts as true. So the central notion is that of what there is according to some sentence or theory. Therefore, we shall follow Quine in using the term 'ontological commitment' primarily for this notion (cf. e.g. Quine 1951a, 203). Restricting attention for the moment to theories, our official gloss on the notion of commitment is then simply this:[5]

(Gl. OC) To say that a theory is ontologically committed to an F is to say that according to that theory, there is an F.

A theory's commitments are then inherited by those endorsing it. If I accept as true a theory that is committed to Fs, then I am committed to Fs.

3 When in what follows I speak simply of commitments, I always mean *ontological* commitments, unless otherwise noted.

4 Strictly speaking, linguistic entities like sentences can of course carry commitments only relative to a language to which they belong. For brevity's sake, I shall often leave relativization to a language implicit.

5 'Gl. OC' stands for 'Gloss on Ontological Commitment'. (Gl. OC) is of course a schema, not a sentence, and as such does not strictly speaking *say* anything. The intended import of the schematic gloss is that each of its instances expresses a truth and serves to partly elucidate the notion of commitment. I shall frequently be resorting to schematic discourse in discussing ontological commitment.

On the ordinary every-day understanding of 'theory', the converse does not seem plausible. Surely, I can be committed to, say, there being a pen on my desk, without there being some *theory* that I accept according to which there is a pen on my desk. However, as is common especially in philosophical and mathematical logic, I shall be employing a much more liberal – and, admittedly, artificial – understanding of the term 'theory'. On this use of the term, *any* set of sentences (of some interpreted language) qualifies as a theory. On this liberal understanding, it does seem plausible that I am only committed to an F if there is some theory that I accept and according to which there exists an F. In addition, on this liberal understanding, we may quite naturally identify the commitments of a single sentence with those of its singleton set.

Note that my use of 'theory' is unlike the probably most common use of the term in logic in that I do *not* require that a set of sentences be closed under logical consequence if it is to count as a theory.[6] There are mainly two reasons for this. Firstly, in practice, the things we want to assess for comparative ontological parsimony, and whose ontological commitments we therefore aim to work out, are the kinds of things we find written down in a book or an article, which seem more naturally modelled by the set of sentences that are thus written down than by the closure of such a set under logical consequence. Secondly, and more importantly, if we were to require closure under logical consequence, we could not adequately represent as a theory any view that endorses a mistaken logic. However, the dispute around which our investigation centers, i.e. the dispute over the legitimacy and nature of higher-order quantification, frequently involves disagreements over the correct view of logical consequence. Presumably, therefore, some of the views we shall be considering and accordingly have a need to adequately represent in some theory, do endorse a mistaken logic.

Having thus clarified the question of the *bearers* of commitment, i.e. the things picked out by terms in the left-hand argument place of '… is ontologically committed to —', we now turn briefly to the *contents* of ontological commitments, specified by the expressions filling the right-hand argument place. English grammar allows us to put a considerable variety of types of expressions into that argument place, including for example 'a number', 'numbers', 'the natural numbers', 'countably many numbers', 'things outnumbering the natural numbers'. For our purposes, it is better to be a lot more restrictive. The general implicit assumption in the literature is that the content of an ontological commitment can be adequately represented by a monadic first-order predicate. Of the phrases just mentioned only the first two, 'a number' and 'numbers', are such that their content can plausibly be represented by such a predicate. Thus, I shall count as legitimate ascriptions of commitment only sentences in which the right-hand argument place of '… is committed to —' is filled with an expression of one of these kinds.[7]

6 A set of sentences is closed under logical consequence iff (if and only if) it has all its logical consequences as members.

7 I do not mean to imply that commitment-ascriptions of the other kinds do not make proper sense

This convention has the effect of ruling inadmissible ascriptions of what we might call singular, or particular commitments to individuals as opposed to our predicational, or generic commitments to entities of some specific kind. For the former commitments are ascribed by filling the right-hand argument place of 'is committed to' with a name (singular term), such as 'Socrates', or 'the Eiffel Tower'. Using the identity-predicate, we can ascribe closely related predicational commitments – a commitment to an object identical to Socrates, for instance – and we may treat phrases like 'committed to Socrates' as short-hands for the corresponding predicational ones. Singular commitments will play almost no role in this thesis, and we may forget about them until they do.

It is worth noting at this point that the linguistic constructions Quine tends to employ when speaking generally of ontological commitments, rather than about a specific example such as the commitment to a prime number, seem at first glance not to fit too well with the schematic talk of commitment to Fs that I tend to use (and that most contemporary writers tend to use as well). Quine usually speaks of the entities to which we are committed, or of us being committed to entities of a given sort. The former phrase seems inadequate where the commitment in question is not satisfied, for then there are no entities for us to be committed to. The latter phrase seems better, and I shall use it myself from time to time. I suspect, however, that the best way of regimenting it that Quine can happily avail himself of is again in terms of our schematic talk. The schematic regimentation, at any rate, seems to be adequate to Quine's intentions and to be well-motivated by the specific ascriptions of commitment he mentions. Moreover, in comparison to the alternative constructions, it seems less problematic. I shall therefore continue to employ it, and I shall assume (with everybody else) that by doing so I am not parting company with Quine over anything significant (or that if I am, the result is an improvement rather than the opposite).

Let me now briefly comment on two features of our gloss (Gl. OC) on the notion of theory's commitments that concern the *extension* that notion is thereby assigned. The first such feature is the use of the construction 'according to'. One might worry that by relying on this phrase, (Gl. OC) assigns an unduly *narrow* extension to the notion of commitment.[8] For example, consider a theory including (regimentations of) the sentences 'every number has a number as its successor', 'no number is its own successor', and 'no two numbers have the same number as their successor'. Intuitively, it seems plausible to describe such a theory as being committed to an infinity of numbers. But it seems much less clear whether we could correctly say of such a theory that *according to* it, there are infinitely many numbers. Indeed, it may not be possible to even *express* this latter claim in the language of the theory.[9]

or are not interesting. The point is simply that for most of the thesis, we may safely set aside those ascriptions of commitment that I have just ruled out, and so it is more convenient to restrict our use of 'is ontologically committed to' to those cases we shall mostly be concerned with.

8 Thanks to Ian Rumfitt for raising this issue. The example to follow is also essentially his.

9 This particular example raises some subtle issues. First, I should note that the commitment to an

My own sense is that the standards we apply in ordinary uses of 'according to' vary quite substantially across different contexts. In particular, I do not think that we are always as strict as the above argument assumes in assessing whether some claim holds according to, say, a newspaper article, a witness's report, or that politician's statement. At any rate, it will be one of the tasks of later chapters to specify a suitable extension for the notion of a theory's commitments in a more precise manner. In pursuing this task, I will refrain from placing weight on intuitions from ordinary usage concerning just how demanding the conditions are for a claim to count as holding according to that theory. Rather, the considerations I shall adduce here will be based on theoretical constraints arising from the specific role we wish the notion of commitment to play. So I think that the use of 'according to' in (Gl. OC) does not prejudge important issues about commitment that should be left open for subsequent investigation.

The second feature of (Gl. OC) that I wish to comment on is the use of 'there is'. As far as I can see, it should be conceded on all sides that the phrase 'there is' is frequently used in an ontologically committing way. At the very least, in ontological debates, the phrase is often so used. Indeed, in such debates, it is regularly used with the express intention of indicating an ontological commitment on the part of the speaker. I am not aware of any good reason to think that such uses of 'there is' are generally unsuccessful. It is not clear, however, that any and all uses of 'there is' are in this way ontologically committing. And so it may be that some instances of the schematic gloss (Gl. OC) admit of a reading on which 'there is' is not ontologically committing, and the instance of (Gl. OC) therefore false. I therefore want to highlight that (Gl. OC) should be read with an interpretation of 'there is' that is in play in ontological debates, and on which it therefore expresses the ontologically important notion of existence.[10]

Given our gloss (Gl. OC) on ascriptions of commitment, everyone should agree that a theory's including among its members a sentence which means *there is an F*[11] is a sufficient condition for a theory's carrying ontological commitment to Fs. If, and only if, a theory is committed to Fs for just this reason, I shall say that it is *explicitly* ontologically committed to Fs:[12]

infinity of numbers is not strictly speaking an ontological commitment, since the phrase 'an infinity of numbers' is not intended to function semantically as a predicate. There is a closely related proper ontological commitment, however, which could be used instead in the argument, namely the commitment to a number which is distinct from infinitely many numbers. Now there is a sense in which indeed no first-order language can express the notion of infinity: no first-order theory is true in all and only models with an infinite domain. However, it is unclear whether one who, like Quine, dismisses higher-order resources as illegitimate, should concede that there is a commitment here – a commitment to a number distinct from infinitely many numbers – that cannot be expressed in a first-order theory. Fortunately, there is no need to decide the issue here.

10 This still assumes that there is such a thing as *the* ontologically important notion of existence, an assumption that has been challenged by a number of authors. I briefly talk about their views in chapter 1. By and large, however, I assume without argument that they are mistaken.

11 Note that the sentence 'there is an F' is here used not in the ordinary way, i.e. to say that there is an F, but to pick out its meaning. I shall always use italics to mark this non-standard use of expressions.

12 'Df. EOC' stands for 'Definition of Explicit Ontological Commitment', accordingly for 'Df. IOC'.

(Df. EOC) A theory T is *explicitly* ontologically committed to Fs $\leftrightarrow_{df.}$:
T includes some sentence that means *there are Fs*

Commitments which are not explicit we will accordingly call 'implicit'.

(Df. IOC) A theory T is *implicitly* ontologically commited to Fs $\leftrightarrow_{df.}$:
T is ontologically committed to Fs $\wedge$
T is not explicitly ontologically committed to Fs.

By our assumptions, explicit commitment entails commitment: if a theory is explicitly ontologically committed to Fs, then it is ontologically committed to Fs.

In order to decide what a given theory, in a given language, is explicitly committed to, we need to know which of the expressions of the language express existence, or more cautiously: express the notion of existence that ontologists are interested in, that is in play in ontological debates over the existence of things of some kind. On the Quinean methodology described earlier, the notion of ontological commitment is applied only to regimented theory, i.e. collections of sentences in a canonical, formal language. Since the symbolism employed in the kind of language that Quine favours for the purpose of regimentation will be used quite frequently throughout this book, I want to take this opportunity to briefly explain that symbolism.

Quine's favoured languages are name-free, extensional, first-order languages without function symbols. Their vocabulary includes, firstly, a countably infinite stock of *individual* or *first-order variables*. We use single lowercase letters, usually from the end of the alphabet, as first-order variables, and augment them with subscripted numerals if necessary. First-order variables are essentially formal counterparts of singular pronouns such as 'it' in natural languages. Secondly, we have the usual stock of sentential connectives and quantifiers which are to be interpreted according to the following table:

$\neg$	it is not the case that
$\wedge$	and
$\vee$	or
$\rightarrow$	if ... then
$\leftrightarrow$	if and only if
$\exists x \ldots x \ldots$	there is something x such that ... x ...
$\forall x \ldots x \ldots$	everything x is such that ... x ...

Third, we have a suitable stock of (interpreted) predicate letters 'F', 'G', 'H', etc., again adding numerals as subscripts if necessary. These correspond to expressions of a natural language like 'is red', 'dances', 'stands next to', 'gave ... to', or 'is heavy and tall'. That is, they correspond to expressions that can be obtained from a sentence by deleting one or more singular terms, like names, pronouns, or definite descriptions. It will always be clear from the context what the adicity of a given predicate is, i.e. the number of expressions it combines with to form a complete sentence. Finally, we have the special purpose predicate '=' to be read as 'is identical with', and the parentheses '(' and ')'.

Clearly, of the expressions in such a language, only the first-order existential quantifier '$\exists x$' could plausibly be taken to express the ontologists' notion of existence. And according to Quine, it does indeed express that notion:[13]

(EE) The first-order existential quantifier '$\exists x$', on its standard interpretation, expresses the ontologically important notion of existence.

This claim then straightforwardly yields the result that any first-order theory that contains the sentence '$\exists x\ Fx$' is explicitly ontologically committed to Fs.[14] For any theory that includes '$\exists x\ Fx$' thereby includes a sentence which means *there are Fs*, and is therefore such that according to it, there is an F.

The truth of (EE) is one of two major Quinean assumptions that I make in this book without offering a comprehensive defence of them.[15] The second such assumption is that the mere presence of a predicate 'F' in a first-order sentence such as '$\exists x\ Fx$' does not render the sentence ontologically committed to a corresponding object such as the property F-hood, or the set of Fs. There is a fairly strong argument for that claim: If '$\exists x\ Fx$' were committed to F-hood, then it would in general be safe to infer from a sentence including 'F' that there is a property instantiated by all and only Fs. Russell's paradox shows that this kind of inference is not in general safe: it leads one to assert the existence of a property instantiated by all and only properties not instantiating themselves, which, on pain of contradiction, cannot exist.

0.3. Second-Order Quantification

The term 'second-order quantification' is used in very different ways in the literature. It is sometimes used to mean, roughly, *quantification over such properties as may be instantiated by individuals*, and it is sometimes used to mean, roughly, *quantification over the powerset of the first-order domain* (or the sets of tuples of the members of the first-order domain). We might call the first one the *metaphysical* use of 'second-order quantification', and the second one the *semantic* use. The term is also often used to mean, roughly, *quantification into the position of (first-order) predicates*, or *quantification that stands to (first-order) predicates as ordinary first-order quantification stands to names (or individual constants)* – we

13 'EE' stands for *E*xpression of *E*xistence.

14 Here and in what follows, when I speak of a theory containing, or including, a given sentence, I mean that the sentence is a *member* of the theory in question.

15 Chapter 1 offers a partial defence of this assumption. I should emphasize that although for the most part of the book, I proceed on the assumption that (EE) is true, many significant results do not depend on this assumption, and for many results that do, there are related and, I think, still interesting and significant claims that would follow even if (EE) were false.

might call this a *syntactic* use of the term.[16] The differences between these uses are important, for it is a substantial question whether all, or even just some of the expressions qualifying as second-order quantifiers on one of those uses also qualify as second-order on the other uses. Throughout this book, the term 'second-order quantification' is used in the last, syntactic way.

The argument

(1) Bob smokes and Bob drinks. So, someone is such that he smokes and he drinks.

and its semi-formal counterpart

(2) Bob smokes and Bob drinks. So, $\exists x$ (x smokes $\wedge$ x drinks).

employ first-order quantification, that is quantification into *name* position. The conclusion is obtained from the premise by replacing both occurrences of the name 'Bob' by a suitable variable or pronoun, and prefixing the result with the quantifier phrase 'someone is such that' or '$\exists x$', binding that variable or pronoun. The inference brings out that the conjunctive premise describes the same person as both a smoker and a drinker. Generally, first-order quantifications are used to express general, or unspecific claims concerning what there is, what kinds of things there are – in this case, someone who both smokes and drinks.

Second-order quantification is supposed to allow us to make analogous inferences with respect to a *predicational* element of a sentence. That is, it is supposed to allow us to formulate arguments such as

(3) Bob smokes and Bill smokes. So, $\exists X$ (Bob X $\wedge$ Bill X).

The quantification in (3) is quantification into predicate position. The conclusion is obtained by replacing both occurrences of the predicate 'smokes' in the premise by a suitable variable and prefixing the result with the quantifier phrase '$\exists X$', binding that variable. The inference is supposed to bring out, roughly speaking, that the conjunctive premise describes both Bob and Bill in the same way. Generally, second-order quantifications are supposed to express general, or unspecific claims concerning what things are like, what there is for things to be or do – in this case, something that both Bob and Bill do, namely smoke.

Judging from surface appearances, some natural language sentences employ quantification almost of that kind. I already employed some pertinent examples at the very beginning of this chapter. We shall now have a closer look at a few more instances of the phenomenon. Consider the following sentence:

[16] We shall see shortly, though, that in one respect this label is perhaps misleading. Unfortunately, I don't have a better one.

(4) If Bob smokes and Bill does not smoke, then Bob does something Bill does not do (namely, smoke).

The occurrence of 'something' in the consequent of (4) is associated with the position of the verb 'smoke' in the antecedent, as witnessed by the 'namely'-rider. Or consider this reformulation of (4)

(5) If Bob is a smoker and Bill is not a smoker, then Bob is something Bill is not (namely, a smoker).

in which the occurrence of 'something' is associated with the position of the general term 'a smoker'. On my use of 'predicate', neither 'smoke' (as used in (4)), nor 'a smoker', are predicates; I apply the term only to expressions that are syntactically like 'smokes', 'does (not) smoke', 'is (not) a smoker', etc., in that they produce sentences when combined with one or more names. Nevertheless, they seem to be very closely related to predicates – so much so that the difference between, say, 'a smoker' and 'is a smoker' is often harmlessly neglected, and the term 'predicate' used for both. The quantifications in (4) and (5), therefore, seem to be at least very close relatives of the formal quantification in (3).

My reason for not broadening my use of 'second-order quantification' so as to properly include the type of quantification in (4) and (5) is practical, rather than principled. For most of the time, I will be focusing on artificial languages which, as is standard, have expressions syntactically like 'is a smoker' or 'smokes' rather than 'a smoker' or 'smoke', and the quantifiers I will be discussing will accordingly be associated with the position of these expressions. It is therefore more convenient to define 'second-order quantification' by 'quantification into predicate position' rather than trying to find a natural definition that covers the natural language quantifications as well. Speaking loosely, we may still call the quantifications in (4) and (5) (putative) instances of second-order quantification in natural language, taking it to be understood that what we mean is 'the closest relative(s) in natural language to second-order quantification'.

Calling my use of 'second-order quantification' *syntactic* is misleading to the extent that it suggests that the classification of an expression as a second-order quantifier has no immediate consequences concerning its semantic features. That is not the case. I count an expression a *predicate* just in case it is both syntactically *and semantically* like predicates in natural language. Second-order quantification, on my use of the term, is therefore quantification into the position of expressions of a specific semantic kind. Moreover, I count a sentence a quantification into the position of expressions of a certain category just in case it is both syntactically *and semantically* related to its instances – sentences obtained by deleting quantifier and variable and filling the resulting gaps with an expression of the pertinent category – in the way in which ordinary first-order quantifications are related to their instances. I shall sometimes summarize these constraints on second-order quantification by saying that second-orderquantification is supposed to be *predicational*.[17]

One last remark to complete this preliminary elucidation of my notion of second-order quantification: Talk of quantification of the first and the second order raises the question whether there might be such a thing as third- and higher-order quantification. First-order quantification is quantification into name position, and second-order quantification is quantification into the position of expressions which form sentences when combined with names as their arguments, i.e. predicates. Third-order quantification, by analogy, is quantification into the position of expressions which form sentences when combined with predicates as their arguments. Calling ordinary predicates *first-level* predicates, we may call expressions forming sentences when combined with first-order predicates *second-level* predicates. Generalizing, we call an expression an $(n+1)$-th level predicate just in case it combines with n-th level predicates to form sentences. Quantification into the position of n-th level predicates is quantification of $(n+1)$-th order. Quantification of orders higher than the second will only play a comparatively minor role in my discussion, however.

Since I have in the previous section described in some detail what a formal language looks like that is of the kind favoured by Quine, let me now also briefly state how a standard *second-order* extension of such a language varies from it. It contains just two new kinds of expressions. Firstly, we have for every (finite) adicity of predicates a countably infinite stock of *second-order* or *predicate variables*. We write these using capital letters from the end of the alphabet, subscripting with numerals if necessary. Secondly, we have the second-order existential and universal quantifiers '$\exists X$' and '$\forall X$'. These behave just like their first-order counterparts, except of course in that they combine with second-order rather than first-order variables. It is unfortunately not straightforward to offer direct English translations of these in the way I have done for the first-order quantifiers. Indeed, I later spend two entire chapters examining attempts at giving precise natural language translations of these quantifiers. The examples I gave above of broadly equivalent semi-formal second-order quantifications and natural language counterparts provide at least a rough grasp of the meaning of these quantifiers, which will have to suffice for the time being.

0.4. The Thesis

There are two main theses that I defend in this book. The first is that second-order quantifications are not just disguised versions of their first-order counterparts. The second is that second-order quantifications do not carry the same ontological commitments as their first-order counterparts. The second thesis entails the first: if second-order quantifi-

17 Others use 'predicative', which I shy away from merely because it sounds like the opposite of 'impredicative', which is standardly used to describe a certain kind of circularity in definitions of second-order notions.

cations carry different commitments than their first-order counterparts, they cannot just be disguised versions of the latter. I shall therefore generally focus on the second, stronger thesis. Since I find it easier to come up with a good label for the claim that I reject rather than the one I defend, let us call *Ontological Collapse* the claim that second-order quantifications carry the same ontological commitments as their first-order counterparts. My aims in this section are: firstly, to describe in slightly more detail what that claim is; secondly, to give some indication of what its falsity might mean for ontology, and thirdly, to explain why I take the burden of proof to lie with my opponent who wishes to maintain Ontological Collapse.

To explain more clearly what the thesis of Ontological Collapse is supposed to say, I first need to clarify the notion of a *first-order counterpart* of a given second-order theory. Consider the second-order quantification: Tom is everything a musician should be. Its import, we might say, is to rule out all the scenarios in which there is a true sentence of the form 'A musician should be F, but Tom isn't F'. Assuming that there are any such things as properties, it seems that roughly the same effect can also be achieved by means of the following first-order quantification over properties: Tom has every property a musician should have. For consider any scenario that is ruled out by the second-order quantification. For explicitness, suppose that a musician should be creative, but Tom isn't. Then given the existence of properties, it seems to follow that a musician should have the property of creativity, but Tom doesn't have it. In that case, the scenario is also ruled out by our first-order quantification.[18]

Properties are not the only candidates for objects one might appeal to in emulating the effect of second-order quantification. Sets might serve the same purpose, as might certain kinds of functions. Essentially, what is required is a kind of object that can be correlated with what there is for objects to be in such a way that any distinct such objects are assigned to any distinct things for objects to be. I shall use the term 'second-order entity' to speak indiscriminately of any such candidate objects, leaving open their exact nature.

By a first-order counterpart of a second-order theory T I shall then mean a first-order theory that says essentially the same as T except in that it uses first-order quantification over second-order entities to achieve the effect of the second-order quantifications in T. This is still somewhat vague, but deliberately so. For just how close such a theory can possibly come to saying the same as the second-order theory, and whether it can possibly emulate exactly the effect of the latter's second-order quantifications, turns on difficult and controversial issues that will be examined later on. We therefore do better here do leave such matters open and merely require a rough match.

[18] If the effect achieved by the first-order quantification is to be the same as that achieved by the second-order quantification, it must also not rule out more scenarios than the latter. This is unproblematic, though. If there is a property that a musician should have but Tom lacks, then there is someting a musician should be – namely, exemplifying that property – that Tom isn't.

Now given (EE), thanks to its including existential first-order quantifications over second-order entities, a first-order counterpart of a typical second-order theory will be ontologically committed to second-order entities. The distinctive claim of the proponent of Ontological Collapse is that these same ontological commitments are already carried by the second-order theory. Our preliminary precisification of the thesis of Ontological Collapse therefore reads as follows:[19]

($\downarrow_{\text{prelim}}$) For every second-order theory T, there is a first-order counterpart T' such that T and T' carry the same ontological commitments.

It should be noted that it is only under the assumption of (EE) that that ($\downarrow_{\text{prelim}}$) captures the thesis that I shall argue against. For just what the claim ($\downarrow_{\text{prelim}}$) amounts to depends on what ontological commitments are incurred by first-order quantifications. If, for instance, contrary to what Quine thinks, first-order quantifications do not in general carry ontological commitments, ($\downarrow_{\text{prelim}}$) may come out true even if, as I maintain, second-order quantifications do not introduce any ontological commitments. It therefore seems worth briefly explaining what my main thesis comes to if (EE) turned out to be false. Even if first-order quantifiers do not carry *ontological* commitments, we may describe them as carrying *some kind* of commitment – call it simply *first-order* commitment. My main thesis is then that the following variant of ($\downarrow_{\text{prelim}}$) is false:

($\downarrow_{\text{prelim}'}$) For every second-order theory T, there is a first-order counterpart T' such that T and T' carry the same first-order commitments.

For the remainder of this introduction, I shall be proceeding on the Quinean assumption of (EE), so that ($\downarrow_{\text{prelim}}$) captures the view that I want to reject.

There are numerous ways in which the truth or falsity of ($\downarrow_{\text{prelim}}$) may impact ontological debates. For the purpose of illustration, I shall simply sketch one pertinent case, as it presents itself in the context of the Quinean methodology that I have described above. Suppose that we wish to produce a regimented version of ordinary arithmetic, to then see what ontological commitments it incurs. We may then want to include in our theory a sentence that captures the principle of mathematical induction for natural numbers. Informally, we can state this principle as follows:

(6) For all properties P, if 0 has P, and every number has P if its predecessor has P, then every number has P.

The most natural rendering of (6) in a formal first-order language of the kind Quine favours involves quantification over properties (sy is the successor of y):

[19] Over the course of the book, we shall have occasion to consider various modifications and refinements of this thesis. To have a systematic way of labelling them while also keeping the labels at a manageable length, I have chosen to represent the general idea of collapse by means of a downward arrow, which I hope will work okay as a mnemonic.

(7) $\forall x\,((x \text{ is a property} \wedge 0 \text{ has } x \wedge \forall y\,((y \text{ is a number} \wedge y \text{ has } x) \rightarrow sy \text{ has } x)) \rightarrow$
$\forall z\,(z \text{ is a number} \rightarrow z \text{ has } x))$.

In order for the principle to have the intended import, there must be such things as properties. Otherwise the quantification would simply be vacuously true thanks to a trivial lack of counter-examples. Our regimented theory of arithmetic thus winds up endorsing an ontology of properties. If, on the other hand, we have second-order quantifiers at our disposal, we may instead regiment (6) by:

(8) $\forall X\,(X(0) \wedge \forall y\,((y \text{ is a number} \wedge Xy) \rightarrow Xsy) \rightarrow \forall z\,(z \text{ is a number} \rightarrow Xz))$.

Now, if Ontological Collapse is true, including (8) in our theory leaves us with same ontological commitments as we incur by including (7) (or some different first-order counterpart quantifying over some other kind of second-order entitity). But if, as I maintain, Ontological Collapse is not true, then it does not. As a result, we may be able to accept our second-order regimented arithmetic without thereby having to include properties on our list of things there are.

Let me now turn to the third task of this section, namely to explain why I take the burden of proof to lie with the friend of Ontological Collapse. Consider the second-order theory that includes just the following sentence: '$\exists X\ \forall x\ Xx$'. It says that there is something that everything is. According to Ontological Collapse, it is committed to some sort of second-order entity, such as a property that is universally instantiated. That is, given Ontological Collapse, we have to say that according to our little theory, there is a property that everything instantiates. Prima facie, the sentence used to state this latter claim, i.e. 'there is a property that everything instantiates' is quite different in meaning from the sentence in our theory. Its internal syntactic structure is significantly different, and it involves various words like 'property' and 'instantiates' to which there appears to be no obvious counterpart in the second-order sentence. I therefore maintain that until we are given some positive reason to accept that according to our theory, there is a property everything instantiates, we should reject this claim. That is, to be justified in rejecting it only requires the rebuttal of any arguments offered in support of it. It does not require that we give any further positive argument for its falsity.

In view of this, my focus in the following chapters will be on what I consider to be the strongest arguments that have been, or might be, advanced in support of Ontological Collapse. If they can be refuted, as I shall try to show they can, we should reject Ontological Collapse. The next section briefly describes what the main arguments for Ontological Collapse are.

0.5. Three Arguments

There are, as far as I can see, three main arguments for the view that, surface appearances to the contrary, second-order theories ontologically collapse into first-order paraphrases.