

**Committing Ourselves to Nothing:
An anti-orthodox view of existential quantifier expressions**

**A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY**

Stephen M. Nelson

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy**

Peter W. Hanks (Adviser), Geoffrey Hellman (Co-adviser)

April, 2013

© Stephen M. Nelson 2013
ALL RIGHTS RESERVED

Acknowledgements

As this dissertation is the culmination of my time as a student in philosophy, I must acknowledge those who contributed to my various educational phases. From my undergraduate years, there were inspiring teachers like Ernie McCullough and Robert Laliberte, but most of all, my mentor Catherine Cater deserves my undying gratitude for cultivating my philosophical interests and encouraging me to go to graduate school, pushing me to get as far away from Fargo as possible.

While at Arizona State University, my cohort of graduate students were invaluable companions through the first few years, especially my great friends Nick Kroll and Mark Budolfson. Also, the instructors who took us under their wings: Tom Blackson, Stew Cohen, John Devlin, Greg Fitch, Peter French, Bernie Kobes, and Michael White.

At the University of Minnesota, another great community of students shaped my work, including Will Bausman, Jeff Johnson, Shay Logan, Leia Rollag, Jack Woods, and especially my writing group with Nathan Gutt, Jason Swartwood, and Ian Stoner. Professors who inspired or encouraged me along the way include Bill Hanson, Doug Lewis, Joe Owens, Sandra Peterson, Valerie Tiberius, and John Wallace. And most of all, my committee: Roy Cook and Geoffrey Hellman gave invaluable help and support at key points, and my adviser Peter Hanks gave me all the guidance and encouragement I could have hoped for—if I'd had my way, I would never have settled on my own view to develop and defend. He convinced me of the need to do this, made me believe that it was well within my powers to do it, and then guided me through it to the end.

Department visits by Matti Eklund, Michael Lynch, Nathan Salmon, and Ted Sider also had significant impact on this project. And in the final months, my current writing group with my colleagues Brent Braga, Mark Decker, and Brian Huschle kept me motivated to keep writing while managing a hefty teaching load.

Dedication

To my family: my parents, David and Sherri, who prepared me for this, and everything else; and my brother Matt, for making me always want to be smarter and showing me what smarter looks like.

To my girlfriend, Beck Schweitzer, for being there with me for the final year of this project, encouraging me and inspiring me, making me feel like this dissertation was something I could actually complete.

To family and friends close to me who had to suffer through being part of the life of someone working on such a long, difficult project as this.

And to the memory of Greg Fitch. Aside from being a close friend and a wonderful man and philosopher, Greg is responsible for sparking my interest in the topic of this dissertation and for guiding me through the beginning stages, helping me figure out what it is I was actually interested in.

Abstract

There is a significant difference between the words ‘is’ and ‘exists’ that has either been overlooked or under-appreciated by many philosophers. This difference comes in sentences that express existential quantification using ‘is’, ‘exists’, or their cognates, such as, “There are cookies in the jar,” or, “There exists a strange species of fish that nobody has studied yet.” Phrases such as ‘there are’ and ‘there exists’ are *existential quantifier expressions*, since they’re used to express existential quantification. The *orthodox* view of these expressions is that they are, in the words of David K. Lewis, “entirely synonymous and interchangeable”. This dissertation presents and argues for an *anti-orthodox* view of meaning of ‘there is’ and ‘there exists’. The root of the difference in meaning between the two expressions is that ‘there is’ turns out to be context-sensitive, on the model of demonstratives like ‘this’ or ‘that’, while ‘there exists’ is invariant in its meaning. These views are motivated through the introduction of a notion called ‘ontological robustness’, which helps us evaluate the level of ontological commitment in our assertions. The anti-orthodox view is defended over orthodoxy through holistic arguments that compare the virtues of each theory, including such metrics as how they fare in accounting for our stubborn desire to talk about and quantify over nonexistent objects.

Contents

Acknowledgements	i
Dedication	ii
Abstract	iii
1 Ontological Robustness	1
1.1 Introduction & Overview	1
1.2 Ontological Robustness	3
1.2.1 Definition of Ontological Robustness	4
1.2.2 Examples of ‘There is’ Claims	6
1.2.3 Examples of ‘There Exist’ Claims	12
1.2.4 Examples About Validity	14
1.2.5 What the Examples Tell Us	18
1.3 Looking Ahead	22
2 Contextualism about ‘There Is’	24
2.1 Introduction	24
2.2 Overview of Contextualism	26
2.3 Two Semantic Approaches to Quantifiers	30
2.3.1 Logical Treatment	30
2.3.2 Linguistic Treatment	36
2.4 On ‘Exists’	46

3	Unrestricted Domains and Nonexistent Objects	48
3.1	Introduction	48
3.2	Why this isn't about Quantifier Domain Restriction	48
3.3	Quantifying Over Non-Existent Objects	53
3.3.1	Non Denoting Quantifier Expressions	55
3.3.2	Using Non Denoting Demonstratives as a Model	57
3.3.3	Believing Gappy Propositions	61
3.3.4	Reference Magnetism and Speaker Intentions	64
4	Contextualism Over Orthodoxy	70
4.1	Introduction	70
4.2	Quine's View of Existential Quantifier Expressions	71
4.2.1	“On What There Is”	73
4.2.2	“Existence and Quantification”	75
4.3	Clarifying Orthodoxy and My Denial	80
4.4	Comparing Anti-Orthodoxy with Orthodoxy	82
4.4.1	Agreement With Usage	84
4.4.2	Agreement With Our Intentions	86
4.4.3	Ontological Neutrality	88
4.4.4	Semantic Simplicity	90
4.4.5	Summing up the Desiderata	91
5	Concluding and Tying Up Loose Ends	94
5.1	Introduction	94
5.2	Contrast With Epistemic Contextualism	94
5.3	Some Other Unorthodox Views	99
5.3.1	Hofweber's Ambiguity View	99
5.3.2	Graham Priest's Denial of Synonymy	102
5.4	Conclusion	105
	References	106

Chapter 1

Ontological Robustness

1.1 Introduction & Overview

There is a significant difference between the words ‘is’ and ‘exists’ that has either been overlooked or under-appreciated by many philosophers. This difference comes in sentences that express existential quantification using ‘is’, ‘exists’, or their cognates, such as, “There are cookies in the jar,” or, “There exists a strange species of fish that nobody has studied yet.” Phrases such as ‘there are’ and ‘there exists’ are *existential quantifier expressions*, since they’re used to express existential quantification. This is just what we talk about in an Intro to Logic course, when we teach the proper usage of the ‘ \exists ’ symbol in proofs and translations.

When we first introduce existential quantification to students, we do not distinguish at all between the various existential quantifier expressions. We tell them that ‘ \exists ’ is the symbol to be used whenever we express the notion of *at least one-ness*, using phrases like ‘there is’, ‘there exists’, ‘something’, etc. The important parts for the students to learn are the inferential role and the association of the quantifier with a domain. The inferential role is the group of inference rules associated with the symbol, and the domain of the quantifier is a set that is provided in some formal way (e.g. stipulating that our domain is ‘U.S. Senators’ or ‘animals on the farm’).

As we all know, though, we frequently simplify certain notions when teaching them in an introductory course. The fact that we teach Intro to Logic students that all the various existential quantifier expressions are interchangeable and work the same

way does not automatically mean that this is a truth about the English language.¹ But, as a matter of fact, the orthodox view among philosophers appears to be that the simplification we teach in an Intro to Logic course is also the correct semantic view of existential quantifier expressions.

David Lewis gives the orthodox view of existential quantifier expressions as explicitly as possible here, while denying Richard Routley’s view that distinguishes between two kinds of existential quantifiers:

Routley sees himself as defying an established orthodoxy; and I am prepared to appoint myself spokesman for the orthodoxy he defies. [...] We of the establishment think that there is only one kind of quantification. The several idioms of what we call ‘existential’ quantification are entirely synonymous and interchangeable. It does not matter whether you say ‘Some things are donkeys’ or ‘There are donkeys’ or ‘Donkeys exist’—you mean exactly the same thing whichever way you say it.²

Routley’s view that Lewis is arguing against falls within the family of ‘Meinongian’ views, which trace back to Alexius Meinong and his distinctions between different kinds of objects, some of which exist, some of which merely subsist, and some of which have no being at all. And Lewis’s view that he takes on as orthodoxy traces back (at least with respect to its orthodoxy) to W.V.O Quine’s reaction to and arguments against the Meinongian position, beginning in his “On What There Is” [36]. For the past six decades, we have had an orthodox view rooted in Quine, supported by philosophers such as Lewis and Peter van Inwagen, with the primary deniers of this orthodoxy coming from the Meinongian tradition, such as Routley, Terence Parsons, Ed Zalta, and Graham Priest.

Although I will argue here that there is a significant difference in meaning between ‘is’ and ‘exists’ in existential quantifier expressions—and thus deny the orthodox view—I do not support the view (commonly associated with the Meinongians) that there are non-existent objects, or that there are different kinds of being. The orthodox view is largely motivated by the lack of a non-Meinongian alternative; orthodoxy proponents

¹ Just think of the complexity that we leave out when first discussing the logic of the material conditional ‘ \supset ’, translating it as ‘if...then...’.

² Lewis [30], pp. 24–25

recognize that their view does not adequately describe our use of existential quantifier expressions, but they see the Meinongian view, which they think is incoherent or crazy, as the only alternative to orthodoxy.

The orthodox view has it wrong on both counts: existential quantifier expressions are not all ‘entirely synonymous and interchangeable’, and a commitment to Meinongianism is not the only alternative to the Quine/Lewis orthodoxy. I will build a case for these claims here, beginning by discussing what I call ‘ontological robustness’. This is a property of existentially quantified claims that will illuminate the difference between those that make use of ‘is’ and those that make use of ‘exists’. After I introduce and explain this notion, I will argue that the significant difference between ‘is’ and ‘exists’ is that existentially quantified claims using ‘exists’ are always ontologically robust, while those using ‘is’ sometimes are and sometimes are not ontologically robust. In subsequent chapters, I will fill out an anti-orthodox view by showing that ‘there is’ is a context-sensitive expression, but not ‘there exists’. I’ll develop this view over the next two chapters, argue for it in the fourth chapter, and tie up some loose ends in the fifth.

1.2 Ontological Robustness

Consider the sentence, “There are superheroes who can fly,” uttered by the same person, but in different settings. In one setting, our speaker is discussing comic books with his daughter. In another setting, our speaker is discussing the ontology of fiction at an APA Conference session. There is clearly a difference between these two events, but what is it? In the first case, the claim sounds natural as a part of a discussion about what kinds of superheroes our speaker’s daughter can expect to find if she reads some comic books or watches cartoons. In the second case (at the APA), the claim will be taken to have a different kind of commitment—something involving existence.

We can characterize the difference between these two scenarios as a difference in whether the existence of superheroes *matters* or not—in particular, whether it matters *to the speaker*. We can say this much without even knowing whether the speaker actually *believes* that superheroes exist or not. The speaker could have a sophisticated philosophical view about the matter, or he could have no view at all. If he does believe that superheroes exist, then his claim would not be problematic in either setting (at

least not automatically). But if he does not believe that superheroes exist, the two situations part ways. The first claim, while talking to his daughter, could still be perfectly appropriate, even if he does not believe superheroes exist. But the same claim uttered at the APA session would *not* be appropriate. In that situation, the existence of superheroes matters to the discussion at hand, and he should not quantify over them if he is not happy being committed to their existence.

1.2.1 Definition of Ontological Robustness

This feature I’ve just described is something I will call ‘ontological robustness’, since it is exhibited by claims or assertions that are robust in their ontological commitments. I will discuss only *positive* existentially quantified claims, since negative existentials (e.g., “Santa Claus doesn’t exist”) carry with them some quite different issues.³ So let’s define an assertion of a positive existential claim P by a speaker S in context C, quantifying over Fs, as ‘ontologically robust’ as follows:

S’s assertion of P is *ontologically robust* in C iff S would not sincerely assert P in C if S believed that Fs did not exist.

Since the right side of the ‘iff’ here is a counter-factual, analyzing an assertion for ontological robustness will involve looking to the closest worlds⁴ where S believes that Fs do not exist and then asking whether S would sincerely assert P or not in those worlds (in the same context as—or as close as possible to—the context in which the claim is uttered). If some of those worlds are ones where S would still assert P in C, then P is not ontologically robust (in C). But if all of the worlds are ones where S would not assert P in C, then P *is* ontologically robust (in C).

³ See, for example, David Braun [6], Nathan Salmon [41], and many others for works on the puzzles around negative existentials.

⁴ I say ‘worlds’ rather than ‘world’ because I assume the Lewisian model of closeness in which there are an infinite amount of ‘closest’ worlds, rather than the Stalnaker model with one single closest world. This seems like a more natural way to think of closeness, since often the change we want to imagine being made is one that could have many different ways of affecting the world that are essentially equivalent in closeness. I will not argue for this further, though I will freely point out that the Stalnaker view would considerably change the interpretation of ontological robustness to a stronger claim. As it is, we need only find one world among the closest where the speaker would continue on with her claim to show that it is not ontologically robust. But if there is only one closest world, it matters considerably what the speaker would do in that world, so we better be able to determine what happens there.

More must be said about what exactly the Fs are that we should take into account in our analysis. A brief example should illustrate this well. Suppose we are making plans for a hiking trip, and I say:

- (1) There are young pine trees growing in the north woods that I want to go see.

Which Fs should we consider to be relevant when looking at the closest worlds where I believe that Fs don't exist? Trees? Pine trees? Young pine trees? Young pine trees growing in the north woods? The last option is the most appropriate, as well as the most useful. Looking to the closest worlds where I believe that trees do not exist is overkill for the inquiry at hand. Also notice that if we were to symbolize (1) as a classic Aristotelian I-type sentence, we would do so as:

- (2) $(\exists x)[(Yx \ \& \ Px \ \& \ Gxn) \ \& \ (Wsx)]$

When we teach such symbolizations in an Intro to Logic course, we tell the students that these sentences are existentially quantified conjunctions. Inside the scope of the existential quantifier, the left conjunct is the group of things that the speaker is talking about, while the right conjunct is whatever the speaker is attributing to that group. It takes some practice to be able to symbolize such sentences properly, but it is not terribly difficult once the instruction sinks in. In sentence (1), I'm clearly talking *about* young pine trees growing in the north woods, and I'm saying about those particular trees that I want to go see them.

The set defined by the left conjunct may be called the *witness set* to the existentially quantified claim.⁵ That is, the witnesses to sentence (1) are young pine trees growing in the north woods, not all trees, or even all young pine trees. The group of objects we're talking about, or the witness set, need not be restricted explicitly. We have a mechanism of quantifier domain restriction that may restrict this set through other contextual mechanisms, such as shared intentions or shared knowledge in the context.⁶

The examples I discuss here will not rely on any particular views of quantifier domain

⁵ The term 'witness set' is defined formally in Barwise and Cooper [4], p. 191. My use of it here will be informal, though it should fit the use for which Barwise and Cooper introduced it.

⁶ I am happy to follow Stanley and Szabo's [44] explanation of how this mechanism functions. The issue of quantifier domain restriction will return in Chapter 3, where I'll have much more to say about it.

restriction—suffice it to say that we will want to be sensitive to whatever restrictions may be in play to fix our witness set for an investigation into the ontological robustness status of a claim.

1.2.2 Examples of ‘There is’ Claims

In this and the following few sections, I will discuss several examples that will help develop the notion of ontological robustness. I’ll begin with some examples of existentially quantified claims that make use of quantifier expressions with ‘is’ (and its cognates), then I’ll discuss some with ‘exists’, and finally, I’ll give some examples of arguments that employ both expressions. I will do some amount of analysis of these examples as I go, discussing why they would or would not be examples of ontologically robust claims (applying my definition from above), and then I will summarize and comment on the examples in more detail afterwards.

Superheroes

Consider first a conversation about superheroes, similar to that which motivated the discussion above. Suppose Nick is talking with his young daughter Mirah. Nick has begun the process of introducing Mirah to comic books and superheroes, but Mirah’s knowledge is still pretty limited. She’s also at an age where she is still learning the differences between make-believe and reality. Mirah’s favorite superhero is Spider-Man, and she’s familiar with some other big ones. Mirah is thinking about how Spider-Man is part human, part spider, and she asks Nick, “Daddy, are there any superheroes who are made of other animals than spiders?” Nick thinks for a minute and says, “Well, let’s see, some are basically just people with animal names, like Batman and Ant-Man. But there are superheroes who are animals but also people. Remember the Teenage Mutant Ninja Turtles?” Now suppose Mirah gets excited and says, “When I get older I want to be part turtle too, just like Leonardo!”

Nick could let this go, but if he wants to help his daughter get more comfortable with the divide between make-believe and reality, he might say, “Well, the Teenage Mutant Ninja Turtles don’t exist—they’re just made up for the comic books. People can’t actually be part turtle.” And then, if he’s a good father, he’ll also say plenty

of encouraging things about how Mirah could still be just as strong and talented as Leonardo when she grows up.

For the sake of clarifying what I take to be the root of a philosophical puzzle, let's highlight a couple claims that Nick makes. He asserted both of the following:

- (3) There are superheroes who are animals but also people.⁷
- (4) The Teenage Mutant Ninja Turtles don't exist—they're just made up for the comic books.

Notice that Nick easily quantifies over superheroes using 'there are' in (3), where the Teenage Mutant Ninja Turtles are obviously intended to be among the things he's quantifying over. But then he points out soon after, in (4), that they do not exist. Clearly, we could not freely interchange the existential quantifier expressions that Nick uses—since Nick doesn't believe that superheroes exist, he would not have said (3) with 'exist' instead of 'are'.

Now let's apply the definition of ontological robustness to Nick's claim of (3). First, if we plug in the variables, we get:

Nick's assertion of (3) is *ontologically robust* when talking to Mirah about superheroes⁸ iff Nick would not sincerely assert (3) when talking to Mirah about superheroes if Nick believed that superheroes did not exist.

The story itself shows that Nick *would* sincerely assert (3) if he believed that superheroes did not exist. To determine this, we look to the closest worlds where Nick believes that superheroes don't exist and ask whether he would assert (3) in those worlds. But the example given shows the actual world to be just such a world, as evidenced by his assertion of (4). In the world of this example—which seems entirely realistic—Nick is willing to sincerely assert (3) even though he believes that superheroes do not exist. So the right-hand side of our ontological robustness definition is false, which means the left-hand side is as well. And this means that Nick's claim is *not* ontologically robust.

⁷ From here on out, I will give in a footnote what I take to be the logical form of each existentially quantified claim that I number in these examples, clarifying what I take to be the structure—and thus also the witness set—of the claim. This sentence would end up as $(\exists x)[Sx \ \& \ (Ax \ \& \ Px)]$.

⁸ I simplified the context a little, but the phrase 'when talking to Mirah about superheroes' is meant to summarize that particular kind of discussion that they were having above.

One important point that we may draw from the case of Nick and Mirah, which shows that our potential class of examples of non-ontologically robust ‘there is’ claims is quite large, is that we already use ‘there is’ as a device for quantifying over things that we believe do not exist. That is, for many things, especially objects like fictional and mythical characters, the actual world is the closest world where we believe that the objects do not exist, and yet we still sincerely quantify over those objects using ‘there is’. (We do not, however, use the expression ‘there exists’ in the same scenarios. More on this shortly.) These are the limiting cases of non-ontologically robust claims, since the actual world is the closest world-of-interest for us.

Fictional Detectives

Another example that is similar to the superhero one, with a small but significant difference, is the following. Suppose Nora and Glenn are discussing detective novels. Glenn is new to the genre, and Nora is a long-time lover of these books. Glenn has just read a bunch of *Sherlock Holmes* stories, and he’s quite enamored. Nora tells him that those are a good start, but that there are tons of other kinds of detectives than just Holmes. He asks her, “Are there any other detectives who are as good as Holmes, though?” Nora responds, “Yes, there *are* other detectives who are as good as Holmes—Hercule Poirot is actually a better detective, in my opinion.”

Nora’s existentially quantified claim, the first part of her assertion, is the following:

- (5) There are other detectives who are as good as Holmes.⁹

To determine if this assertion is ontologically robust, we need to examine the closest worlds in which the witnesses to this quantified claim don’t exist and ask whether Nora would assert (5) in those worlds. This is a bit trickier than the superhero example, in which the witness set for (3) was clearly superheroes. Here we do not want to simply take the witnesses to be *all* detectives. This would point us to the closest worlds in which detectives don’t exist at all, which would be orthogonal to our interests. Looking at those worlds would be overkill in determining what to say about Nora’s claim, throwing the baby out with the bathwater.

⁹ Initially, we would translate this as $(\exists x)[(Dx \ \& \ x \neq h) \ \& \ Gxh]$. See below for a modification that will, I hope, seem reasonable.

Instead, we have plenty of cues to determine the witness set intended by Nora, which is narrower and more useful than the set of all detectives. In this context, Nora and Glenn are discussing detective *novels*, so they're clearly talking about fiction. Nora would take Glenn's question to be about *fictional* detectives, and she would be intending her response to match. So the witnesses here would just be fictional detectives, with the real claim under consideration being:

- (6) There are other fictional detectives who are as good as Holmes.¹⁰

The worlds we should consider here are just the closest worlds in which Nora believes that fictional detectives do not exist. Would Nora still assert (5) in these worlds?

The closest worlds where Nora believes that fictional detectives do not exist would very likely be the actual world, or maybe some other worlds extremely similar to our own. That depends a bit on Nora and what she thinks about fictional detectives. Many people, if asked, may hold the belief that fictional characters exist. (Not Nick in our previous example, though.) That belief may be mistaken,¹¹ but it's still one that people hold. So if Nora is one of those people, we'll have to find the closest worlds in which she has the opposite belief. It likely would not be very hard to change her mind, though,¹² so these worlds are not very far away. In such worlds, Nora would (justifiably) see the philosophical issue of the existence of fictional detectives as irrelevant to her discussion with Glenn. In these worlds, Nora and Glenn still read detective novels, and they still care about things like whether Holmes or Poirot is the better detective.

Since Nora would still sincerely assert (5) in the closest worlds where she believes that fictional detectives don't exist, (5) is *not* ontologically robust. What this means, roughly and intuitively, is that Nora's claim did not have any ontological weight to it. Or we might also say that any ontological weight that it appears to have is weight that would dissolve on closer inspection, if such a thing mattered. The information she's trying to convey and the conversation she's trying to have do not depend at all on the existence of the witnesses to her quantifier. And in spite of that, it still seems like her claim is sincere as well as appropriate in this conversation.

¹⁰ In this case, the translation would be $(\exists x)[(Dx \ \& \ x \neq h \ \& \ Fx) \ \& \ Gxh]$.

¹¹ As I think it is. See my [33] where I argue for an anti-realist view of fictional characters.

¹² She could read my paper, cited in the previous footnote, for example.

Tables

Let's now consider an example that goes further beyond the limiting case where the actual world is also the closest world we need to consider. Suppose that Martha is planning a trip to visit her friend Bergit, and Martha wants to bring with some craft projects she's working on. She's talking to Bergit on the phone about whether there will be enough space for her to spread out and work on her crafts, and Bergit says:

(7) There is a table that will work just fine.¹³

This claim has both an ontologically robust and a non-ontologically robust reading, and I'll discuss both briefly. The witnesses for this claim are *tables*, and so we imagine the closest worlds to the actual world in which Bergit believes that tables do not exist. What do such worlds look like? This is now an exercise in counterfactual reasoning. I'll discuss two different kinds of worlds that might be appropriate, and why they give us different results with respect to ontological robustness.

One kind of world we might consider is a world where Bergit believes that tables don't exist because there aren't even any table-shaped objects. The closest such worlds would likely be ones where Bergit believes that all the tables in the world have been destroyed.¹⁴ Bergit believes that all kinds of other pieces of furniture exist, but tables are just not part of the furniture of the world (quite literally). Let's call these the *deprived* worlds, since they are completely deprived of the functional usefulness that comes with having tables.

A different kind of world could be a world that, from the perspective of functional usefulness, is exactly like our own, but there is no such object as a table. Describing this world requires a bit of familiarity with debates in ontology, especially in mereology. We may see a table as being the composition of all the simples that are arranged tablewise. Someone who is a nihilist about composition, for example, will argue that there is no composite object that constitutes a table over-and-above the simples arranged tablewise.¹⁵ We need not decide this particular debate, or even advance it, here. The

¹³ This would be symbolized simply as $(\exists x)(Tx \ \& \ Wx)$.

¹⁴ There must have been some tables to begin with, since Bergit has beliefs about them in her world. Thanks to Brent Braga for pointing this out to me.

¹⁵ This strange kind of vocabulary can be found in plenty of discussions about mereology, such as Peter van Inwagen [46] saying, "Of course, if our proposed answer to the Special Composition Question

fact that there is such an ontological question about tables is enough to motivate our consideration of a world where Bergit believes that no such composite objects exist. Let's call these the *non-composition* worlds.

The closeness of worlds is notoriously difficult to determine with a great amount of confidence or stability.¹⁶ Whether the *deprived* or the *non-composition* worlds are closer to the actual world depends a bit on Bergit and her inclinations. If Bergit has had no exposure to philosophy and is not inclined to be gripped by metaphysical puzzles, the *deprived* worlds seem closer. But if she is a bit more open to abstract philosophical worries, then the *non-composition* worlds seem closer.

In the *deprived* case, where Bergit is not philosophically inclined, we should say that the assertion of (7) *is* ontologically robust, since if Bergit believed that all the tables had been destroyed, for example, she wouldn't say that there is one—this would obviously contradict her belief that they'd all been destroyed. In the *non-composition* case, however, we should say that Bergit's assertion of (7) *is not* ontologically robust. If the reason for Bergit thinking that tables do not exist is rooted just in philosophical worries about mereological composition, we could reasonably expect her and others to still go around talking about tables, even sincerely, since it is such a useful way to talk. It may be a claim that she would take back if pressed about her ontological commitments, but that does not mean she's insincere when she says (7) to her friend Martha.

Free Tables

Now let's consider an example of a claim that *is* ontologically robust. Suppose a restaurant host says to a new customer:

(8) There is a free table near the window.¹⁷

In this case, it is reasonable to see the witnesses to this claim as being *free tables*, so we need only consider the closest worlds in which the host believes that free tables do not exist. These worlds would be ones where the host believes the restaurant to be full, and

is right, then it is doubtless true that, for any *x*s, if those *x*s are arranged tablewise or chairwise, then those *x*s are all simples," p.109.

¹⁶ See, for a classic example, Quine's 'Caesar' counterfactuals in Quine [38], p. 222.

¹⁷ A standard symbolization of (8) would be $(\exists x)[(Tx \ \& \ Fx) \ \& \ (N_{xw})]$

so he would not sincerely assert (8) to the newcomers. So (8) is clearly ontologically robust.

Wolves in the Park

Another example of an ontologically robust claim comes with the following example. Suppose Olga is thinking of going for a hike in Scenic State Park in northern Minnesota and Ralph, the park ranger, says to her:

(9) There are wolves in the park, and they're a little dangerous.¹⁸

This is clearly an existentially quantified claim, and the witness set for it is *wolves in Scenic State Park*. Now consider Ralph and Olga in the same context, where Olga is getting ready to go for a hike in Scenic State Park but, in this case, Ralph the ranger believes that wolves do not exist in Scenic State Park. This would not be a world where no wolves exist—a much closer group of worlds would be those where Scenic State Park just isn't a place where wolves go, and Ralph is aware of this. In these worlds, would Ralph sincerely assert (9) to Olga? Certainly not. He would need to be insincere if he were to say it.¹⁹ This is the kind of case where the existence of the witnesses really is important to whether a speaker will sincerely utter the claim, since the existence of the wolves is the whole point of the warning.

1.2.3 Examples of 'There Exist' Claims

So far we have only discussed examples that make use of existential quantifier expressions with 'is' and its cognates. Let's now take a look at some examples with 'there exists' and 'exists'. The usage of these expressions outside of philosophical contexts is infrequent, but there are certain kinds of settings where we tend to find them. In some scientific settings, for example, we may use 'there exists' somewhat naturally when describing the world.

¹⁸ We would translate this claim as $(\exists x)[(Wx \ \& \ Ixs) \ \& \ Dx]$.

¹⁹ We could imagine cases where he might still say (9) to her—if, say, some other dangerous animals were around the park, but he knows that only the threat of wolves would keep Olga from hiking. But that would be an insincere assertion of it.

Eyeless Fish

Consider Jane, a biology teacher, telling her class:

(10) There exists an eyeless fish that lives in deep waters.²⁰

Some science teachers do speak this way, and we understand them when they do; and there's nothing grammatically wrong with the claim. Is Jane's claim ontologically robust? Yes, it certainly is. The witnesses to this claim will be eyeless fish, and the closest worlds where Jane believes that eyeless fish don't exist will likely be worlds where eyeless fish either don't exist or they just haven't been discovered yet. Or they could be worlds where Jane is just a bad biology teacher and has mistaken beliefs about the world. In any of these cases, it seems clear that Jane would not sincerely assert (10).

With this example, the setting of the biology class would also make the 'there is' version of this example ontologically robust. Suppose Jane said instead:

(11) There are eyeless fish that live in deep waters.

We would have the same witness set here and the same context, so the closest worlds would be the same. She wouldn't sincerely assert (11) any sooner than she would (10) in these worlds.

One thing to notice about these two variations of the eyeless fish example is that if we're explaining *why* either of the claims (10) and (11) are ontologically robust, we will naturally give slightly different explanations. With the first one, we might say that Jane wouldn't say (10) simply because she believes they don't exist—so obviously she wouldn't say that they exist. But with the second one, we're more likely to appeal to the fact that she's teaching a biology class as the primary explanation, since this clarifies that the existence of the fish actually *matters* to what she's saying.

In this case, it doesn't even matter if the closest worlds where Jane believes that fish don't exist is one where she is a nihilist about mereology. If this is the case, she would not use the word 'exist' when talking about fish—it would be too obviously contrary to her beliefs. In this world, however, she likely would be comfortable uttering (11) in many settings, though probably not in the middle of a discussion about mereology.

²⁰ We would translate (10) as $(\exists x)[(\sim Ex \ \& \ Fx) \ \& \ Cx]$, where 'Ex' is 'x has eyes', 'Fx' is 'x is a fish', and 'Dx' is 'x lives in deep waters'.

Small Cottage

An example from a more ordinary setting is a bit more difficult to manufacture and still have it sound like something anyone would actually say. We do use the word ‘exist’, but when it appears outside of the phrase ‘there exists’, it becomes a bit contentious whether it should be read as a quantifier expression or as a predicate. In order to avoid such contention, I’ll attempt to discuss only examples that are obviously quantified claims.

Suppose Caroline and Anthony are talking about places they would go on vacation or to live if they won the lottery and didn’t need to worry about money anymore. Suppose also that they enjoy using somewhat dramatic and poetic language. In this setting, Caroline says the following:

- (12) There exists a small cottage on an island off the North Shore of Lake Superior that I would want to live in for a while.²¹

As with the fish case above, we should say that if Caroline believes that cottages do not exist on islands off the North Shore (which, under a reasonable reading, are the witnesses to the claim), then she would not sincerely assert (12). This claim is clearly ontologically robust, since the existence of the right kinds of cottages really matters to her claim in this context.

Also, as with the fish case, if we altered the sentence to a ‘there is’ claim and had it uttered by Caroline in the same context, it would be reasonable to see that as being ontologically robust as well.

1.2.4 Examples About Validity

I’ll now discuss a few examples that involve our intuitions about the validity of some simple arguments. As David Kaplan maintains, sometimes our intuitions about the truth of a sentence may be less stable than our intuitions about the logical consequence relation holding between some claims.²²

²¹ We would translate (12) as $(\exists x)[(Sx \ \& \ Cx \ \& \ Nx) \ \& \ Wcx]$, where ‘Sx’ is ‘x is small’, ‘Cx’ is ‘x is a cottage’, ‘Nx’ is ‘x is on an island off the North Shore of Lake Superior’, and ‘Wxy’ is ‘x would want to live in y’.

²² In Kaplan [24], he says: “The *important* point is that although we may have differing, even shaky, intuitions about truth, we—or at least, I—have more stable intuitions about logical consequence. These have been ignored because of the nearly universal, and according to me, fallacious, assumption that the notion of logical consequence is *derivative* from the more secure notion of truth,” p.11

To see the contrast between ‘is’ and ‘exist’ in a slightly different way, consider the following arguments:

Argument 1

P₁ There exist superheroes who can fly.

C₁ Superheroes exist.

Argument 2

P₂ There are superheroes who can fly.

C₂ Superheroes exist.

The only difference between these two arguments comes in the premises, where the first uses ‘there exist’ and the second ‘there are’ to existentially quantify over superheroes. *Argument 1* seems trivially valid—it sounds as simple as an example of reasoning from a conjunction to one of its conjuncts. *Argument 2*, on the other hand, strikes us as invalid, since counterexamples immediately come to mind.

This may sound tricky, but I think that with *Argument 2*, the kind of setting we’re in will make a difference to our intuitions about its validity, just as the context affects our intuitions about individual claims. If we’re having a conversation in a comic book store, or with our children, we would not be willing to allow that **C**₂ follows from **P**₂. If we’re having a philosophical discussion, though, our intuitions would be a bit different. In that setting, we wouldn’t be willing to assent to **P**₂ unless we’re also willing to assent to **C**₂. In any case, it seems clear that it would be a mistake to say flat-out that *Argument 2* is valid, since non-philosophical contexts are legitimate places to have discussions and make inferences. But we don’t have any similar kind of concern with *Argument 1*.

The notion of ontological robustness can help clarify the difference between these arguments, and thus also the difference between ‘is’ and ‘exists’. Ontological robustness must (by definition) be evaluated with respect to a context of utterance, but we should have some idea of what we would think of these arguments in various contexts. If we just focus on the premises, we see that **P**₁ will be ontologically robust regardless of its context of utterance. That is, anytime someone asserts **P**₁, it will be true that in the closest worlds in which they believe that superheroes don’t exist, they would not sincerely assert **P**₁. Or, to put this another way, if someone believed that superheroes didn’t exist and they wanted to assert **P**₁, they would have to do so insincerely.

On the other hand, \mathbf{P}_2 may be asserted sincerely sometimes, even when the speaker believes that superheroes don't exist. It may be completely beside the point of the conversation whether superheroes exist or not—the speaker may just be discussing the different kinds of super powers that various comic book characters have, completely aware that she is discussing fiction. If the conversation turns to a discussion of the ontology of fictional characters, however, then the speaker cannot still sincerely assert \mathbf{P}_2 if she still believes that superheroes don't exist. Their existence now matters to the discussion.

The connection between the ontological robustness of these premises and our intuitions about the validity of the arguments is that they are both linked to what we would be willing to assent to in a variety of worlds where we have different beliefs about the existence of superheroes. With *Argument 1*, since we could not sincerely assert \mathbf{P}_1 while believing that superheroes do not exist, the conclusion \mathbf{C}_1 is a claim that we would accept in any scenario where we would accept \mathbf{P}_1 . With *Argument 2*, the natural contexts in which \mathbf{P}_2 comes out non-ontologically robust (such as simply discussing super powers) will track those scenarios that shape our intuitions that this argument is invalid, since we will vividly consider what we take to be solid counterexamples (like Superman, Iron Man, etc.). In these cases, we wouldn't accept that we could rationally move from \mathbf{P}_2 to \mathbf{C}_2 .

This difference in our intuitions about the validity between the two arguments, connected to the difference in ontological robustness of the premises, leads us to the difference between 'is' and 'exists'. We find in these cases (as we would in many others) that the context is crucial for determining whether 'there is' claims are ontologically robust or not, but it is not crucial for determining whether 'there exist' claims are ontologically robust. The importance of context is seen in the examples in the previous section with evaluating individual claims, and it is also seen in the way we evaluate the validity of these arguments. Since 'there is' claims will be ontologically robust in philosophical settings but not ordinary settings, we will not accept that the corresponding existence claims follow from them, the way that the existence claim would appear to follow from the 'there exists' claim.

Before moving on, I'll consider a couple more pairs of arguments we can use to compare our intuitions about validity and how they relate to the ontological robustness

status of some of the claims involved. Not every pair of arguments will draw out the difference that I think is there, but my aim is just to show that the difference between ‘is’ and ‘exists’ is substantial enough that we can find it in a variety of places.

First, here is a pair of arguments where an existentially quantified claim is supposed to follow from a different existentially quantified claim, where all that changes is the existential quantifier expression being used:

Argument 3

P₃ There exist superheroes who can fly.

C₃ There are superheroes who can fly.

Argument 4

P₄ There are superheroes who can fly.

C₄ There exist superheroes who can fly.

With these examples, the analysis will be quite similar to the previous pair. *Argument 3* strikes us immediately as a valid argument, while our intuitions about *Argument 4* will depend much more on the kind of conversation we’re having. In a non-philosophical setting, *Argument 4* will very likely strike us as invalid, for the same reasons discussed above.

Next, consider a pair of arguments that follow the standard pattern of an ‘existential introduction’ rule. In both arguments, we begin with a premise about a few superheroes who can fly, and the conclusion drawn from it is an existentially quantified claim, differing only in which existential quantifier expression is used:

Argument 5

P₅ Superman and Thor are superheroes who can fly.

C₅ There are superheroes who can fly.

Argument 6

P₆ Superman and Thor are superheroes who can fly.

C₆ There exist superheroes who can fly.

As with the other pairs, the first of this pair strikes us immediately as a valid argument, while the second one strikes us at least sometimes as invalid. We cannot exactly evaluate the premises for their ontological robustness status, since that notion is defined as a

property had (or not had) by positive existentially quantified claims. With the conclusions, the analysis I've given of these claims is that C_5 is sometimes ontologically robust, sometimes not, depending on the context, while C_6 is always ontologically robust. This helps explain our intuitions about the validity of these arguments, since the flexibility of C_5 explains why the first argument sounds valid—we may assent to P_5 because we're just talking about comic books, or we may actually think that superheroes exist, but either way, we'd be happy assenting to C_5 on whatever basis we have for assenting to P_5 . But in the situations where we would assent to P_6 while also not believing that they exist, we would not assent to C_6 , which explains why we likely feel that *Argument 6* is not valid.

1.2.5 What the Examples Tell Us

The examples discussed in this section should make it clear that the notion of ontological robustness, though stipulated here, is a useful notion that tells us something interesting and significant about existentially quantified claims. Analyzing an assertion for its ontological robustness status in a context gives us a fruitful way to discuss our ontological commitments in that context.

A method for uncovering our ontological commitments naturally brings to mind Quine's criterion of ontological commitment, but applying the notion of ontological robustness to a case is considerably different. Quine's criterion says that we are committed to those objects that we would willingly quantify over if we were pressed to be serious and honest about what we think exists.²³ But these two processes involve different kinds of philosophical discussions. Hashing out our ontological commitments in Quine's sense is a metaphysical project, whereas determining the ontological robustness status of an assertion is a project in semantics or the philosophy of language, something we engage

²³ Peter van Inwagen [47] describes the process behind Quine's criterion as follows: "The strategy is this: one takes sentences that the other party to the conversation accepts, and by whatever dialectical devices one can muster, one gets him to introduce more and more quantifiers and variables into those sentences. (Or, if you will, one gets him to accept new sentences, sentences that come from the sentences he initially endorsed by the progressive replacement of devices and constructions belonging to ordinary English by devices and constructions belonging to the canonical language of quantification. [...]) If, at a certain point in this procedure, it emerges that the existential generalization on a certain open sentence F can be formally deduced from the sentences he accepts, one has shown that the sentences that he accepts, and the ways of introducing quantifiers and variables into those sentences that he has endorsed, formally commit him to there being things that satisfy F," pp. 246–7

in when our interest is in determining the meaning of our assertions in the contexts in which we asserted them.

To contrast these two projects, we can look at cases where we routinely quantify over objects that we ourselves believe not to exist. Quine's claim is that when we're pressed about these claims, if we really don't believe they exist, we'll retract our earlier assertion of the quantified claim. We'll say that we must have just been talking loosely, that we just weren't paying close attention to the status of these objects. Quine is no revisionist, though—he has no problem with us making assertions in our ordinary discourse that quantify over nonexistent objects, since those claims can accomplish other tasks than just to state the existence of certain objects. He only would have a problem with it if, when pressed, we did not go some way towards retracting our earlier claim, while still professing to believe that the objects don't exist.

My project, on the other hand, only disagrees with part of Quine's picture. Quine seems to imply that we're not sincere when we quantify over objects that we believe not to exist, but I think that sometimes we are. I can accept the rest of Quine's picture, even the part about us retracting our earlier claims when we're pressed to think about what we actually believe to exist. I just think that the fact that we retract our earlier claims while we're talking metaphysics doesn't show that the earlier claims were loose or insincere. My interest is not in what we would commit ourselves to when pressed; rather, I'm interested in the meaning of our existential quantifier expressions in a variety of contexts (i.e., not just in the context of discussing ontology).

Pointing Towards a Contextualist View

Before going on to discuss them in more detail, I'll first give a quick summary of the examples given above. Section 1.2.2 includes two 'is' claims (about superheroes and detectives) that clearly are not ontologically robust, one (about tables) that could go either way, depending on some attributes of the speaker, and two 'is' claims (about free tables and wolves) that *are* ontologically robust. The examples with 'exist', given in section 1.2.3, are both clearly cases of ontologically robust claims. And finally, section 1.2.4 discusses a few examples of arguments and the effect that the quantifier expression chosen may sometimes have on our intuitions about the validity of an argument. Included in that discussion is some analysis of how the notion of ontological robustness

can help us explain that difference.

What is shown by these examples? One lesson we can clearly glean from them is that claims with ‘is’ in the quantifier expression sometimes are ontologically robust, and sometimes aren’t. Some of the examples of ‘there is’ claims are clearly ontologically robust, and some are not. Furthermore, a single ‘there is’ claim uttered in one context may change its ontological robustness status if uttered in a different context, such as (7). We can shift that status just by changing features of the context, such as the philosophical openness of the speaker, or the setting in which it’s uttered.

When we turn to the word ‘exist’, however, we get a different story. All of the examples included above that involved claims with ‘exist’ in them (i.e., in 1.2.3 and 1.2.4) were clearly ontologically robust. And if we were to go back into the examples of ‘is’ claims in 1.2.2 and change the quantifier expressions to ‘exist’ ones, I believe we would find them all to be ontologically robust regardless of the context. Those claims would then be:

- (13) There exist superheroes who are animals but also people.
- (14) There exist other fictional detectives who are as good as Holmes.
- (15) There exists a table that will work just fine.
- (16) There exists a free table near the window.
- (17) There exist wolves in the park, and they’re a little dangerous

The last two of these were found to be ontologically robust even in their ‘is’ forms above. With the other three, the ‘is’ versions had non-ontologically robust readings in addition to potential ontologically robust readings. It’s possible that this is my own personal semantic failing, but I can’t find a scenario that would allow any of the sentences (13) through (17) to be non-ontologically robust. We would have to imagine a speaker who believes that the superheroes or wolves do not exist, but then goes on to quantify over them using the expression ‘there exist’ anyway. The problem is that the word ‘exist’ just appears to have more power than ‘is’ when it comes to fixing the ontological robustness status of a claim.

I will address *why* it is so next, but here I just want to point out *that* it appears to be the case that context is crucial for determining whether ‘there is’ claims are

ontologically robust or not, but it is not crucial for determining whether ‘there exist’ claims are ontologically robust. The former will frequently be ontologically robust in philosophical settings, but not in ordinary settings. And the latter, the claims with ‘exist’, will routinely be ontologically robust regardless of context.

Why We Should Expect These Differences

There is a rather obvious and intuitive reason why examples with ‘exists’ should be routinely ontologically robust. The expression ‘there exists’ and its cognates that actually make use of the word ‘exist’ are strongly and immediately linked to the notion of existence, so we will not use these expressions to quantify over objects that we believe do not exist. It sounds immediately contrary or contradictory to us.

We have constraints on the kinds of claims where it is appropriate to use ‘exists’ to existentially quantify—constraints that do not apply to ‘there is’. These constraints are such that, even in ordinary settings, it is only appropriate to use ‘exist’ when there is actually some ontological, or quasi-ontological, issue at hand. For example, we might naturally say, “It turns out that life exists on Mars,” since there’s actually an issue of existence at hand. Now, this issue of existence is a scientific one, not a philosophical one. But the focus of this claim is on the existence of some object.²⁴

The immediate link between expressions that employ ‘exists’ and the notion of existence prevents us from using these expressions in ordinary contexts where we do not care about the existence of the objects—especially ones where we want to make claims that at least appear to refer to objects that we ourselves believe to be non-existent. This happens all the time, and it turns out that we just don’t end up using the word ‘exists’ all that much in ordinary conversation. We *do* need some device in our language, however, to communicate our quantified thoughts about nonexistent objects. The expression ‘there is’ works as just such a device, since it does not strike us as immediately contrary or contradictory to quantify over known nonexistent objects with it.

²⁴ It would be a mistake to think that there is a sharp line dividing existence questions that are to be answered by scientists and those that are to be answered by philosophers. The different disciplines provide us with different methods for approaching and attempting to answer existence questions. When the objects in question are ones that can possibly be discovered by scientific means, then we use those means to investigate, but when they are not so discoverable, we employ other means. And sometimes the same objects will be investigated through both disciplines, usually to satisfy very different interests (e.g., questions about the existence of tables), but not always (e.g., questions about atoms and gunk).

This phenomenon is no accident—it makes sense that we would have a preference for communicating with more flexible expressions that don't always pack in the notion of existence, when we're so frequently using existentially quantified claims to communicate something other than the existence of some objects. The device of 'there is' allows us to make the claims we want to make without worrying about existence—or puzzles about nonexistent objects—at all.

1.3 Looking Ahead

I aimed to establish in this chapter that there is a significant difference between 'is' and 'exists' as they appear in existential quantifier expressions. The difference is that existentially quantified claims that employ existential quantifier expressions with 'exists' are always ontologically robust, while those that employ existential quantifier expressions with 'is' sometimes are, and sometimes are not, ontologically robust.

The orthodox view expressed by Lewis must be inadequate, since it cannot be the case that all existential quantifier expressions are 'entirely synonymous and interchangeable', as he maintains. I have only discussed 'is' and 'exists', but as we see, these two expressions differ significantly in their meaning. The device of ontological robustness helps clarify the difference, giving us a method for teasing out the ontological thickness of positive existential claims. When we analyze a variety of examples, we see that 'there exist' claims are routinely ontologically robust, while 'there is' claims may or may not be robust, depending on the context in which they're uttered. It should be clear that this difference provides us with a significant reason to reject the orthodox view.

If we reject the orthodox view, though, what is the alternative? Proponents of orthodoxy may maintain that the only alternative is a Meinongian view that commits us to non-existent objects as part of our ontology, but I reject the notion that Meinongianism is the only alternative. I will offer an account of the meaning of existential quantifier expressions, beginning in the next chapter, that respects the difference between 'is' and 'exists' without requiring us to add non-existent objects to our ontology, or to make any robust distinctions between different kinds of being.

The anti-orthodox view that I suggest replacing the orthodoxy with may not be surprising at this point, given how I've discussed this issue. I maintain that 'is', as it

appears in existential quantifier expressions, is context-sensitive, while ‘exists’ is not. A context-sensitive view of ‘is’ explains the variance across contexts of the ontological robustness status of ‘there is’ claims, and a view of ‘exists’ that treats it as invariant explains the lack of change of the ontological robustness status of ‘exists’ claims in different settings.

Two expressions cannot be synonymous if one of them is context-sensitive and the other is not—this is a significant semantic difference. This difference with respect to context-sensitivity also prevents these expressions from being entirely interchangeable. So my proposed view amounts to a denial of both parts of Lewis’s description of the orthodox view—I claim that ‘is’ and ‘exists’ are neither interchangeable nor synonymous. Next, I’ll begin filling out this anti-orthodox view by developing my context-sensitive view of existential quantifier expressions with ‘is’.

Chapter 2

Contextualism about ‘There Is’

“This difference in meaning is admittedly somewhat subtle, but it should be the goal of semantic analysis to elucidate subtle differences, not ride roughshod over them.”

-Edward L. Keenan, “Quantifier Structures in English”

2.1 Introduction

The view I present here is that the existential quantifier expression ‘there is’ is a context-sensitive expression. I’ll call this *contextualism about ‘there is’*. What does it mean to say that ‘there is’ is a context-sensitive expression? Just as with demonstratives and indexicals, it means that the content of the expression, when uttered in a context, is sensitive to some feature of that context. If that feature of the context changes, then so does the content.

With indexicals like ‘I’ and ‘today’, the feature of the context is found in the character of the expression. For example, the character of ‘I’ could be described as ‘the speaker of the context’. This character is a function from the context of the utterance to the speaker in that context. With demonstratives like ‘this’ or ‘that’, the feature of the context to which the expression is sensitive is some kind of a demonstration in that context, either physical or intentional. When I say ‘that glass’ and point to a glass, my pointing (whether the physical pointing or the mental pointing, or intention, that goes along with it) appears to be doing the work that the character of the indexical did.

My view is that ‘there is’ has some similarity with demonstratives, picking up on the intentions of the speaker in a context. Specifically, what changes from context to context is the *intended domain* of the speaker, which is essentially the domain that we would consider part of the interpretation that determines the truth or falsity of a statement.

In this chapter, I’ll present my contextualist view, beginning with an overview and moving on to some of the more gritty semantic details. Although there will be some arguments here and there, the primary focus of this chapter will be to present the view and develop enough of the details to be able to say how this context-sensitivity works and where it’s located. The previous chapter presented a preliminary argument for this view, based on the analysis of the ontological robustness status of several examples. The argument will continue in subsequent chapters, primarily through holistic comparison of this view against its competitors.

Before I begin the overview, I’ll give a brief caveat about the discussion that follows: though it will appear from some of the things I say that I am assuming a Meinongian position about objects—that in a substantive sense there are nonexistent objects—I do not in fact hold this position. My view is about context-sensitivity, so the context of a discussion in which the relevant expressions get used matters. And in the context in which I’m writing, I am doing philosophy, thinking about existence, etc. And in this context, I would not seriously assert, “There are nonexistent objects,” nor would I think that claim is true. I think that the set of all the objects that there are is one and the same as the set of all the objects that exist.

That said, I will be making claims that involve things that I will refer to as, for example, ‘a domain that includes nonexistent objects’. Since I don’t think that there are nonexistent objects, I also don’t think that there are domains that *include* nonexistent objects. I do, however, think that in many settings and conversations, speakers (including me) assume or presuppose or help themselves to such domains. This is an error of those speakers, at least in some sense, but I believe it is a common part of linguistic practice. So in this section, when I speak about domains that include nonexistent objects, my claims should really be taken as being about *purported* domains that include nonexistent objects. Anything I say here that may appear to commit me to a Meinongian position, or any kind of ontology that includes nonexistent objects in any

substantive way, are merely a matter of my speaking a bit loosely in order to effectively explain what I think is happening with the meaning of existential quantifier expressions. In the next chapter, I will discuss in more detail the sorts of errors I think speakers make in this regard, and this will (I hope) help placate concerns that a non-Meinongian may have about my ontological commitments I seem to have in reading what follows here in this chapter.

2.2 Overview of Contextualism

In Chapter 1, I introduced a notion called *ontological robustness*, and I discussed several examples of existentially quantified claims to illustrate this notion. My conclusion was that existentially quantified claims that use ‘there is’ may be ontologically robust in one context but non-ontologically robust in a different context.¹ I take this as both motivation and support for the view that ‘there is’ is a context-sensitive expression. The specific contextualist view I present, then, must do its part to explain and illuminate the phenomena we see surrounding the notion of ontological robustness.

My claim is that the context-sensitivity of ‘there is’ lies in the mechanism that determines the domain with which the quantifier is associated. This mechanism involves the intention of the speaker, so that if the speaker intends to quantify over non-existent objects, then the quantifier will be associated with a domain that includes non-existent objects (if there is one). If the speaker intends to speak only about objects that exist, then the quantifier will be associated with a different domain, i.e., the domain of objects that exist.

A standard Kaplanian account of demonstratives holds that the speaker’s intentions play a role in fixing the content. The intentions are not part of the semantics of the expression, i.e., they are not part of the content. But they play a contextual, meta-semantic role in helping to fix what was said. They are not the same as a rule that appears as part of the character of an indexical, but they do play a similar role. If a speaker had different intentions accompanying her use of a demonstrative expression, the content of her utterance could be different.²

¹ I also believe that ‘there exist’ claims are invariably ontologically robust, and I will discuss that further below, at the end of this chapter.

² Kaplan [23] [22] are the classic works on this issue, though it’s in *Afterthoughts* specifically where

It is helpful to separate here the metaphysical/semantical issue from a related, but distinct, epistemological issue. The metaphysical/semantical part of this view of demonstratives says that intentions play a role in helping to *fix* the content of the utterance. A related epistemological issue involves how we might know the speaker's intentions, or perhaps also how we might know that a speaker has been successful in referring to something via particular intentions. In discussing my view, I will address both of these kinds of issues, but it is important not to conflate them into one, since the metaphysical/semantical part is actually part of my view of the meaning of 'there is', while the epistemological question is a secondary issue that comes up when worrying about how the view fits with our common knowledge and practice.

Now, revisiting the issue of ontological robustness, we noticed that the following sentence is sometimes ontologically robust, sometimes not:

- (1) There are superheroes who are animals but also people.

My claim was that when we make this claim in an ordinary setting, it seems clear that we intend to quantify over nonexistent objects. The example I gave was of Nick talking to his daughter Mirah, where he also said, nearly in the same breath, that superheroes do not exist. There is no contradiction or even any cognitive dissonance in thinking and asserting that superheroes don't exist while also quantifying over them, so long as we think it is coherent to intend to quantify over nonexistent objects. Nick does not need to be a Meinongian—he doesn't have to have any well-worked-out metaphysical views at all. He may just not care about the kinds of worries that philosophers bring up when we propose an ontological theory that includes nonexistent objects. This doesn't mean that he's *successful* in attempting to quantify over nonexistent objects. My claim is just that something about the meaning of 'there is' allows us to talk the way Nick talked in a sensible, coherent way.

If we imagine Nick as a philosopher, however, engrossed in a discussion about the ontology of fictional characters, then we see (1) in a different light. We can imagine situations where his utterance of (1) is intended to reflect his beliefs about what exists, where the domain of discourse in his conversation is intended to only include objects that exist.

he discusses the use of intentions that is important to my view.

To illustrate my contextualist view about ‘there is’, I’ll walk through what I take to be the difference between these two scenarios in which Nick utters (1). I’ve already mentioned that I think the key here lies in Nick’s intentions. It may be helpful, though, to walk through, in a metaphorical way, the path from utterance to content. I’ll discuss this path first in an informal way, and then a bit more formally with functional language.

When Nick utters (1), he does so in a context. Two contexts are described above, and we’ll call the first one *ordinary* and the second one *ontological*. It is natural to think of a conversation having a domain of discourse associated with it, or some set of objects intended to be used for the purpose of interpreting the statements made in the conversation. My claim is that this domain can be fixed, or altered, through the intentions of a speaker. A domain of discourse can be a dynamic set, though I do not think it is completely dynamic, constantly changing in small ways. It seems reasonable to think that there are some privileged sets that tend to get fixed onto more than others, such as the set of all objects that exist, and perhaps others.³

In walking through the path from utterance to content, we begin by looking at the linguistic meaning of (1). Nick is speaking English, and the words and phrases he utters all have meanings in English. The linguistic meaning may be the character of an expression, or it may just be something in the neighborhood of a character—some kind of information that a competent speaker of the language would know. So in this case, the linguistic meaning of the utterance would include some minimal information about what superheroes, animals, and people are. It would also include minimal information about the meaning and usage of the other words: ‘who’, ‘are’, ‘but also’, and, of course, ‘there are’. The kind of information in the linguistic meaning of ‘there are’ would be whatever a competent speaker needs to know about it to have a handle on it. So it would include some information about the words expressing what I like to call *at-least-oneness*, that whatever it operates over must have at least one representative somewhere. A competent speaker needn’t have any kind of sophisticated thoughts about sets, but the notion of a non-empty group is pretty basic, so it’s reasonable to say that the involvement of something like that notion is part of the linguistic meaning of ‘there is’. In addition to this information, my claim is that the broadness of the group in question

³ I’ll revisit this notion in the next chapter when discussing Ted Sider and his views about ‘reference magnetism’ and quantifier meanings.

is flexible, and we don't appear to want to limit it just to those objects that exist. This flexibility is the heart of the context-sensitivity, since it allows us to talk about different broad groups of objects in different settings.

From the linguistic meaning, we follow the path along to the semantic meaning, or the content. There was an utterance of (1), which was an utterance in a language, where the words have linguistic meaning, and so, through that meaning, the utterance expresses some kind of content that, since (1) is a declarative sentence, will be the sort of content that could be true or false, *viz.* a proposition. We can see the constituents of the proposition as corresponding to the different words or expressions used in the sentence, so that we have the word 'superheroes' leading to the property of being a superhero, and the same with 'animals' and 'people'. The contents of 'who' and 'but also' are essentially just whatever is needed to evaluate the proposition in the right way (since 'who' is basically functioning as a variable here, and 'but also' is just an expression of conjunction). When it comes to 'there are', the entity that shows up in the proposition is essentially a function—one that takes properties as its input and gives other functions as its output. I'll give more details about the functions shortly, but to continue with a more informal discussion, we can describe this entity as a *quantifier*, specifically as an *existential* quantifier. As a semantic entity, a quantifier has certain kinds of features associated with it, such as an inferential role and an attachment to a domain over which it ranges. It is in this way (as part of this informal explanation) that the domain of discourse finds itself attached to the semantics of the quantifier expression, or finds its way into the content. And since, at the level of linguistic meaning, the flexibility was in the different groups of objects we could be talking about, we will find that reflected in the content as the domain associated with the quantifier.

Now, to compare the two contexts in which Nick utters (1), what can we say is different? The general atmosphere is changed, since instead of talking with little Mirah, he's now talking with some philosophers. Most importantly, though, the *goal* of saying something like (1) changes for Nick between these two contexts as well. In the ordinary context, his goal is to educate Mirah and arrange her beliefs in a way that is best-suited for reality. This includes both reminding her that superheroes do not exist, but also giving her the information that certain superheroes do happen to be both animal and people. If she's going to know something about superheroes, she should have

accurate information, and arranging her beliefs in that way will give her an accurate picture of superheroes, regardless of whether they exist or not. In the ontology context, however, the goal is more ontologically robust. In the midst of that discussion, Nick could not coherently both maintain that superheroes do not exist and that there are superheroes who are both animal and people. The goal of that discussion is to present and defend views about what exists, and so the domain of discourse must be in line with conversational goals. The explicit concern with which objects exist would override Nick's intentions to quantify over nonexistent objects, which he would not have in this context anyway. So the goals and intentions of the speaker make a big difference, and that difference relates to the domain of objects associated with the quantifier.

2.3 Two Semantic Approaches to Quantifiers

The superhero example gives us an intuitive grip on the contextualist view about 'there is', but now I will widen the view a bit by discussing some semantic details of the view. The existential quantifier is one of the central logical notions in first-order predicate logic, and a good deal of work has been done developing the semantics of it in that setting. But in more recent work in linguistics, the existential quantifier tends to be seen as one determiner among many, especially due to the influence of work on generalized quantifiers by Barwise and Cooper [4]. These different perspectives give us a couple ways to approach quantifier semantics, and I'll discuss both here, applying them back to the superhero case as well.

2.3.1 Logical Treatment

I'll begin by discussing the kind of treatment we give of the existential quantifier in a logic course, a treatment rooted in Peirce, Frege, Russell, and others. The project here is to translate quantified sentences into their symbolic form, revealing the logical structure of the claims, and then evaluate them using model-theoretic truth conditions.

Let's begin with a simple example of an existentially quantified claim and use this to explain what needs to be explained about the logical treatment. Consider:

- (2) There are black swans.

This sentence is made up of a few different parts. First of all, we have ‘there is’, which is the existential quantifier expression. It’s a *quantifier* expression because it’s expressing a *quantity*. And it’s the *existential* quantifier expression because the quantity it expresses is the quantity of *at least one*.⁴ Next, we have the noun phrase ‘black swans’, with the adjective ‘black’ modifying the noun ‘swans’. The words ‘black’ and ‘swan’ both express properties that objects may or may not have.

When we symbolize a quantified claim like this, we make use of *variables* to connect the quantifier to the predicates in appropriate ways. We may unpack such a sentence, before symbolizing it, by putting in variables to display the connections. When we do this, we also reveal the form of a sentence such as (2) to be an existentially quantified *conjunction*. That is, saying that there are black swans is *really*, in a sense, saying that there is some object that is both black and a swan. So we get:

(3) There is an x such that x is black and x is a swan

The first occurrence of x here is attached to the existential quantifier, and the two subsequent occurrences of x are attached to the two predicates into which we break up the noun phrase ‘a black swan’. Since the same variable, x , is used throughout, the convention tells us that there is a link between them all. This link is such that the subsequent uses of x refer back, similarly to the phenomenon of pronoun anaphora, to the occurrence of x linked to the quantifier. To illustrate the similarity to pronoun usage, we could give this sentence another way:

(4) There is an object _{x} such that it _{x} is black and it _{x} is a swan.

This formulation still makes use of the variable x , but it treats it as a subscript to make explicit the connection between the two occurrences of ‘it’ and the original ‘object’.

To continue in the direction of formalizing our sentence, we introduce a symbol for the existential quantifier, ‘ \exists ’. This way we may express existential quantification at the logical form, regardless of which existential quantifier expression may be used in English (i.e., ‘there is’, ‘some’, ‘there exists’, etc.). To generate a grammatical formal sentence,

⁴ In standard First Order Logic, we have just two quantifiers, the existential and the universal (expressing the quantity of *everything*). Other expressions of quantity have always been recognized as quantifiers, but in FOL they are either derived (such as ‘two’, ‘at least three’, etc.) or else absent (such as ‘most’, ‘many’, etc.).

a quantifier symbol must always have a variable attached to it, and that variable must appear inside the scope of the quantifier. Following this, our sentence would next get written as:

$$(5) (\exists x)(x \text{ is black and } x \text{ is a swan})$$

This is symbolic enough for our purposes. We may also go further and assign predicate symbols to shorten our predicates, but the formulation in (5) will suit our purposes just fine.

Next, there is the matter of the objects that we're making use of in this discussion. When we say 'there is an x such that', or 'there is an object such that', we need to specify the group of objects we're drawing from. Variables such as x may be said to have a *range*, or a set of objects that are potential values for them. A rough way of explaining this is by noticing the difference between 'it' and 'she'. The potential referents of these pronouns are not identical—any object at all could be referred to using 'it', but only females (or certain other objects, such as cars and boats) may be referred to with 'she'.

We may see the variable as really being part of the quantifier, and from this it seems natural to say that the notion of an object is packed into the quantifier itself. We will talk about the set of objects that the variables range over as the *domain of the quantifier*. A domain is necessary in order to evaluate quantified claims, since they really are claims *about* some set of objects. An existentially quantified claim, asserting at-least-oneness of some property, is really telling us that some subset of a domain is non-empty.

How does a domain get associated with a quantifier? That depends. If we're doing exercises in a logic class, learning how the quantifiers work, we just stipulate a domain as part of an interpretation. An interpretation is an entity that gives meaning to our language, so it includes a domain of objects plus assignments of the linguistic items to objects, subsets, or sets of n -tuples from the domain. Names get assigned to objects from the domain, one-place predicates to subsets of the domain, and so on.

The stipulation of an interpretation can also happen in a natural conversation, though most often it is tacit or presupposed. A great deal of the interpretation comes along with the language we speak. When we have a conversation in English, the meaning of the names and predicates, though not necessarily precise, are defined in the way

natural languages get their meaning—organically, over a long time, through usage, style manuals, dictionaries, etc. There’s nothing preventing us from stipulating our own (partial) interpretations in an ordinary conversation, though, even if it’s not something we commonly do.

Now that we have the general idea of the notions relevant to quantifier meanings, we are in a position to discuss the truth conditions of an existentially quantified claim. Truth conditions, relative to some interpretation \mathfrak{I} , are generally given in a form that is relevantly similar to this:

$$(6) \mathfrak{I} \models (\exists x)F(x) \text{ iff for some } o \text{ in the domain, } \mathfrak{I} \models F[o]^5$$

To be thorough, I’ll unpack this definition at least enough for our purposes here. First of all, the ‘ \models ’ symbol just means that what’s on the left makes what’s on the right true. So the notation ‘ $\mathfrak{I} \models F$ ’ can be read as ‘ F is true in interpretation \mathfrak{I} ’, or ‘ \mathfrak{I} makes F true’. The truth of a sentence is always relative to an interpretation. Next, the formulation ‘ $F[o]$ ’ is an abbreviation for the result of plugging o ’s name into the free variable slots in $F(x)$. This way of stating the truth conditions allows for the possibility that the object that satisfies (or makes true) a sentence may not even have a name in the language of the interpretation. In that event, it is packed into this treatment that the interpretation would be altered with the effect of adding a name to the language and assigning that name to the object, so we could construct the sentence needed. We will use ‘ o ’ metalinguistically, to designate the object itself, and ‘ \mathfrak{o} ’ as the name that the object o would have in the language. Finally, a standard definition of an ‘interpretation’ is as follows:

An *interpretation* \mathfrak{M} for a language L consists of two components. On the one hand, there is a nonempty set $|\mathfrak{M}|$ called the *domain* or *universe of discourse* of the interpretation, the set of things \mathfrak{M} interprets the language to be talking about. When we say ‘for every x ’ or ‘for some x ’, what we mean, according to interpretation \mathfrak{M} , is ‘for every x in $|\mathfrak{M}|$ ’ or ‘there exists an x in $|\mathfrak{M}|$ ’. On the other hand, there is for each nonlogical symbol a *denotation* assigned to it. For a constant c , the denotation $c^{\mathfrak{M}}$ is to be some individual in the domain $|\mathfrak{M}|$. For an n -place nonlogical predicate R , the

⁵ This formulation comes from Boolos, Burgess and Jeffrey [5], p. 117, with some insignificant changes to suit my preferences.

denotation $R^{\mathfrak{M}}$ is the be some n -place relation on $|\mathfrak{M}|$ (which is officially just a set of n -tuples of elements of $|\mathfrak{M}|$, a one-place relation being simply a subset of $|\mathfrak{M}|$).⁶

Informally, we may think of the semantics of English as being an interpretation for the English language; it may not be a well-defined or precise model, but we can see how an English sentence will be true relative to it. Names and predicates are all assigned denotations, but they could have been assigned different ones under different interpretations.

This account of truth for the existential quantifier is just what we expect—an existentially quantified claim is true in an interpretation just in case at least one object from the interpretation’s domain has the features that the claim says are had by at least one thing. Using our previous example, we can say:

- (7) $\mathfrak{J} \models (\exists x)(x \text{ is black and } x \text{ is a swan})$ iff for some o in the domain, $\mathfrak{J} \models o$ is black and o is a swan.

So our original claim (2) is true in an interpretation just in case the interpretation’s domain has at least one object in it that is both black and a swan.

Sentence (2) is written in the English language, and we’re really just discussing interpreting it according to the standard semantics and definitions of English. That is, we mean the standard things that we normally mean when using words like ‘black’ and ‘swan’. We can use the symbol ‘ \mathfrak{E} ’ for what we might call the ‘standard interpretation’ for English, in the way that we use ‘ \mathfrak{N} ’ as the symbol for the standard model for arithmetic when doing mathematical logic. That is, \mathfrak{E} has some domain of objects, and it assigns just what we’d expect as denotations for the names, predicates, etc., of the English language. This will not be simple or precise like \mathfrak{N} , but we can still discuss \mathfrak{E} as a heuristic at least.

We might ask if there is one default domain of objects for English, or if that varies in different conversations, depending on some feature, such as the context of utterance, or the intentions of the speaker. Since interpretations are individuated by their domains as well as their assignments, this question could be seen as a follow-up to the discussion

⁶ Boolos, Burgess and Jeffrey [5] again, pp. 103-104.

of \mathfrak{E} in the previous paragraph: namely, is there really a standard interpretation \mathfrak{E} for English, or do we have different domains (and thus different interpretations) in different settings?

As mentioned above, it is part of my view that we have different domains in different contexts, at least in association with the expression ‘there is’. In this sense, we will also have different interpretations in different contexts, though I’m not suggesting any difference in assignments (other than what is forced by the changing domains). This would not preclude English from having a *standard* interpretation, though. In fact, I think that there is a standard interpretation that we make use of some of the time when we use ‘there is’, and it is also the interpretation whose domain is associated with ‘there exists’ invariantly.⁷

The Logical Approach Applied to Superheroes

Before moving on to the other approach to quantifier semantics, I would like to look briefly back at the example about superheroes and apply to it what we’ve discussed about the logical treatment. We begin by looking at the logical form of the sentence, translating the sentence, “There are superheroes who are animals but also people,” into a kind of hybrid predicate logic notation, such as:

$$(\exists x)[x \text{ is a superhero and } (x \text{ is an animal and } x \text{ is a person})]$$

We can, using parentheses as punctuation, reflect a certain grouping, showing that ‘x is a superhero’ is in a sense part of the subject, while the predicate includes ‘x is an animal’ and ‘x is a person’. This is how we read a standard Aristotelian I-type sentence—as an existentially quantified conjunction, where we’re saying about the left conjunct that it has the features described in the right conjunct. With the tools of first-order predicate logic, though, we cannot capture the contrastive element of the ‘but also’.

My claim about context-sensitivity is essentially a claim that in different contexts, we intend different interpretations to be used for evaluating our claims. In both cases, the rules of the English language are the ‘interpretive’ part of the interpretation, but

⁷ This point falls in line with the discussion that will occur in the next chapter on ‘reference magnetism’.

the domain of the interpretation shifts depending on our interests. Intuitively, this makes sense—Nick intends his utterance of (1) to be interpreted much differently by Mirah than he would by ontologists, and he’s speaking English in both settings, so that difference really must come in the domain of discourse he intends. Again, the intentions of the speaker come in at a pre-formal level, but the place where they make a difference in the formal sense is when we bring an interpretation to bear on the sentence uttered.

I’ll return to this example again after discussing the linguistic treatment of quantifier semantics in the next section.

2.3.2 Linguistic Treatment

Next we move to a discussion of the existential quantifier as treated in linguistics, as one determiner among many. Quantifiers are forms of determiners, but determiners do much more than just quantify. A general way of looking at determiners is as functions that take properties (the referents of noun phrases) as their arguments. The result is another function that takes properties again (though now the referents of verb phrases) as its argument and results in a truth value. Equivalently, we could just see the determiner as a two place function, taking in two properties as arguments, resulting in a truth value.⁸

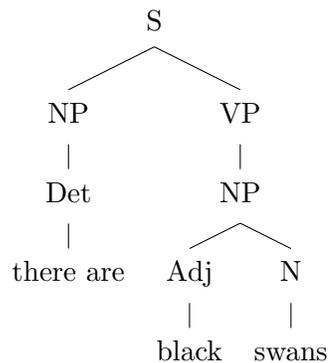
Seen in this general functional way, it is clear that many other expressions besides just our usual quantifiers are determiners. Aside from the quantifiers like all, some, most, many, two, etc., we have several kinds of categories of determiners. Demonstratives (this, that, these, etc.) count as determiners, as do proper names. There are possessive determiners (John’s, my, etc.), and there are anaphoric determiners (fewer, half again as many, etc.).⁹ Although my discussion here will just be about existential quantifier expressions, it is good to be aware that such expressions are seen a bit differently from a linguistic perspective than from a logical one. In logic, whether we are aware of it

⁸ Cresswell [15] gives the following explanation: “Assume that an individual variable like x has the category of *name* (n). Then *shouts* x will be a *sentence* (s), and nouns like *author* and verbs like *shouts* will make a sentence out of a name—they will be in category (s/n). *not* will be in category (s/s) since it makes a sentence out of a sentence, and so on. The abstraction operator turns a sentence back into a (one-place) predicate, i.e. into a (s/n). A determiner can be thought of as in category ((s/(s/n))/(s/n)), i.e. it takes a predicate (a noun) and turns it into a one-place operator which makes a sentence out of a predicate (a verb),” p. 137.

⁹ See Keenan and Stavi [27] or Keenan [26] for examples of large listings that give a sense of the varieties of natural language determiners.

or not, we tend to think of the existential and universal quantifiers as quite special, holding a privileged position in first-order logic as our two quantifiers, each definable in terms of each other (plus negation). But first-order logic clearly does not capture all of our natural language—it cannot even adequately capture the logic of some relatively easy-to-comprehend determiners like ‘most’.

So let’s now return to the example given above as (2), “There are black swans,” to illustrate how we might break down an existentially quantified claim with the treatment given to us in modern linguistics, rather than in logic. First of all, we might give a diagram of the sentence, categorizing the parts of it, as follows:



In this sort of diagram, we have S=Sentence, NP=Noun Phrase, VP=Verb Phrase, Det=Determiner, Adj=Adjective, and N=Noun.

We see that the structure of a sentence like (2) gets represented a bit differently in the linguistic framework than in the logical one. The treatment of the truth conditions for such a claim is also a bit different. Keenan [26] gives us a uniform way for describing the truth conditions for a variety of determiners. For example, he gives these conditions, which are intended to be general for all properties A , B :

$$(8) \text{ (EVERY)}(A)(B) = \mathbf{T} \text{ iff } A \subseteq B$$

$$(9) \text{ (NO)}(A)(B) = \mathbf{T} \text{ iff } A \cap B = \emptyset$$

I’ll discuss this terminology in more detail shortly, but the rough idea is that, in the first case, an Aristotelian A-type sentence such as, “Every man dies,” is true just in case the

set of men is a subset of the set of objects that dies. And an E-type sentence like, “No man is an island,” is true just in case the intersection of the sets of men and islands is empty. Keenan does not explicitly give the same sort of formula for the existential quantifier, but it is easy enough to see how that would go:

$$(10) \text{ (THERE ARE)}(A)(B) = \mathbf{T} \text{ iff } |A - B| \geq 1$$

Again, the basic idea is that an I-type sentence such as, “There are men who enjoy pain,” is true just in case the result of subtracting the set of things that enjoy pain from the set of men gives us a set with cardinality of 1 or more, i.e., the set of men who enjoy pain is non-empty.

Now, to move through this formula a bit more slowly and carefully, let’s recall our example (2) from above, “There are black swans.” Using the above formulations, there are two ways of representing the sentence. They’re equivalent, but we may prefer one way over the other, depending on how we view the quantifier and its potential restriction. These two formulations are the following (where D in (12) is just the domain of objects associated with the claim, which I’ll discuss more shortly):

$$(11) \text{ (THERE ARE)}(\mathbf{BLACK})(\mathbf{SWANS})$$

$$(12) \text{ (THERE ARE)}(D)(\mathbf{BLACK} \cap \mathbf{SWANS})$$

Let’s take (11) first. We can read it as follows: **(THERE ARE)** is a function from properties to functions, while **(BLACK)** and **(SWANS)** are both properties. **(THERE ARE)** takes **(BLACK)** as its argument, and it gives another function, which we could call, in this case, **(THERE ARE BLACK)**. This is now a function from properties to truth values, taking **(SWANS)** as an argument and giving the truth value True or False, depending on whether there actually are black swans or not.

To put this in set-theoretic lingo, we can say that **(BLACK)** and **(SWANS)** are associated with subsets of our domain D, the first just being the set of black objects in D, and the second being the the set of swans in D. The function **(THERE ARE)** grabs onto the first subset, the set of black objects, and compares it to the other subset, the set of swans. If the intersection of these two subsets is non-empty, then the value of (11) is True; if it is empty, then the value is False.

Sentence (12) differs as follows. If we want to be clear that the quantifier itself is unrestricted, we might fill the first argument with D , representing the entire domain of objects associated with the claim. Then (**THERE ARE**), which is a function from properties to functions, takes (D) as its argument, and it gives another function, (**THERE ARE ENTITIES**). This is now a function from properties to truth values, taking (**BLACK** \cap **SWANS**) as an argument and giving the truth value True or False, depending on whether there actually are black swans or not.

To put this in set-theoretic lingo as well, we can say that (D) and (**BLACK** \cap **SWANS**) are associated with subsets of D , the first just being D itself, and the second being the intersection of two subsets, one the set of black objects in D and the other the set of swans in D . The function (**THERE ARE**) grabs onto the first subset, D itself, and compares it to the other subset, the set of black swans. If the intersection of these two subsets is non-empty, then the value of (12) is True; if it is empty, then the value is False.

Keenan discusses the equivalence of these two kinds of formulations in the following passage:

We have been interpreting Ss like *Some swans are black* directly as **SOME** (**SWAN**)(**BLACK**). But in standard logic it is represented as $(\exists x)(\text{Swan}(x) \ \& \ \text{Black}(x))$, where the quantification is over all the objects in the model, not just the swans. In our variable free notation this would be **SOME**(E) (**SWAN** \cap **BLACK**).¹⁰ This formulation eliminates the restriction to swans in favor of quantifying over all the objects in the universe, and it preserves logical equivalence with the original by replacing the original predicate property **BLACK** with an appropriate boolean compound of the noun property and the predicate property, **SWAN** \cap **BLACK**. Thus *some* does not make essential use of the domain restriction imposed by the noun argument. The same equivalence obtains if we replace *some* by e.g. *exactly two*. *Exactly two swans are black* is logically equivalent to *Exactly two objects are both swans and are black*.¹¹

¹⁰ Keenan uses E the way I've been using D , as the domain of objects associated with our claim or model.

¹¹ Keenan [26], pp. 58-59.

Keenan calls quantifiers that have this feature ‘sortally reducible’, when we could replace the first argument with the entire domain D and the second argument with a boolean function of the noun and predicate properties. Quite a few determiners turn out to be sortally reducible, and he discusses them in Keenan [25].

Notice that this domain is the same as that which shows up in the logical treatment—the universe of discourse, or the set of objects that underlies our conversation. So this notion is common and crucial to both treatments. I’ll say more about domains as we go along, continuing into the next chapter. Now it will be useful to stop and say a little more about functions and how we might think about them in this setting.

Picturing the Functions

In the Linguistic Treatment, we talk about determiners as functions from properties to other functions (which go from properties to truth values). It is helpful when discussing functions to look at them two different ways: intensionally (i.e., descriptively) and extensionally (i.e., set-theoretically). I’ll do both here. Before getting to determiners, though, I’ll discuss a simpler example of the meaning and the operations of a function.

The *successor* function—let’s call it s —is one of the simplest functions there is. This function takes any integer as an argument and gives the successor of that number, i.e., it adds one to that number. So the successor of 1 is 2, the successor of 45 is 46, the successor of 1,911 is 1,912, and so on. We’ll write equations such as these as: $s(1)=2$, $s(45)=46$, and $s(1,911)=1,912$.

Now, I’ve already explained everything that needs to be explained about s , but we can see that there are a couple different ways that we go about explaining it. One is to describe what the function *does* (usually in active, dynamic language, as if the function is an agent), and another is to provide *examples* of particular cases. The intension of the function, then, is the descriptive meaning of it; the sort of meaning that we would see if we looked in a dictionary. The extension of the function, on the other hand, is a set of objects, where each object is a particular example of the function’s input/output matrix. So we get:

Intension of s : A function that takes an integer as argument and gives as output the successor of the argument, i.e., the argument integer plus one.

Extension of s : $\{\langle x,y \rangle: y=x+1\}$

We read the set in the second of these as, essentially, ‘the set of ordered pairs $\langle x,y \rangle$ such that the second item (y) is just the first item (x) plus 1.’ So this set would start off like this:

$$\{\langle 1,2 \rangle, \langle 2,3 \rangle, \langle 3,4 \rangle, \langle 4,5 \rangle, \langle 5,6 \rangle, \dots\}$$

The only real restriction in the extension of a function is what we can call ‘uniqueness-of-output’—the restriction that we cannot have two members such that they contain the same first item and different second items. That is, for one input, a function must always give at most one output. If you put 4 into the successor function, you will get 5 out every time. This uniqueness-of-output restriction holds for all functions. The converse holds for the successor function (since each number also has a unique predecessor), but it does not hold for all functions. Take addition, for example. Both $3+5=8$ and $4+4=8$. That is, the extension of the addition function contains both $\langle \langle 3,5 \rangle, 8 \rangle$ and $\langle \langle 4,4 \rangle, 8 \rangle$. But uniqueness-of-output still holds, since the input of $\langle 3,5 \rangle$ will always only give you the same output of 8.

When we move to a determiner, these functions get a bit more complicated. A determiner like ‘there is’, as mentioned above, is a function from properties to functions. And the functions spit out by the first kind are functions from properties to truth values. We can view this formally in the following way. Let our function symbols be $f, f^*_1, f^*_2, f^*_3, \dots$, and let our property symbols be p_1, p_2, p_3, \dots . We’ll let f be our determiner function and f^*_n be the function that is the result of operating f on some property p_n . We can then see these functions extensionally as follows:

$$f: \{\langle p_1, f^*_1 \rangle, \langle p_2, f^*_2 \rangle, \langle p_3, f^*_3 \rangle, \dots\}$$

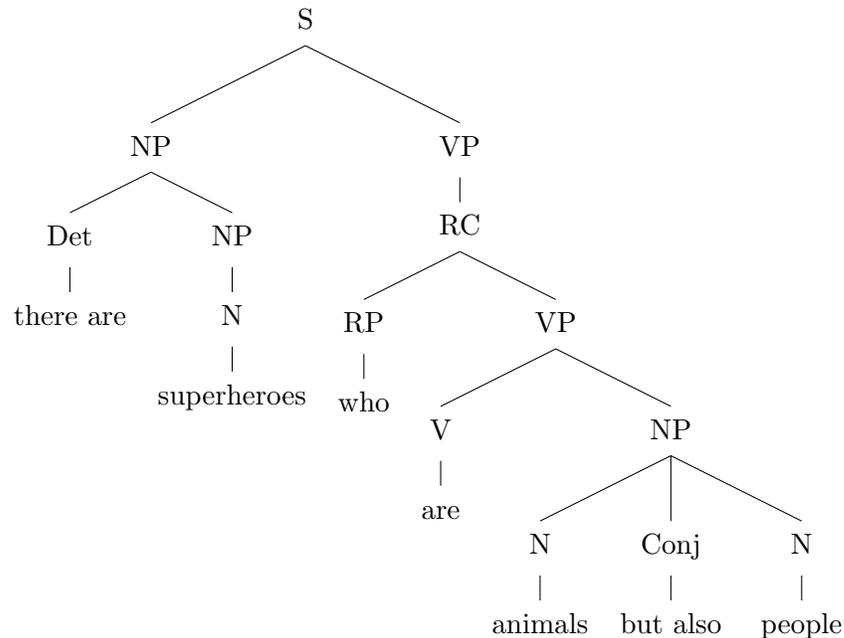
$$f^*_n: \{\langle p_1, T \rangle, \langle p_2, T \rangle, \langle p_3, F \rangle, \dots\}$$

In the second of these, the second slot of the ordered pairs would be either T or F, depending on whether the sentence in question is actually true or false—these are just filled in to show a complete structure. Another way we might represent the functional structure of a sentence with a determiner and two predicates, then, is: $[f(p_n)](p_m)$, or, equivalently, $f^*_n(p_m)$.

Returning to (2), we could say that a function t is the determiner expressed by ‘there is’, p_b is the property of being black, and p_s the property of being a swan. The determiner made out of applying t to p_b would be t^*_b , so we could express (2) in functional language either as $[t(p_b)](p_s)$, or, equivalently, as $t^*_b(p_s)$.

The Linguistic Approach Applied to Superheroes

Let’s return now to the superhero example and view it through the lens of the linguistic treatment, focusing on the functional aspect of the words involved. The linguistic approach begins with a different take on the representation of the syntax of the sentence from that given by the logical treatment. Rather than translating it to its first-order representation, we might diagram the sentence to reveal its structure as follows:



This approach may look quite a bit different, but it does not exactly reveal to us anything new in what’s happening with the sentence. We see that ‘superheroes’ goes under the initial NP node, which groups it in a way that was reflected with parentheses above in the first-order translation. The main difference we might notice is the absence of

variables in a parse tree like this one. This structure is displaying something slightly different to us than a translation—rather, it’s showing us the grammatical ordering and dependencies, and we can read off of this the functional structure of the sentence. Recall that a determiner is treated as a function from properties to functions, so the initial NP node shows us the determiner expression (‘there are’) and an expression of its argument (‘superheroes’). The result of applying the *there are* function to the property of being a superhero is another function that takes the property (of being both an animal and a person) expressed by the predicate as its argument. We can represent this in functional notation as follows:

(13) **(THERE ARE)(SUPERHEROES)(WHO ARE ANIMALS BUT ALSO PEOPLE)**

But recall that it is equivalent to write this same sentence in the following, slightly different, functional notation, making the domain an explicit component as the argument of the determiner function:

(14) **(THERE ARE)(D)(SUPERHEROES \cap WHO ARE ANIMALS BUT ALSO PEOPLE)**

It is this second formulation upon which I’d like to focus at the moment. This formulation makes it clear that there must be some domain D accompanying the utterance of the sentence, since the sentence includes a function that takes D as its argument.¹²

So where does this D come from? It must just be the domain of discourse in the conversation in which the sentence is uttered. And that domain of discourse, I claim, is supplied by the intentions of the speaker.

When we are talking about the actual functions and properties displayed in this functional representation, we are again talking about the *content* of the utterance on certain occasions. So again, my claim is that a meta-semantic feature of the context (the speaker’s intention) plays a determinative role in fixing a crucial part of the content of the utterance, *viz.* the domain that gets associated with the determiner function. Here

¹² The previous formulation ultimately requires this too, since the property (**SUPERHEROES**) will be seen as being associated with the set of all superheroes, which must, on a standard treatment, be seen as a subset of some larger domain or other.

the domain works as the initial argument of the determiner function. This is a little different in structure to the logical treatment, where we talk about an interpretation that gets used to give meaning and truth values to the expression. But the effect is the same, including a domain at the level of content that is crucially associated with the quantifier.

In fact, since (13) and (14) are equivalent,¹³ it is reasonable to see the domain either as being the argument for the function or as *part of the determiner function itself*. In this sense, it is not just the proposition expressed that changes from context to context, but the quantifier itself, i.e., the content of the existential quantifier expression ‘there is’.

In the example of Nick uttering (1) in different contexts, my claim was that he intended different domains at the two different times/places of utterance. Let’s call the first domain D_1 and the second D_2 . At this point, I am not making any strong claims about which sets these domains are; my claim is just that they are not identical. At a minimum, D_1 is intended to have nonexistent objects in it, and D_2 is not. And since the domains are fixed via the speaker’s intentions, it is reasonable to say that these two domains would not be identical. What we must say, then, is that when Nick utters (1) in the ordinary context, the content of his utterance could be represented in either of the following ways:

(15) **(THERE ARE)_(D₁)(SUPERHEROES \cap WHO ARE ANIMALS BUT ALSO PEOPLE)**

(16) **(THERE ARE_{D₁})(SUPERHEROES)(WHO ARE ANIMALS BUT ALSO PEOPLE)**

But when Nick utters (1) in the ontological context, the content of his utterance is a bit different, represented in either of the following two ways:

(17) **(THERE ARE)_(D₂)(SUPERHEROES \cap WHO ARE ANIMALS BUT ALSO PEOPLE)**

(18) **(THERE ARE_{D₂})(SUPERHEROES)(WHO ARE ANIMALS BUT ALSO PEOPLE)**

¹³ See the discussion of this above.

In each case, we see that some constituent of the proposition differs, depending on the context in which the sentence was uttered. Between (15) and (17), the domain, seen as the argument for the determiner functions, while between (16) and (18), the determiner functions themselves differ, since the domains—here seen as part of the functions—differ.

This is what it means when we express the view that ‘there is’ is context-sensitive—the content of utterances made using ‘there is’ as an existential quantifier expression varies depending on the context. And that variance in content is reasonably seen to be located in the determiner function itself, giving us different functions in different context. As mentioned above, I see the speaker’s intention, or the intended domain of the speaker, as that feature of the context upon which the variance depends, or upon which the content of the quantifier expression is fixed.

To express this in slightly different functional notation, we could say (as above) that a function t_n is a determiner expressed by ‘there is’, associated with domain D_n . We will have p_s , the property of being a superhero, and p_{ap} the property of being both an animal and a person. The determiner made out of applying t_n to p_s would be $t_n^*_s$, so we could express (1) in functional language as $t_n^*_s(p_{ap})$.¹⁴

Now we are in a position to express the variance between contexts in this kind of functional notation. Let t_n be the determiner function expressed by ‘there is’ with intended domain D_n . In the ordinary context, Nick’s utterance of (1) will have a functional content represented as $t_1^*_s(p_{ap})$. In the ontological context, Nick’s utterance of (1) has the similar, but slightly different, functional content represented as $t_2^*_s(p_{ap})$. Again, this is just a different way of expressing the variance at the level of content that occurs, depending on which context Nick is in when he utters (1).

This notation corresponds nicely to the treatment given above in (16) and (18). If we like, we could also modify our notation to match the equivalent treatment of (15) and (17). Just let p_{d1} and p_{d2} be the properties associated with being members of D_1 and D_2 , respectively,¹⁵ and let p_{sap} be the property associated with the set that is the union

¹⁴ As mentioned above, there are a couple equivalent ways of expressing this in functional language. The alternate way we might express (1) is $[t_n(p_s)](p_{ap})$. I see these options as giving us a way to highlight the choice of looking at the domain as being *part* of the determiner function or as being an *argument* of the determiner function. For my purposes, I don’t believe this choice really makes a difference, but it may for other inquiries.

¹⁵ If it’s not clear by now, I’m helping myself to a non-minimalist view about properties. I think that very little hinges on a precise view about properties here, so I will not discuss them any further than

of the sets associated with p_s and p_{ap} . Then we may represent the functional content of Nick's utterance of (1) in the ordinary context as $t^*_{d1}(p_{sap})$, and in the ontological context as $t^*_{d2}(p_{sap})$.

2.4 On 'Exists'

So far in this chapter, I have focused entirely on the expression 'there is'. In the previous chapter, I built a case for a difference between 'there is' and 'there exists', arguing that the notion of ontological robustness gives us reason to see such a difference. My claim was that the difference we find in examples that make use of the two existential quantifier expressions is such that 'there exists' claims appear to be invariantly ontologically robust, while a 'there is' claim may be ontologically robust in one setting but not in another.

This chapter has extended my view about the difference between 'there is' and 'there exists' by filling out the details of a contextualist view of 'there is'. I have less to say about 'there exists' at this point, since the purpose of this chapter is not to extend substantive new arguments for the view. Rather, the focus here is on explication of the details of the view. The preliminary argument lies in Chapter 1's claims about ontological robustness, and there will be further argument to come, primarily in Chapter 4, when this view is compared holistically with its competitors.

At this stage, it will at least be useful to fill out my positive view by making clear that, while I hold that 'there is' is a context-sensitive expression, I maintain that 'there exists' is an invariant expression. That is, its meaning is fixed from context to context. Aside from this difference, 'there is' and 'there exists' are, of course, very similar expressions. Since they are both existential quantifier expressions, they both denote determiners. And the functional operation of the determiner picked out by 'there exists' will be the same as that picked out by 'there is', except for the behavior of the function with respect to the speaker's intentions.

This difference can be seen as analogous to the kind of difference we see between other context-sensitive expressions and their invariant counterparts or neighbors. Consider

just to say that each property will be associated with the set of all objects that have that property, and for each set of objects, it makes sense to speak of an associated property, e.g. the property of being a member of that set.

the difference between ‘yesterday’ and ‘Monday’. As I write this, today is Tuesday, so I could interchange those two expressions in many of my sentences without any problems. I could say, “I took my truck to the shop yesterday/on Monday,” or, “I had a much better day yesterday/on Monday than I’m having today.”¹⁶ The expressions ‘yesterday’ and ‘Monday’ both denote a day of the week, but ‘Monday’ is invariant in that it always refers to the same day of the week, while ‘yesterday’ denotes different days of the week, depending on the day on which it’s uttered.

One way of viewing this difference is that context-sensitive expressions have a big hole in their meaning that needs to be filled—a day-shaped hole for ‘yesterday’, a speaker-shaped hole for ‘I’, or a domain-shaped hole for ‘there is’—while invariant expressions have no such hole. A bit less poetically, we can use the language of functions and say that context-sensitive expressions are like functions in search of an argument, while invariant expressions are like constant functions—there’s a nominal sense in which they take in an argument, but since they have a constant output, the input really plays no role in their meaning. So any time we speak, we do so in a context, and some features of that context get used as needed arguments for context-sensitive expressions, but that context plays at most a vacuous role in invariant expressions.

These analogies should sound plausible when we compare ‘there is’ with ‘there exists’, since the word ‘exists’ carries with it its own connotations that are not present in ‘is’. The expression ‘there is’ is flexible and adaptable, while ‘there exists’ is essentially a specification of ‘there is’. That is, the word ‘exists’ fills out the meaning of ‘there is’, making it clear—or at least clearer—what we are quantifying over. We may have plenty of questions about which things exist (evidenced by the popularity of physics and metaphysics), but there is a presumption when we use ‘there exists’ that there is some domain that includes all existing objects, and that’s the domain we intend now. There is no such presumption packed into ‘there is’, as we saw in the examples in Chapter 1.

¹⁶ For those worried about the fact that ‘Monday’ here is ambiguous, or could be referring to multiple different past Mondays, we could fill out the example and use ‘Monday, December 3, 2012’ instead.

Chapter 3

Unrestricted Domains and Nonexistent Objects

3.1 Introduction

In this chapter, I'll take up two topics that follow close on the heels of the contextualist view presented in the previous chapter. These topics are the following:

1. Although it may be tempting to think so, my contextualist view *is not* about quantifier domain restriction, but rather about the *unrestricted* domains of the quantifiers, and
2. Although it may be tempting to think so, my view *does not* require me to be committed to there being nonexistent objects, since the demonstrative model of contextualism I hold provides for a natural way to make sense of reference failure in claims that purport to quantify over nonexistent objects

I will discuss these topics in turn, and they each should help clarify contextualist view while also placating some potential concerns about it.

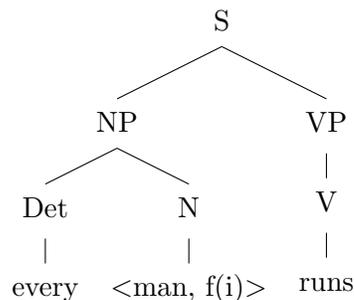
3.2 Why this isn't about Quantifier Domain Restriction

Quantifier domain restriction, discussed recently in works such as Stanley and Szabó [44], *is not* what I'm talking about when I give the contextualist view of 'there is' presented

in the previous chapter. To say it plainly, contextualism is a view about the *unrestricted* domain of the existential quantifier, and how that varies from context to context. Quantifier domain restriction, on the other hand, is a view about how, once we're in a context and an unrestricted domain is fixed, those unrestricted domains get further restricted. It will be tempting to conflate these two issues, so for clarification, and to pre-empt objections that my view is nothing more than a view about quantifier domain restriction dressed up to look a bit differently, it will be helpful to address it head on.

I'll briefly explain Stanley and Szabó's view about quantifier domain restriction here, since we need to have a few details of a representative view in order to explain why it is a different sort of issue. Their view holds, among other things, that nominal expressions have a hidden component at the semantic level that gets fixed by context. This component is a set that creates a restriction on the domain of quantifiers under which the nominal expression appears. Outside of any context, this component is constructed of a function variable and an individual variable, e.g., ' $f(i)$ '. Upon a use of the nominal expression, the context of its use fills in these variables with a function and an individual, respectively.

For example, consider the claim, "Every man runs."¹ On a particular utterance of this sentence, chances are good that the intention is not to say that every man in the whole world runs, but rather that every man within some particular domain runs. The Stanley/Szabó view represents this claim as follows:



Here what fills the noun slot is the pair $\langle \text{man}, f(i) \rangle$, which is the set that is the

¹ Examples of quantifier domain restriction are more common, or at least clearer, with universally quantified claims, so we'll use this example to begin illustrating the view.

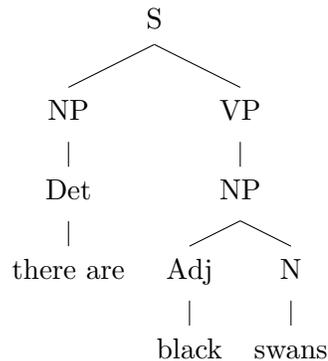
intersection of the set of all men and the set determined by the result of filling in the variables ‘f’ and ‘i’ in the expression ‘f(i)’.

So if the speaker intends to be talking about, say, just those men who are in the University of Minnesota Department of Philosophy (UMDP), then the context would replace the function variable ‘f’ with a function—let’s call it ‘g’—from organizations to their members. And context would also supply the UMDP as the particular organization—let’s call it ‘p’—to replace the individual variable ‘i’. The set picked out by ‘<man, g(p)>’ would be the intersection of the set of men with the set of members of the UMDP. That intersection is just the set of male members of the UMDP. So the claim would now say that everything in the set of male members of the UMDP runs. This lines up nicely with the intentions of the speaker.

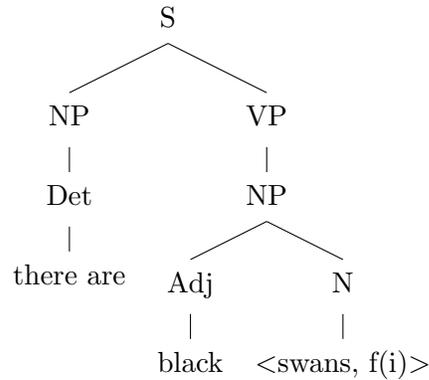
For a case with the existential quantifier, let’s revisit an example from the previous chapter:

- (1) There are black swans.

The parse tree given there looked like the following:



Modifying this tree in light of Stanley and Szabó’s view, we would get:



Notice that the difference comes just in the nominal slot, replacing ‘swans’ with ‘<swans, f(i)>’. What is being added is a function that, along with features of the context, serves to restrict the domain. But this means that in order for that variable to do any work, there must already be a domain that will be restricted, i.e., there must already be an unrestricted quantifier domain. The unrestricted domain is conceptually prior to any restriction placed on it. These two activities—fixing an unrestricted domain and then placing some kind of restriction on it—are distinct in their purpose. The unrestricted domain includes, we might say, all the kinds of things that the speaker is helping herself to in the conversation. The restricted domain is much more specific, and it only involves one particular kind of thing. It may answer a question such as, “Which men are you talking about? All men?” or, “There are black swans where? Here in the U.S.?”

Examples of quantifier domain restriction are easy to come by with the universal quantifier. When someone says, “All the beer is in the fridge,” she is very likely intending a restriction on which beer is in play—it would be very odd if she meant all the beer in the universe.

With the existential quantifier, examples are a little less straightforward, but we can still see how they work. Consider the following situation. Michael is a member of the Perth Association of Black Swan Travelers. (Members of PABST go off and travel the world, looking for other places that have black swans.) Michael is traveling in the U.S., and he spots black swans in Thief River Falls, MN, so he buys a postcard with a photo of the Thief River and big lettering that says “Thief River Falls” on it. He writes, “There are black swans!” on the back and addresses it to PABST, sending it from the

local post office. When PABST receives this card, they will get the message that there are black swans *in Thief River Falls, MN*, even though Michael didn't explicitly restrict his claim.

Now, there are multiple ways to explain the semantics of Michael's postcard claim. Quantifier domain restriction à la Stanley and Szabó is one of them, saying that the context of the utterance (which here would include the front of the postcard and the postmark) will provide the information needed to fill the appropriate variable in the nominal, and the domain of the quantifier will then be restricted in some reasonable way. So Michael wouldn't actually be making a flat-out ontological claim about there being at least one black swan in the universe—rather, he'll be saying that there's at least one black swan in Thief River Falls, MN. Competitors to Stanley and Szabó may argue that the restricted claim is just on implicature—i.e., that the restriction is done at the level of pragmatics, not semantics. This is not important for the point I am making in this section. I am simply trying to point out what quantifier domain restriction is in order to distinguish it from my own contextualist view.

Notice that whether the restriction of the domain happens one way rather than another, the dispute between Stanley/Szabó and their competitors in this case isn't about the unrestricted domain. They are concerned with how the unrestricted domain becomes restricted, whether that's a semantic or a pragmatic process. If we focus on the unrestricted domain itself, we will ask different questions, such as whether the claim is ontologically robust or not. In this case, it seems clear that, regardless of what we have to say about the restriction of the quantifier domain, Michael's claim is ontologically robust. In the closest worlds where Michael does not believe that black swans exist, he would not send a postcard back from Thief River Falls like the one described—at least not in a way that counts as a sincere utterance of the claim.

This difference between the unrestricted domain and a restriction of that unrestricted domain is a subtle one, and it's especially confusing if we want to hold context-sensitive views about both of them. (I'm not arguing for that combination, but I'm also not quibbling with the Stanley/Szabó view, so the combination is a live possibility for me.) The basic difference, though, is that the unrestricted domain includes whatever objects we consider fair game for our entire conversation, whatever may help us communicate all of the thoughts we'd like to communicate, or whatever *kinds* of objects we would

like to help ourselves to. A restriction of that domain, on the other hand, includes just whatever objects we would like to focus on *right now*, at this particular point in our conversation, to make the point we would like to make with the appropriate amount of precision.²

3.3 Quantifying Over Non-Existent Objects

The other issue raised in the previous chapter but postponed until now is that of domains that include nonexistent objects. I stressed that I don't actually believe in nonexistent objects or, by extension, domains that include nonexistent objects. This may sound disingenuous or tricky, so it is worth unpacking this issue in some detail here. In his recent book *Talking About Nothing*, Jody Azzouni frames the kind of problem I'm dealing with quite clearly:

One thing going on is that we have a powerful intuition that when we talk we always talk *about* things. Our talk is always directed toward one or another topic. This is not the mere grammatical impression that every sentence has a “subject.” The intuition runs deeper than that—it's the impression not only that our talk is about things, but that how it is with those things determines the truth or falsity of that talk. If this description of our intuitive impression (of how it is with what we say) is right, then a kind of desperation naturally sets in when we engage in conversation about the nonexistent, for there is nothing for that talk to be directed toward. Further, there is nothing true or false to be said about nothing. If we have nothing as a subject of what

² An issue that I am purposely side-stepping here is that of absolute generality, or whether there is such a thing as an absolutely unrestricted domain available to speakers. This is a serious issue, discussed at length in Rayo & Uzquiano [40], but it does not need to be settled for anything I have to say here. The distinction between what Stanley/Szabó are doing and what I'm doing should still be clear—I'm talking about domains that are initially associated with our quantifiers, and they're talking about nominal restrictions of those domains. I do talk about domains as if there is a domain of all existing objects that always gets associated with 'there exists' as its unrestricted domain, and often also with 'there is', while speakers sometimes intend larger domains to be associated with 'there is'. This should not be taken to be a robust claim about absolute generality, though. All I really need to make the claim that I want is the idea that speakers sometimes intend to include nonexistent objects in the unrestricted (or initial) domain when using 'there is', but not when using 'there exists'. I speak more loosely than that, operating as if it makes sense to talk about the domain of all existing objects, but if that troubles you, consider the same claim in the weaker language, and it should be fine.

we say, then we have nothing meaningful that we can say about *it*.³

The first part of this passage is about our ordinary grammatical intuitions—the desperation discussed in the second part is a desperation felt by philosophers when we try to grapple with discourse about the nonexistent. I take the intuitions Azzouni discusses to be fairly stable, and I want to respect this notion that we feel strongly as if we are directing our talk towards objects like superheroes and other nonexistent objects. So in this section, I’ll expand on my view in a way that addresses these issues, specifically when it comes to the domains associated with the quantifier meanings—this will be where the action is for existentially quantified claims. The notion of truth is also central to Azzouni’s point, and I will discuss this below, both our intuitions about the truth of quantified claims and my view about the truth of those that purport to quantify over nonexistent objects.

The contextualist view I’ve presented treats quantifier expressions on the model of demonstratives (as opposed to, say, indexicals). On this view, the reference of the quantifier expression, and thus also the domain associated with it, is fixed by the intentions of the speaker. There may be no physical pointing to accompany the utterance and lead us to a domain, but a physical pointing is not necessary. Kaplan’s remarks about ‘directing intentions’, in cases where there is no physical pointing, gives us a model that appears to be a natural fit.

Kaplan originally thought a physical act of pointing was crucial, but he changed his mind and came to think of “the directing intention, at least in the case of perceptual demonstratives, as criterial,” and “the demonstration as a mere externalization of this inner intention.”⁴ My use of this notion is a further extension of Kaplan’s, since existential quantifier expressions need not be perceptual in a way analogous to perceptual demonstratives. For example, we say that when Nick, in the ordinary setting, uttered:

(2) There are superheroes who are animals but also people.

he intended to quantify over nonexistent objects. There is nothing perceptual here about the superheroes involved—Nick is not relying on our ability to see any superheroes in the context, or anything like that. But a quantified claim like (2) that appears to be

³ Azzouni [3], p.7.

⁴ Kaplan [22], pp. 582.

straightforward and non-figurative reveals to us some information about the domain of the quantifier meaning the speaker intends. When Nick asserts (2), the noun phrase (‘superheroes’) after the quantifier exhibits his intention to use a quantifier whose domain includes objects of the sort picked out by that noun phrase (‘superheroes’). Since superheroes do not exist, they are not in any domains; so we can say that Nick’s intended domain must be one that includes nonexistent objects. This does not mean, however, that Nick is *successful* in his reference attempts.

My account of such assertions is similar to what we might say about unsuccessful uses of true demonstratives—they are instances of some kind of reference failure. In what follows, I will discuss what that reference failure looks like for quantified claims like (2) in an ordinary setting.

3.3.1 Non Denoting Quantifier Expressions

The main issue here is what to say about the content of a claim like (2), uttered in an ordinary setting. The only part of (2) at issue is the domain associated with the quantifier meaning, which is just one constituent of the proposition expressed. Because of this, I believe the content of (2) in these settings is relevantly similar to the content of claims that include empty names or non-referring demonstratives.

David Braun [6] [7] has done extensive work in this area, particularly in describing and defending a view about the content of claims that include empty names. He argues that these claims express ‘unfilled’ or ‘gappy’ propositions, i.e., propositions that just have a gap where the object denoted by the empty name should have gone. I will accept the main thrust of Braun’s *Gappy Proposition* view and align myself with his defenses of the view. His view seems plausible in the case of empty names, and the defenses he offers, adapted to the issue of non-referring quantifier expressions, will placate a great many concerns someone might have about the view I’ve presented.

In the case of claims like (2), my account from the previous chapter of the content of this claim was that it could be represented using functional notation in either of these two ways:

- (3) **(THERE ARE)(SUPERHEROES)(WHO ARE ANIMALS BUT ALSO PEOPLE)**

- (4) **(THERE ARE)(D)(SUPERHEROES \cap WHO ARE ANIMALS BUT ALSO PEOPLE)**

My refined or amended claim now is that, in contexts such as the ordinary one in which Nick is talking to his daughter Mirah, the intended domain associated with the quantifier must be a domain that would include nonexistent objects. Nick may intend this to be so, but since there is no such domain (which is my belief), then there can be no domain to act as argument for the function here. With the two ways of representing the content of Nick's claim, we see two ways for displaying the gap in the content on this particular utterance of it:

- (5) **(____)(SUPERHEROES)(WHO ARE ANIMALS BUT ALSO PEOPLE)**

- (6) **(THERE ARE)(____)(SUPERHEROES \cap WHO ARE ANIMALS BUT ALSO PEOPLE)⁵**

We see that one way of looking at it shows a gap in the content where the quantifier meaning (i.e., the determiner function) itself should be, and the other shows a gap where the unrestricted domain should be. My claim is that the unrestricted domain is really part of the content of the quantifier expression, so the second way of displaying it is really just a way of making explicit the fact that the unrestricted domain is there in the content. These two ways may be equivalent ways of representing the content, but I believe (5) to be the more accurate representation of the content expressed by Nick in this ordinary context.

There are different ways we may choose to represent the content of a claim like Nick's. As I represent it here, we have three sections, each in parentheses. The first represents a function and the other two represent properties. The determiner function takes a property (the first of the two) as its argument and gives as its value another function. That second function takes the second property as its argument and gives as its value a truth value, i.e., the truth value of the proposition.

⁵ The notation used here is kind of a hybrid of two different notations. I follow Keenan [26] in representing the content of an Aristotelian I-type sentence as '**(THERE IS)(P1)(P2)**', and I follow Braun [7] in using an underline to represent a gap in this content. For more on the latter convention, see Braun [7], footnote 6, as well as Salmon [41], footnote 53.

The initial place in the proposition is empty, since Doris’s use of the expression ‘that polar bear’ is non-referring in this context. But the sentence she utters has a certain structure and linguistic meaning to it, which corresponds naturally to the structure of a singular proposition. So we may hold that this structure, along with whatever pieces might be available, will show up at the level of semantic content. We essentially have the structure of a regular singular proposition, and we also have the property *being handsome* in its appropriate place.

I’m not going to argue for this particular view of gappy propositions, defended by the likes of David Braun [7] and Nathan Salmon [41], at least with respect to empty names. I will, however, discuss in more detail some of the defenses Braun gives against certain kinds of objections that will also apply in kind to my view about the gappy propositions expressed by certain existentially quantified claims.

What we say about situations such as Nick’s utterance of (2) in the ordinary setting parallels this view about what happens with non-referring demonstratives. In both cases, we have the intentions of a speaker doing some work in helping to fix the reference of one of their expressions. For Doris, it’s ‘that polar bear’; for Nick, it’s ‘there are’. But in both cases, the speakers intend to express something that, at least at the level of semantic content, they cannot express—there is no such content available to them. Doris’s polar bear does not exist, and neither does a domain that includes superheroes. In both cases, the speaker expresses something linguistically meaningful, using normal English words that all competent speakers understand, so there is no confusion caused by the fact that the proposition expressed is missing something.

A natural thing to worry about with gappy propositions is whether they come out true or false. My view is that a sentence like (2), uttered by Nick in an ordinary situation like the one described above, lacks a truth value. The gappy proposition expressed by Nick in this case cannot have a truth value, since any truth value a quantified claim like this one would have is the result of a function operating on a property. But the relevant function is absent here, so there is no operation that could produce a truth value.

such as Jeffrey King [28] may see as being a quantificational expression. We could get the same kind of failed reference in examples with simple demonstratives, with a little more setup to the example.

This view parts ways with David Braun somewhat. He claims that gappy propositions are all false, though his inquiry is about singular propositions expressed by sentences with empty names in them.⁷ The fact that his view is about singular propositions and mine is about general, quantified propositions may be enough to justify a difference in views about truth values. My reason for saying that gappy quantified propositions lack a truth value is primarily due to the functional aspect of them, which is not present in Braun’s examples. But also, Braun is not actually firmly committed to the view that they are false—he argues that it does not matter all that much whether they are false or lack a truth value.⁸

The view I am presenting is an error theory, positing a relatively widespread error among competent speakers of English. Certainly Nick takes himself to be saying something true when he asserts (2), even if he doesn’t actually care about the existence of superheroes. But if we point out to him the tension between his assertion of (2) and his belief that superheroes do not exist, and ask him to reflect on (2), he will realize as well that there was something defective about his claim. I believe my view offers an intuitive explanation of that defect.⁹

When people like Nick are in ordinary settings and make claims like (2), they intend to quantify over superheroes and attribute some properties to them. They intend to help themselves to a domain that includes superheroes, and so includes nonexistent objects. This is their intention, even if they are not successful in what they set out to do. We know it is their intention because of what they say, how they say it, and how they respond if questioned about what they say. Alternatively, if someone says, “There exist superheroes who can fly,” they are displaying their intention to limit themselves to the group of existing objects by using the word ‘exists’.¹⁰

⁷ This claim is restricted to positive singular propositions—negated gappy propositions may come out true on his view. See Braun [7], pp. 604 ff.

⁸ See Braun [7], pp. 607 ff.

⁹ Notice that this error theoretic aspect of my view draws a contrast between it and a standard contextualist view about ‘knows’. On a contextualist view such as Stewart Cohen’s, there is no problem with reference failure—the epistemic standards of the speaker of a sentence like, “S knows that P,” fixes the semantic value of the knowledge claim. So in a low-standards context, an utterance of, “S knows that P,” may express something true, even though the same sentence may express something false if uttered by the same person in a high-standards context. See, for example, Cohen [12] for details of his view, and I will also pick this topic up again in Chapter 5.

¹⁰ Other ways we may linguistically make our intentions clear is to modify our claims with expressions such as ‘really’, ‘actually’, ‘in our minds’, and so on. Such modifications may alter the context of the

Recall that the data discussed in Chapter 1 about ontological robustness, data that motivates this entire view, is about the *appearance* of truth and falsity, *viz.* the appearance of truth in one context and falsity in another, with respect to the same claim. We don't begin with the data that claims like (2) are true, but rather that they seem *prima facie* to be true. We hear them as true, and we assert them and treat them as if they are true, even if we might be mistaken about that.

My demonstrative view has a nice, elegant answer to why a claim like (2) appears true in ordinary contexts and false in ontological ones. This answer is that in ordinary contexts, we intend to refer to a quantifier meaning whose domain includes nonexistent objects like superheroes. But then, when we shift to discussing ontology, we now intend to refer to a quantifier meaning whose domain includes only objects that exist. This explains our inclination to see (2) as true in some settings and false in others, regardless of whether (2) actually is ever true, or even whether (2) ever actually expresses a complete, non-gappy proposition.

It should not be a surprise that our intuitions about the truth-value of some of our ordinary claims may be misguided or mistaken. Of course it is always a goal of a speaker, at least when making sincere assertions, to say something true. But propositions and truth-values are mysterious enough for linguists and philosophers to have serious disputes about their workings and natures, so there's no reason to expect ordinary speakers, however competent, to get everything right in this respect. This excuse is not a catch-all that allows us to explain away all problems, but I do believe it's appropriate here.

Suppose we asked Nick whether he thought he was saying something true to his daughter Mirah when he uttered (2) in his conversation to her. His initial response would be something like, "Sure, I suppose so." Truth would have been a goal of his at some level, but not the primary concern. If we pressed him on the matter—the way we might if we were interested in Quine's ontological commitment process—he may retract his claim. But he may say other things that lead in different directions, too. He may maintain that superheroes exist in our minds, or that they exist as stories. He may say things that will line up with any number of different ontological views, and he may

discussion or the target quantifier meaning, though some of them may also be idiomatic expressions that serve simply as emphasis of an assertion that need not change the context.

have no confident beliefs about any of them. But one thing that will not change is that, faced with a similar situation, he would be happy to say (2) again.

In light of this, it seems reasonable to say that when it comes to Nick's intentions that go along with his assertion of (2), there may be some complications. He may have all kinds of intentions—the intention to inform his daughter, to placate her, to say something true, etc. His *linguistic* intentions are all I'm concerned with—what is it that he intended to say? Even here, it's clearly complicated. But it should be reasonable to say that the words Nick chooses display some linguistic intentions on their sleeves. If he uses an existential quantifier expression in a way that quantifies over superheroes, then there is at least a *prima facie* intention to quantify over superheroes that we cannot deny. This intention gives us a starting point—we look at the words used and begin with the most straightforward and minimal way of interpreting it. And this starting point helps us explain why we have the intuitions about them that we have, *viz.* the intuition that we're saying something true (even if we're not). We have some kind of intention to quantify over superheroes, and we make a claim that would be a true one if superheroes did exist, so the claim has an appropriate kind of ring of truth to it.

3.3.3 Believing Gappy Propositions

Another kind of explanation we might offer about why we take our quantified claims to be true even when they do not actually have a truth value is that we *believe* them. I will offer a brief explanation of this path, though I do not want to get too side-tracked from the primary issue. Since I have been accepting much of David Braun's machinery of gappy propositions and the ability to believe gappy propositions is central to his view, it seems useful to discuss this issue.¹¹

Continuing with the superhero example, we may say that Nick believes the proposition represented by (5), even though it's a gappy proposition. He holds this belief, in part, because he does not realize that it's gappy. The content of his belief has a hole in it, but he's not aware of this, so he believes it anyway. This explanation fits with the fact that, as we can see by the entire context and the other things Nick says to Mirah, he must be taking himself to be quantifying over nonexistent objects, helping himself

¹¹ Braun [7], p. 608, discusses the belief of gappy propositions in the context of his discussion on empty names.

your believing that Aristotle was a philosopher, but which underlies your doppelgänger's believing that Twin Aristotle was a philosopher. You and Twin You differ in what you believe, despite the intrinsic similarity in your mental states, because your mental states stand in different causal relations to Aristotle and Twin Aristotle. These causal relations involve utterances of proper names such as 'Aristotle': Your utterances of the name 'Aristotle' refer to Aristotle, whereas your Twin's utterances refer to Twin Aristotle.¹²

On this kind of view, we see how we might make sense of someone believing a gappy proposition. If we accept Braun's account here, we say that the two twins are in the same mental state with respect to their own world's propositions about Aristotle and Twin Aristotle, even though these are different propositions. The intrinsic state they are in, then, must not rely crucially on the constituents of the proposition itself, at least with respect to the actual identity of Aristotle. If we can change the Aristotle who populates the proposition without changing the intrinsic state of the believer, then some aspect of Aristotle does not matter to the intrinsic state. Looking back at the Napoleon case, too, we see that sometimes our beliefs are formed due to misinformation, but the internal state we are in may be the same, regardless of whether Napoleon himself even shows up in the proposition at all. So while the content of our belief may be a proposition, the intrinsic state we are in is a different animal from the proposition, and gappy-ness does not necessarily affect it.

If we accept this account that makes sense of what it might be to believe gappy propositions, then we can see a way to placate the worry about our intuitions that Nick is saying something true when he utters (2) in an ordinary context. Our intuitions are rooted in the fact that speakers like Nick may actually believe (2). It is neither incoherent nor even all that strange for us to have beliefs that purport to be about nonexistent objects. Now, when pressed, we may choose to give up that belief, in which case we may also retract our previous assertion. But that does not mean that we didn't believe it earlier. In fact, such retractions lend some amount of support to the view that 'there is' claims are context-sensitive: we happily assert them in ordinary settings, backed up by beliefs we may be expressing, but when we are pushed to more critical

¹² Braun [7], p. 601, emphasis in original.

(possibly philosophical) settings, where our beliefs are changed or more refined, we may not be willing to assert the same claim anymore. One important part that has changed, I maintain, is the domain of objects that we’re attempting to help ourselves to from context to context.

3.3.4 Reference Magnetism and Speaker Intentions

Why are Nick’s intentions so important? You may agree that Nick intends to quantify over nonexistent objects when he utters (2) in an ordinary setting, but you may also think that his intentions are not the most powerful thing in play here. One way this line of objection could run is to follow Ted Sider’s use of David Lewis’s notion of ‘reference magnetism’.

This discussion begins with a disagreement between Ted Sider and Eli Hirsch about the consequences of accepting some notion of Quantifier Variance.¹³ Quantifier Variance is essentially the view that there are many acceptable candidate meanings (i.e., quantifiers) for our existential quantifier expressions. Hirsch thinks that when, e.g., David Lewis says, “There are tables”, and Peter van Inwagen says, “No, there are not”, the best way to make sense of this dispute is to attribute to each philosopher a different quantifier meaning for their utterances of ‘there are’. Hirsch takes charity considerations to push us towards making a strong effort to interpret people as speaking truly, if we can, and attributing quantifiers with different domains (Lewis’s larger than van Inwagen’s) will do just that.

Sider’s response is to admit that these other quantifiers are proper candidates for the meaning of ‘there are’ in their sentences, but also to claim that charity considerations don’t trump the ‘reference magnetism’ of the best candidate meaning, *viz.* the quantifier with the domain of objects that exist. Sider thinks Lewis and van Inwagen’s existential quantifier expressions refer to the same quantifier: the plain, unrestricted existential quantifier. The notion of ‘reference magnetism’ is meant to parallel Lewis’s views about naturalness in his explanations for how our predicates successfully refer to certain (more natural) properties rather than other (less natural) ones.¹⁴ Sider extends Lewis’s notion of ‘naturalness’ to logical notions. He claims that the unrestricted existential quantifier

¹³ See Sider [43] and Hirsch [18] for full discussions of this matter.

¹⁴ See Lewis [29], especially pp. 13 ff.

‘carves the world at its logical joints’ better than any other quantifier with a restricted or possibly gerrymandered domain.

The notion of reference magnetism itself is fairly plausible—or at least easy enough to understand—in many cases. It is rooted in a particular view of meaning that Ted Sider and David Lewis both accept. The following two passages will explain and illustrate this general view. First, Ted Sider gives a general statement about this theory of meaning:

On this view, meaning is determined by two sources. Nearly everyone agrees that *we* play a role. Perhaps this role is played by convention: a complex pattern of behavior and dispositions of the speaker and her language community. Or perhaps a more individualistic account focused solely on the speaker is correct. The second source is the controversial one: of the many candidate meanings (whatever their ontological status), not all are created equal. Some are intrinsically more eligible to be meant than others. This eligibility is starkly metaphysical in nature: some candidate meanings ‘carve nature at the joints’ more than others, and it is part of the nature of reference and meaning that candidates that carve nature at its joints are more eligible to be meant. The meaning of a word, then, is the best candidate, where strength of candidacy is based on (1) fit with meaning-determining facts about the speaker or her linguistic community, and (2) intrinsic eligibility on the part of the candidate. For short, meaning supervenes on *use* plus *eligibility*.¹⁵

And next, we have an example from David Lewis:

The thesis that indeterminacy in nature is multiplicity fits well with the thesis that indeterminacy in language is semantic indecision: there are many clouds, and we haven’t decided just which one to call ‘the cloud’. Nevertheless, the two are separable. Maybe instead ‘the cloud’ refers to the entire multiplicity; or maybe just one of the many—but it is a secret which one—is a mighty reference magnet, and so ‘the cloud’ refers to that one.¹⁶

¹⁵ Sider [42], p. xxi.

¹⁶ Lewis [32], footnote 2.

The idea here is that if it turns out that there really is one privileged collection of particles that constitutes an object that would most properly be called ‘the cloud’, then when we utter the phrase ‘the cloud’, it may turn out that this privileged object will magnetically draw the reference of our utterance to it. It may be that if we only had better powers of observation or better understanding of clouds, we would realize that, even though clouds get kind of hazy around the edges, there really are sharp boundaries that demarcate the cloud from the air around it. And so it may be that there is just one well-defined object floating through the sky, changing its shape but remaining that one object. So, whether we know it or not, when we point to it and refer to it as ‘the cloud’, we’re really referring just to that one object.

This notion may get extended to solve problems such as how it’s possible, and even likely, that we are talking about the same things when we use the same terms. Given that our introductions to words like ‘tiger’ may be radically different from each other, we may have all kinds of different ideas about tigers, maybe even conflicting thoughts about some of the details about the animals. But if I say, “Tigers are ferocious,” and you say, “No, tigers aren’t ferocious—tenacious maybe, but not ferocious,” we’d like to think that it’s pretty obvious that we’re talking about the same thing. So this view could say that the property of being a tiger is a very natural property, and thus a ‘mighty reference magnet’. This can help explain why, among all the (complex, gerrymandered) properties we could possibly be referring to when we use the word ‘tiger’, we both end up referring to the same property.

Now, as mentioned above, Lewis’s view is about *properties*, specifically about how some properties carve nature at its joints closer than others. Sider’s extension of this view is to *quantifier meanings*, claiming that, since these are ‘logical’ entities, it can make sense to say that some of them might carve nature at its ‘logical’ joints more closely than others. To put this in terms that fit with the discussion we’ve had so far (which, I believe, is a fair characterization of Sider’s [43] view), Sider says that the domain of objects that exist is the most logically natural one. Sider also thinks that there are no nonexistent objects, so the domain of all the objects that exist would be the largest one available to us. So when we use a quantifier, he maintains that the unrestricted domain that will magnetically draw our reference is the domain of everything that exists.

It’s unclear exactly what Sider would say about Nick’s assertion of (2). In the case

of two philosophers arguing about ontology, Sider's claim is that they both mean the same thing when they use existential quantifier expressions like 'there is' or 'there exists' (which Sider follows Lewis in not distinguishing between). It's reasonable in this case to think that, since both philosophers are arguing about which objects exist, their use of 'there is' will be associated with the unrestricted domain of all the objects that exist. In this case I agree with Sider, since, in accordance with my contextualist view, if we are in a context where the explicit task at hand is to discuss existence and which objects exist, then not only will our use of 'there exists' refer to a quantifier meaning with the domain of all objects that exist (as it always does, I maintain), but so will our use of 'there is', since it adapts to the intended domain of the speaker in the context.

But what about Nick's assertion in the ordinary setting? In addition to his assertion of (2), recall that he also asserted, nearly in the same breath, that superheroes do not exist. Is it fair, then, to think that, against his own will and intentions, the domain of everything that exists magnetically draws the reference of Nick's use of 'there are'? I think that we should say that Nick's intentions trump the reference magnetism of the domain of objects that exist here. Recall Sider's own view of meaning in the passage quoted above. His claim is that the content of an expression involves a balanced effort between the intent of the speaker (or community) and the fitness of the candidate meanings. It doesn't follow that the most elegant candidate meanings will always be the ones to which our expressions refer—if intentions play a role here, they must be able to trump reference magnetism at times. It seems reasonable to think that this is one of those times, especially when Nick explicitly states that the objects he's quantifying over do not exist.

I am happy to accept the general picture of quantifier meaning given by Sider. But the picture as it works out in this case of Nick in the ordinary setting appears to be as follows: Nick has a clear intention to quantify over superheroes, displayed by the claims he makes. But Nick also has an explicit belief that superheroes do not exist. The domain of all objects that exist, being the largest domain available for reference (and also quite 'logically natural', if we follow Sider there), will have some amount of magnetic pull.

We can anthropomorphize this situation (perhaps a bit ridiculously) and imagine the domain of objects that exist having ears that perk up at the use of an existential

quantifier expression. It then reaches up with its magnetic fingers, trying to latch on to Nick's use of 'there is' and pull it down to itself. But in this case, Nick's intentions are not interested in this particular domain, since he doesn't believe superheroes exist, so they swat away the outstretched limbs of the domain of existing objects and then reach out further for a larger domain that might include nonexistent objects. Unfortunately, they are grasping at nonexistent straws, and they won't be able to find anything beyond the domain they just rebuffed. But to give in to the natural domain of what exists would mean that Nick said something false in this case—since there are no superheroes in the domain of existing objects, there will be none in there who are both animal and human. His intentions will reach out for a domain that will at least give his claim the potential of saying something true, even if it ends up having to settle for a gappy proposition that has no truth value.

Before moving on to the next chapter, I'll briefly tie together the moving pieces of this section. My goal here has been to fill out some of the details of my view, specifically about the issue of the domains associated with quantifier meanings. Since I claim that 'there is' is context-sensitive in a way that results in a certain kind of semantic failure in some contexts, this clarification seems necessary. As Azzouni points out, we can feel a kind of desperation kick in when we realize that we think we're talking about something, but really we're talking about nothing.

The way I have dealt with these issues in this section are as follows: first, my contextualist view of 'there is' is modeled on the context-sensitivity of demonstratives; second, what fixes a domain in a context where a 'there is' claim is uttered are the intentions of the speaker; third, reference failure with existentially quantified claims, when it happens, looks relevantly similar to reference failure with demonstratives; fourth, the machinery of gappy propositions gives us a tool to explain the content of the problematic existentially quantified claims; fifth, linguistic meaning allows us a way of saying how these claims are meaningful, understood, etc.; sixth, the gappy propositions in question are truth-value-less, but this fact is opaque to ordinary speakers who (reasonably enough) take themselves to be saying something true; seventh, at least on some views about beliefs, it may be reasonable to see speakers as believing gappy propositions too; eighth, the phenomenon of reference magnetism may add to our picture of how existential quantifier expressions latch on to domains; ninth, reference magnetism reinforces

the thought that ‘there exists’ is invariant, since the domain of existing objects will magnetically attract it; and tenth, even with the reference magnetism picture, we can still make sense of how the speaker’s intentions play a role in breaking the magnetic attraction of the domain of existing objects in a case where the speaker clearly wants to quantify over nonexistent objects.

I believe these ten moving pieces fit together in a relatively harmonious way to give us a picture of one small part of our discourse about nonexistent objects. In ordinary settings, we want to talk about things that we ourselves don’t think exist, and this is inherently confusing. But we go ahead with it anyway and let the chips fall where they may. These problems arise with existentially quantified claims, just as they do with empty names and failed demonstratives. My primary view I am presenting and defending in this dissertation is the contextualist view developed in Chapter 2, but with that as my central view, I also must pay heed to its implications and say how we can make sense of our ordinary discourse about nonexistent objects in a way consistent with the view. That has been the task here. Now, moving forward to the next chapter, I will discuss the main competitor to my contextualist view and give a holistic argument that my view fares better on balance than the competitor.

Chapter 4

Contextualism Over Orthodoxy

4.1 Introduction

First, let's briefly recap the previous three chapters. In Chapter 1, I introduced the notion of 'ontological robustness', which is essentially a tool to help us think about the ontological thick- or thinness of a claim. Going through several examples, I argued that 'there is' claims sometimes are and sometimes are not ontologically robust, depending on the context in which they're uttered. Additionally, I argued that 'there exist' claims are always ontologically robust, regardless of context. In Chapter 2, I presented my contextualist view of 'there is', discussing the semantics of quantifiers and the demonstrative-like features of 'there is'. In Chapter 3, I addressed a few issues that were raised earlier, elaborating on them and addressing some potential objections to my view, centering on issues about the domain associated with the existential quantifier.

Next, in this chapter, I will discuss the primary opposition to my view, where it comes from and in what ways I disagree with it. In doing this, I will be establishing the main argument of this dissertation. Up until now, I have primarily been presenting and defending my view about the context sensitivity of 'there is'. A companion view that I've discussed some and argued for a little is the view that 'there exists' is *not* context sensitive. One of the main purposes of articulating these views, however, is to build an argument against the orthodox view about the meaning of existential quantifier expressions. The orthodox view says that, as David Lewis puts it, all of the existential

quantifier expressions are “entirely synonymous and interchangeable”.¹ A result of my view that I’ve been presenting and defending is that this claim of orthodoxy is false. If ‘there is’ is a context-sensitive expression and ‘there exists’ is not, then they are neither synonymous nor (entirely) interchangeable.

The current form of this orthodox view is rooted in Quine, so I will spend a little time here discussing his statement and explanation of his view. His view appears to be at least tacitly accepted by a majority of contemporary metaphysicians, so there are not a great deal of variations on the orthodox theme. There is a little bit of the logical space of competing views that has been explored by other philosophers in the past several decades, so I’ll discuss some of that here as well. But since the Quinian view is the primary competitor, it will be the main focus of this chapter. Once we have the details of this view on the table, I will then give a new argument for my contextualist view, on the grounds that it gives as good as or better of an account of the data that needs to be accounted for by a view like these than the Quinian view. This will be a holistic argument, comparing two general accounts of the meaning of the existential quantifier expressions ‘there is’ and ‘there exists’ and weighing their relative merits.

4.2 Quine’s View of Existential Quantifier Expressions

The roots of the orthodox view about existential quantifier expressions appear to be in Quine’s metaontological views. Peter van Inwagen gives a nice summary of five tenets of Quinian metaontology as follows:

1. Being is not an activity
2. Being is the same as existence
3. Being is univocal
4. The single sense of being or existence is adequately captured by the existential quantifier of formal logic
5. General strategy for answering the ontological question, i.e., the criterion of ontological commitment²

¹ Lewis [30], p. 24.

² These are the headings of the sections of van Inwagen [47], except for 5, which isn’t given a heading.

Since my focus is on the meaning of the expressions ‘there is’ and ‘there exists’, the second and third tenets are the most relevant ones to this project. The first is an Austinian point that is not a claim about language or quantifiers; it’s about whether existence or being is something we do or not. The way J. L. Austin expressed it was that we should not think of existence as merely, “like breathing, only quieter.”³

The fourth is primarily a point about regimentation and the purpose of the existential quantifier symbol ‘ \exists ’. I don’t have any problems with translation into formal logic itself, though this regimentation can sneak in or rely on some of the other tenets. That is a problem with the other tenets, though, not the translation one.

The fifth tenet is one that I discussed some in Chapter 1, when introducing and using the notion of ontological robustness. There I discussed how, though evaluating a claim for its ontological robustness status does bear some similarity to the process of ferreting out our ontological commitments, it is not quite the same as the process suggested by Quine’s criterion. The ontological robustness status of a claim indicates our current ontological intentions, regardless of whether we actually have views about ontology. Quine’s method could be seen as a sort of ontological therapy, where we can dredge up our deep dark secrets (say, a latent nominalism or a suppressed idealism) and examine them in the light of day, where we’re forced to come to terms with or reject them. I don’t have any issues with his criterion, especially since it is in most respects orthogonal to my interests.

The two tenets that I’m interested in, the second and third, combine to form the view that I’m arguing against, though not completely transparently. By ‘being is the same as existence’, I take Quine to be saying that ‘there is’ and ‘there exists’ mean essentially the same thing. And by ‘being is univocal’, I take Quine to be saying not just that ‘there is’ is not ambiguous (which he is saying), but also that the expression is not context-sensitive. Neither of these amendments or interpretations are explicit in van Inwagen’s summary of Quine’s metaontology, but I think it is fair to attribute them to Quine as well as to the general view I take to be the orthodox view, found in Lewis and others.

To discuss Quine directly first, I’ll focus on two articles of his where he discusses the relevant aspects of his metaontological view. The first is “On What There Is” (1948),

³ Austin [2], p. 68n

and the second is “Existence and Quantification” (1960), and I’ll look at them each briefly, in turn.

4.2.1 “On What There Is”

Quine [36] was an influential article that helped reinvigorate metaphysics in the wake of the bad taste for it left in everyone’s mouths by the logical positivists. In terms of tone and style, this article was one of the shapers of the orthodox practice of metaphysics for the past six decades. It is not the main thrust of the article, but one thing that Quine does here is to articulate some of his metaontological assumptions. The most relevant bit of this for our current interests comes in the following footnote:

The impulse to distinguish terminologically between existence as applied to objects actualized somewhere in space-time and existence (or subsistence or being) as applied to other entities arises, in part, perhaps, from an idea that the observation of nature is relevant only to questions of existence of the first kind. But this idea is readily refuted by counter-instances such as ‘the ratio of the number of centaurs to the number of unicorns’. If there were such a ratio, it would be an abstract entity, viz. a number. Yet it is only by studying nature that we conclude that the number of centaurs and the number of unicorns are both 0 and hence that there is no such ratio.⁴

For the sake of brevity, let’s give the two senses of existence he discusses here the following names:

E1 existence as applied to objects actualized somewhere in space and time

E2 existence (or subsistence or being) as applied to other entities

Quine does not himself think that these are two legitimate senses of existence. He is attributing this to his opponent, which here is the Meinongian, who holds that some objects have being even if they don’t exist. Quine is also being a little unclear here about whether the distinction in question is between the two properties *existence* and *being*, or between words associated with these properties, such as ‘is’ and ‘exists’. He

⁴ Quine [36], p. 3, n. 1

does say that this impulse is to distinguish *terminologically* between the two senses, which makes it sound like his focus is really on expressions like ‘is’ and ‘exists’. He may be intending both, but since my focus here is just on the meaning of these words, I’ll limit my discussion to that aspect. It’s unclear whether Quine is really intending to give an argument for his position, but if he is, it would appear to be the following:

P1 The motivation for distinguishing between **E1** and **E2** is that observing nature only tells us about what exists in the **E1** sense.

P2 But observing nature can also tell us about what exists in the **E2** sense (e.g., ratio of number of centaurs to number of unicorns).

C So the distinction between **E1** and **E2** is unmotivated.

Now this argument could be read in a way that connects with either of the two tenets of Quinian metaontology that I’m interested in. On the one hand, this could be about Quine’s denial that there is a difference between ‘is’ and ‘exists’; on the other hand, this could be about Quine’s claim that ‘existence’ is univocal, so there is not more than one sense of the word. I read Quine as intending his claims to do double-duty, based mainly on the way he introduces **E2** as “existence (or subsistence or being)”. That seems to indicate that we might read this as an argument about two senses of existence, or, if someone wants to say this other sense of existence is captured by other words like ‘subsistence’ or ‘being’, the argument will apply to them as well.

The argument itself is, indirectly, an application of Leibniz’s Law. Quine is basically saying that the Meinongian thinks there’s a property that **E1** has but **E2** lacks (a property about the fruitfulness of observing nature). And since they have different properties, they must be different meanings. But Quine claims (in his second premise) that **E2** also has the property in question. So we can’t use Leibniz’s Law and this property about observing nature to drive a wedge between **E1** and **E2**. And, if that’s the only property that could plausibly distinguish between **E1** and **E2**, then Leibniz’s Law will tell us that they really are the same meanings.

This argument is a bit confusing, but there are a couple ways to approach it that will undercut its effectiveness in relation to my own interests. If our focus is on why we might want to distinguish *terminologically* between existence and being, then Quine presumably means to attack a difference in meaning between expressions like ‘exists’

and ‘is’. If this is the case, then **P1** must be false, since we have other motivations for why we may want to distinguish between them aside from issues about observing nature. The case of discourse about fiction provides such an example—we may believe that superheroes don’t exist, but we may still want to make claims that purport to quantify over superheroes. That on its own is a motivation for distinguishing between ‘exists’ and ‘is’, or ‘there exists’ and ‘there is’.

If Quine’s focus is more on the *properties* involved, the worry about distinguishing between *existence* and *being* in a metaphysical discussion, then this argument is not in the realm of what I’m interested in. I am not arguing for a Meinongian ontology, just a point about the meaning of existential quantifier expressions. These two issues can be conflated too easily, as they appear to be (at least somewhat) in this passage from Quine.

I should also note that if Quine’s points are about language—as they should be, given the inclusion of ‘terminologically’—but only intended to be about how we might talk during a philosophical discussion, then I don’t have grounds to disagree with him here either. On my view, since ‘there is’ is context-sensitive, if we are having an explicit discussion of ontology, ‘there is’ will take on the same meaning as ‘there exists’, since the speaker will be intending to help herself only to the domain of objects that exist. So in a discussion of ontology, it is reasonable to resist someone’s attempts to make a distinction between ‘exists’ and ‘is’ (unless the discussion is about Meinongianism, say, and it is part of the issue at hand in that discussion).

4.2.2 “Existence and Quantification”

In Quine [39], he is responding to Carnap’s worries that we cannot meaningfully ask existence questions that make use of Allwörter, or category words. Quine has two responses to this. One is that we have no principled way to determine what is a category word and what isn’t. The other response is basically the intuition that we have to be able to make sense of what Carnap would consider external or categorical existence questions or claims in order to make sense of what people are doing when they try to get out of being committed to various things.⁵ It seems reasonable to think that someone who goes through a lot of trouble creating a system of paraphrasing so they don’t have to

⁵ Quine [39], pp. 92–93

actually use, say, number terms, has a belief that numbers categorically do not exist. So basically, Quine thinks that ontological disputes or questions are intuitively substantive, and he thinks that intuition counts against Carnap's view.

My current interest in Quine [39] is his discussion of and argument for the views that 'there is' and 'there exists' mean the same thing and that these expressions are univocal.⁶ Quine begins this discussion by entertaining the notion that there are two kinds of existence, a philosophical and a commonsensical sense. The proponent of this kind of view may be someone like an immaterialist, who thinks that objects like rabbits don't exist in the philosophical sense, but they do in the commonsensical sense. Basically, what this would mean is that they think it's perfectly fine to talk about rabbits and quantify over them in ordinary talk, but when it comes to the philosophical room, the immaterialist no longer allows quantification over rabbits. When things get serious, they claim that only "qualities, times, and places"⁷ exist, no rabbits or other ordinary objects.

The target of Quine's view is a sort of ambiguity view that is not something all immaterialists would be committed to; it's just one option for someone who wants to deny the existence of rabbits, say, but still wants to say we can make true claims where we're quantifying over rabbits. We could do this if our quantifier expressions were ambiguous in the right way, so the sentence, "There is a rabbit in the yard," could be read as having two legitimately different meanings.

Quine is somewhat comfortable with the immaterialist going on as if there is an ambiguity here. This is because the two different meanings would not really both be quantificational meanings. That is, on the common sense reading of the claim, the immaterialist would have to say that there isn't really quantification going on—the meaning of 'there is' there would have to be serving some other function.⁸ So Quine

⁶ Leading up to this argument, Quine discusses a few other of the claims that van Inwagen points out as tenets of Quinian metaontology, discussed above. For example, he says, "Existence is what existential quantification expresses. There are things of kind F if and only if $(\exists x)Fx$." (p. 97) In this context, Quine doesn't see treat this as a substantial thesis—he is basically just pointing out how we generally explain the symbolic notation.

⁷ Quine [39], p. 98

⁸ Quine is not too clear about how this would work, but that's forgivable, since it's not his view. The main thing he says about it is that for the immaterialist, the common sense reading would have the same "stimulus meaning" for them as for us (98), but not the same meaning with respect to quantification. A view like Hofweber [20] could likely fill out this sort of distinction in a way that would maybe work for the kind of immaterialist Quine has in mind.

is not worried about there being two different kinds of quantification. He doesn't think there's an ambiguity here at all, but even if there were, he thinks, it wouldn't be between two different quantifier meanings. It would be between two different meanings of the quantifier expressions, one of which was a quantifier meaning and the other of which was something else, something distinct from quantification.

It is important to note here that what Quine is after is not particular to the issue of materialism vs. immaterialism. These comments are general claims about the meaning of quantifier expressions, and he means to apply them across different debates. The other example he brings up is that of nominalism about numbers, where people who don't believe in the existence of numbers may still want to talk as if they exist when they're doing math, but then be more careful in the philosophy room.

The discussion about the immaterialist and the nominalist is Quine's setup for the argument he gives in the following passage:

It has been fairly common in philosophy early and late to distinguish between being, as the broadest concept, and existence, as narrower. This is no distinction of mine; I mean "exists" to cover all there is, and such of course is the force of the quantifier. For those who do make the distinction, the existent tends to be on the concrete or temporal side. Now there was perhaps a reminder of the distinction in the case of the rabbit and the immaterialist. At that point two senses of "there is," a common and a philosophical, threatened to diverge. Perhaps the divergence which that sort of case suggests has been one factor in making philosophers receptive to a distinction between existence and being. Anyway, it ought not to. For the point there was that the rabbit was not a value of the immaterialist's variables; thus existence, if this were the analogy, would not be a species of being. Moreover, we say that the sensible materiality of the rabbit was inessential to the example, since the prime numbers between 10 and 20 sustained much the same point.⁹

What Quine seems to be saying here is that the sort of motivation some (though not Quine) might make use of to posit an ambiguity could just as well motivate the same people in a different direction, towards distinguishing between existence and being.

⁹ Quine [39], p. 100

Notice that this move has a parallel in “On What There Is”, at least on my reading of his footnote argument there as being meant to do double duty, counting in favor of two separate tenets of his metaontology. Here as well he is lumping them together, saying essentially that the motivation someone might have for positing an ambiguity in ‘there exists’ could also lead one to think that ‘there is’ and ‘there exists’ have different meanings. His aim is to undercut this motivation for both purposes.

Quine appears to be saying here that if we were to distinguish between existence and being (which he himself does not want to do), the natural view we might choose is to say that being is more encompassing than existence. That is, the things that exist are a subclass of the things that are, a subclass such as the one of just concrete, present objects. The specifics of which subclass is the right one (i.e., do they really have to be present?) do not matter to our current worry—it’s enough to say that the existent group would be smaller in some natural or intuitive way than the group of all things that have being. This would capture the immaterialist and nominalist worries, as well as many others. Given this setup, I see Quine’s argument for his position in this passage as follows:

- P1** If ‘there is’ and ‘there exists’ picked out different quantifier meanings, the difference between them would be that the domain of the latter would be smaller than the former in some natural way (e.g., including only present concreta).
- P2** But then people (philosophers especially) who think certain things don’t exist should deny ‘there exists’ claims while assenting to the corresponding ‘there is’ claims.
- P3** This isn’t what happens—when pressed to say what they really think (e.g., “in the philosophy room”), they don’t want the objects in question to be values of their variables, and they tend to deny both the ‘there exists’ and ‘there is’ claims.
- C1** So ‘there is’ and ‘there exists’ must not pick out different quantifier meanings.
- C2** So existence is not a species of being.

I take **P1** to come from Quine’s discussion of those people who want to make the sorts of distinctions that he himself does not want to. Whether we’re talking about an ambiguity view or a view that separates existence and being, he thinks that the split they’re positing must be along some natural or intuitive line, such as the line between

physical and non-physical objects or between present and non-present objects.¹⁰ The point is that **P1** is based both on what sort of distinction would be reasonable and also on what sort of distinction has been common with Quine's opponents like the Meinongians. I don't see **P1** as particularly problematic.

Quine doesn't explicitly say **P2**, but it seems reasonable to attribute it to him, since it seems to be needed for this argument. For Quine, the issue of whether some object is the value of a variable for someone is an issue of what we can get them to assent to when they're pressed on it. We can't just attribute objects to them based on what we think they believe—we have to ask them and see what sentences they would agree to when all the cards are on the table. And he thinks we should only hold them accountable for the objects they admit at that point.¹¹

Quine seems to think that the conclusion (that existence is not a species of being) follows from the fact that the objects in question would not be considered as values of variables for the people who don't think they exist, which is **P3** above. For example, in the rabbit case, Quine thinks that the immaterialist, when pressed to say what he really thinks is true, will want to paraphrase away any quantification over rabbits. That is, he won't want rabbits to be values of his variables.

From these premises, Quine thinks it follows that existence is not a species of being, but that claim, for him, seems to be just a specification of what I've given as **C1**, put in terms of the meaning of the existential quantifier expressions. Notice that **C1** is essentially the same as the Lewisian position that all of the existential quantifier expressions are 'entirely synonymous and interchangeable'.

The main problem with this argument is that even if people behave in the way Quine thinks they will (illustrated in **P2** and **P3** of the argument), there may be other explanations for their behavior. It may be true that when we're in the philosophy room, the immaterialist will decide to be careful and not quantify over rabbits using either 'there exists' or 'there is'. But, as Quine himself points out, outside of the philosophy room, the immaterialist goes around quantifying over rabbits in ordinary discourse, at

¹⁰ We might also include the split between actual and non-actual objects here, since it fits with the sort of characterization we're giving of Quine's opponents, even if Quine doesn't bring up modal worries here.

¹¹ This is that fifth tenet of Quine's metaontology, his criterion of ontological commitment. See, for example, Quine [38], pp. 242-43.

least using ‘there is’. This is not an insignificant point—we need to worry about why this should happen.

This is where my discussion in Chapter 1 about ontological robustness becomes important. Quine appears to think that the best way to decide what the existential quantifier expressions really mean is to look at how we use them once we’re engaged in the process of determining our ontological commitments. It seems to me that it’s better to look at the whole picture of our usage, when we’re in ordinary settings and when we’re in ontological settings. The activity of digging into our ontological commitments is a kind of technical activity, in the sense that we may end up using technical or simplified or sharpened versions of certain words. If we really want to know what they mean, we should look at both these technical uses and also the normal, everyday uses of competent speakers.

4.3 Clarifying Orthodoxy and My Denial

The orthodoxy that is rooted in Quine and continued in Lewis is not a view that is frequently explicitly revisited, argued for, explained, etc. It is a kind of tacit orthodoxy of assumption, accepted by ontologists as a default view. It is not accepted by *all* ontologists—people like Graham Priest in the Meinongian tradition fight against this orthodox view. I’ll discuss Priest more shortly, but for now, I will briefly summarize orthodoxy in general.

This summary will be brief, since David Lewis gives us a clear statement of the orthodoxy I’m denying. I’ve included this passage before, but I repeat it here to aid in clarity:

Routley sees himself as defying an established orthodoxy; and I am prepared to appoint myself spokesman for the orthodoxy he defies. [...] We of the establishment think that there is only one kind of quantification. The several idioms of what we call ‘existential’ quantification are entirely synonymous and interchangeable. It does not matter whether you say ‘Some things are donkeys’ or ‘There are donkeys’ or ‘Donkeys exist’—you mean exactly the same thing whichever way you say it.¹²

¹² Lewis [30], pp. 24–25

This passage contains two primary points that, taken together, constitute the orthodox view. Altering them slightly from Lewis’s wording, we get the following:

Orthodox View of Existential Quantifier Expressions:

- (O1) There is only one kind of existential quantification.
- (O2) Existential quantifier expressions are synonymous and interchangeable.

The first of these, (O1), matches roughly with the tenet of Quinian metaontology that van Inwagen expressed as “being is univocal”, while (O2) matches (closely enough) the tenet expressed as “being is the same as existence”.¹³ Assuming these links are reasonable enough, then it’s easy to see why we should be comfortable calling this the orthodox view of existential quantifier expressions. It is explicitly endorsed as such by Lewis, and it also matches up with the view endorsed by Quine and van Inwagen.

These three philosophers have clearly shaped the field of analytic metaphysics more than most others in the past six decades. Quine’s papers discussed above reinvigorated metaphysics in the mid-20th century, and Lewis’s *On the Plurality of Worlds* and van Inwagen’s *Material Beings* were two of the biggest shapers of the conversations in metaphysics in the latter part of the 20th century into the present. We can add to this list Ted Sider, saying in his influential book *Four-Dimensionalism: An Ontology of Persistence and Time*:

What could it mean to say that there is a single correct meaning for ‘ \exists ’ (beyond the triviality that just one meaning is *mine*)? This is a serious question for anyone who takes ontology seriously. I will answer by sketching a picture according to which there is a single privileged meaning for (unrestricted) quantificational expressions.¹⁴

Here Sider clearly aligns himself with the orthodox view, capturing (O1) and (O2) together in one breath by claiming that there is, “a single correct meaning for ‘ \exists ’.”

¹³ It may be worried here that Quine’s stance on synonymy prevents this match. If Quine does not believe there is a *synonymy* relation, as we can reasonably interpret him saying in Quine [37], then he would not agree that quantifier expressions like ‘there is’ and ‘there exists’ are synonymous. My response is: that’s just fine. Quine clearly thinks it’s acceptable to interchange the expressions without any substantive loss or alteration in meaning. That is close enough. If he backs off even further from the interchangeability claim, then he begins to agree more with me and less with Lewis and others. That would also be acceptable—I’d be happy to have Quine on my side.

¹⁴ Sider [42], p. xv.

The purpose of pointing out the stance as well as the relative influence of Quine, Lewis, van Inwagen, and Sider is to support the notion that this view expressed in **(O1)** and **(O2)** really is the orthodox, received view in metaphysics about the meaning of existential quantifier expressions. There are certainly unorthodox views out there, but with proponents like these, it is reasonable to say that this view is the primary or default view. And since it is the orthodox view, it gives us even more motivation to point out where it goes wrong, if indeed it does.

In response to the orthodox view, I have offered a contextualist view of ‘there is’, and I maintain that ‘there exists’ is invariant in its meaning. The last three chapters have been developing these two claims, and they can be summarized in the following way to see the contrast with the orthodox view:

Anti-Orthodox View of Existential Quantifier Expressions:

(A1) ‘There is’ is a context-sensitive expression.

(A2) ‘There exists’ is invariant in meaning.

These two claims, taken together, amount to a denial of **(O2)**, since an invariant expression and a context-sensitive expression, no matter how similar, cannot be entirely synonymous or interchangeable. And though I do not claim that there are different *kinds* of quantification in any robust sense, I disagree with **(O1)** insofar as it is offered in support of a view that there is one fixed meaning for all of the existential quantifier expressions. My claim of **(A1)** denies this, as do other views that claim existential quantifier expressions to be ambiguous.

In what follows, I will now do some holistic comparison between the orthodox view and my own contextualist view, determining how they fare relative to one another on the kinds of desiderata we would like satisfied by a view about the meaning of existential quantifier expressions.

4.4 Comparing Anti-Orthodoxy with Orthodoxy

What is it we really want out of a view about the meaning of ‘there is’ and ‘there exists’? Ideally, we would like the view to be compatible with most pre-theoretic intuitions we

have, while also providing elegant explanations of their semantics. We want a view to rest on the weakest, least-controversial assumptions possible.

My dissatisfaction with the orthodox view is that its proponents must explain away some of the data that I take to be the most important data for understanding the meaning of existential quantifier expressions. Much of this data is about how we use the expressions, both in ordinary and in philosophical settings. Words like ‘is’ and ‘exists’ are not technical terms—they are extremely basic expressions that lie at the core of the language, and any competent speaker must be able to use these expressions properly with ease. So focusing on the way these expressions actually get used is crucial to filling out a theory of their meaning.

In what follows, I’ll discuss some points of comparison between the orthodox view and my own. This will be a sort of scorekeeping exercise, since we will be discussing where each theory has its strengths and weaknesses. The claim I intend to establish here in this chapter is that, on balance, my unorthodox view fares better than the orthodox view. I will only be viewing these theories through the lenses of a small number of possible comparison points, but my goal is to choose the most relevant and illuminating comparisons. In the sections that follow, I will take up various desiderata in turn. These are:

1. Agreement With Usage—the correct view should account for obvious differences in how we use ‘there is’ and ‘there exists’
2. Agreement With Intentions—the correct view should provide an explanation that fits with what we intend to say when we use these expressions
3. Ontological Neutrality—the correct view about the meaning of these expressions should remain as neutral as possible about ontological matters
4. Simplicity—the correct view should be as simple as possible, avoiding needless complexity

With each of these, I’ll explain them and discuss the approaches available to orthodoxy and contextualism, comparing them as I go. At the end, I’ll summarize why I see contextualism as coming out on top in this comparison.

4.4.1 Agreement With Usage

The issue of the usage of existential quantifier expressions was taken up in Chapter 1, exploring the notion of ontological robustness. I revisit it here now that my view and the orthodox view have been more clearly developed, and we can say what each view can offer as an explanation of usage.

It's clear that we use 'there is' and 'there exists' differently. We rarely use 'there exists' in ordinary settings. When we do, the 'exists' generally appears to be present for the purpose of emphasis, such as emphasizing the stark reality of something, or the discovery of life somewhere, or something we may not expect to exist. We use 'there is' much more routinely, and the expression generally takes a backseat to the subject matter at hand, whether it's superheroes, swans, or anything else. But we certainly *can* use 'there is' more like 'there exists', emphasizing existence or subsistence or things like that. Sometimes we do that with verbal emphasis, sometimes by adding other modifying words like 'really', or sometimes just by letting the context make it clear what we're saying.

Anti-orthodoxy—my view that is a combination of **(A1)** and **(A2)** above—offers a relatively straightforward explanation of the variance in our usage of these expressions. If 'there exists' is invariant in meaning, associated explicitly with the domain of objects that exist, it makes sense that we will use it infrequently, saving it for those times when we actually want to accentuate existence. The variance we see in our usage of 'there is' can be explained quite well with the contextualist view I've offered. This variance was actually a primary motivation for the view in the first place, so this shouldn't be a surprise. Sometimes when we use 'there is', we're intending to quantify over nonexistent objects, and other times when we use it, we're in the middle of an ontological dispute. Recognizing this expression as context-sensitive allows us to see why this makes sense, why we can be competent speakers of English and still use 'there is' as a quantifier, regardless of whether we believe in the existence of the things we're purporting to quantify over.

The orthodox view, on the other hand, cannot explain the variance in 'there is' as straightforwardly. If 'there is' and 'there exists' are entirely synonymous and interchangeable, then we should be able to interchange them without altering the meaning of our claims in any situation, especially in flat-out assertions of quantified claims (leaving

aside any concerns about what happens when they are embedded inside belief reports and such things). The examples discussed in Chapter 1 show that we clearly cannot interchange these expressions in all normal situations.

My primary example throughout the previous chapters has been a sentence about superheroes. This is clearly a case where we cannot plug in ‘there exists’ for ‘there is’, since part of the example is that the speaker does not believe that superheroes exist. I use this example because discourse about fiction gives us the clearest case of a situation where we have beliefs that these objects don’t exist, but we try to quantify over them anyway.

Some may worry that I’m ignoring an entire literature about fictional discourse and the ontology of fiction. I am not ignoring it, but I have been bracketing it. There are puzzles about fictional discourse and views designed to solve those puzzles that are certainly relevant to the issues I’m interested in, but it’s not clear that they help the orthodox view in the discussion at hand. For example, one way to deal with quantified claims about fictional characters is to adopt the view that that these quantified claims are really elliptical for ‘in the fiction’ claims. That is, when Nick says, “There are superheroes who are animals but also people,” he is *really* saying, “In the fiction, there are superheroes who are animals but also people.” Or, instead of ‘the fiction’, the operator may be more specific, depending on what Nick is talking about. It may be ‘in the Marvel universe’ or ‘in comic books’.

This ellipses explanation only works for object-fictional claims, or those that are the kinds of claims that may make it sound like we’re pretending the fictional characters are real. Aside from the fact that there are complicated object-fictional claims that don’t lend themselves well to a single ‘in the fiction’-style operator (like those that cross fictions or genres in one single claim), there are also meta-fictional claims to deal with. These are claims that sound more like they belong in the realm of literary criticism, where speakers talk about fictional characters in a way that recognizes their fictional (or nonexistent) status. Nick’s claim may even be read as a meta-fictional claim, since he is speaking more to fact that we may find characters of a certain sort if we look in comic books, while at the same time pointing out that they don’t exist.

The problem here is that views like this elliptical one do not help explain why there is a clear difference between using ‘there is’ and ‘there exists’ in our quantified claims.

They do not make it any more likely that we will use ‘there exists’ when we’re purporting to quantify over things that we clearly do not believe exist. All this kind of view does is address the issue of why we may be happy to say such things in the first place, which is an issue my view comfortably explains as well.

Now, the proponent of orthodoxy may appeal instead to pragmatics to explain why we may choose ‘there is’ over ‘there exists’. A potential view would be to say that ‘there is’ and ‘there exists’ are synonymous, but they have some different associations external to their content that lead us to use one over the other. This view would be coherent, and it may say some things that are very similar to the account I’ve given of how ‘there is’ and ‘there exists’ differ, just placing these differences at the level of pragmatics instead of semantics. It may say that since the word ‘exists’ appears in ‘there exists’, people may choose not to use that since they know it may draw the listeners’ attention to existence more than they might like.

There are two problems with this line—one is that it gives up on the notion that ‘there is’ and ‘there exists’ are interchangeable, as Lewis maintained, and the other is that it must say something even less palatable about the truth of the claims than my own view does. I maintained that a claim like Nick’s expresses a gappy proposition that is neither true nor false. The orthodox view must maintain here that Nick’s claim is false (at least if taken to be meta-fictional). This seems worse than having no truth-value, when we’re discussing a claim that the speaker himself thinks is true.

4.4.2 Agreement With Our Intentions

In Chapter 3, we discussed the intentions of the speaker and the role they play in the contextualist theory of ‘is’. On this demonstrative view, intentions play an important meta-semantic role in fixing the meaning of the quantifier expression. Whether my particular view is correct or not, it seems clear that we do actually have the kinds of intentions I discussed, i.e., the intention to quantify over nonexistent objects. There are clearly times where we expressly believe that certain objects don’t exist, yet we are still happy to quantify over them using ‘there is’, but not ‘there exists’.

It may be obvious, but I’ll just point out the difference between this desideratum and the previous one. The point about usage has to do with how we actually use expressions like ‘there is’ and ‘there exists’—what we actually say and what we consider

to be acceptable or felicitous usage of them. The point at hand is about our intentions, so it is about our mental states or about what we want or hope to express during a conversation. The previous desideratum can be satisfied by a theory that adequately explains our usage, giving an account that fits with and makes sense of how we talk. This desideratum about intentions can be satisfied by a theory that adequately explains the fit between our language and our intentions—that is, what our language says about our intentions, and why we say what we do, given what we intend.

Both views clearly get most of what we want with respect to our intentions—when we use ‘there is’ or ‘there exists’, we intend to quantify over some objects. We intend to do what quantifiers allow us to do, described in detail in Chapter 2, when we discussed truth conditions and semantics. But the main part where orthodoxy and anti-orthodoxy diverge is when it comes to accounting for our intention to say something true when we assert existentially quantified claims.

Now, it’s part of my view that speakers who purport to quantify over objects that they themselves believe to be nonexistent are making a certain kind of mistake. They intend to quantify over nonexistent objects, and they are unsuccessful, since there aren’t any nonexistent objects. Quine’s own criterion of ontological commitment can help us see that they must have such an intention and that they are making a mistake—if pressed on his assertion and invited to be careful about ontological commitments, a speaker like Nick will likely retract his claim, realizing that he wasn’t entirely successful in what he intended to do. I say ‘entirely’, because he *was* successful in communicating what he wanted to communicate to his daughter Mirah. He failed to say something *true*, but he did say something *meaningful*, and he would say the same thing again in a similar conversation.

The fact that I’m attributing a mistake to ordinary speakers in this situation may seem like it works against my view, but it should be noted that orthodoxy must say something similar, and even less appetizing. If the problem stems from the fact that ordinary speakers intend to quantify over nonexistent objects, then the proponent of orthodoxy will also need to attribute an error to ordinary speakers. The error will be different—instead of the speakers expressing a gappy proposition (as they do on my view), they would be expressing a proposition that includes the one true existential quantifier meaning, whose unrestricted domain is just those objects that exist. So

instead of being truth-value-less, as the claims are on my view, they would be false on the orthodox view.

Now, the same machinery is available to both the orthodox and the anti-orthodox view to deal with this problem. The proponent of orthodoxy need not appeal to linguistic meaning in order to explain why we feel something meaningful is said—there is a complete proposition expressed on this view. But that proposition is false, so the proponent of orthodoxy will need to address that issue instead. It may be maintained that some other proposition is pragmatically conveyed, and this other proposition is true, so this satisfies our intention to say something true. This route is open to both the orthodox and the anti-orthodox view, so it is a bit of a wash.

I think that both sides are on relatively equal footing when it comes to dealing with the problem of speakers making a mistake in the way described here. The primary difference is that my view has speakers expressing gappy propositions without truth values, while the orthodox view has speakers expressing complete, false propositions. I take this to be a slight advantage for my view, since the mistake of expressing a gappy proposition is slightly more understandable than the mistake of expressing a false one. This is certainly debatable, and it may depend on each individual case of a mistake, but I am more comfortable saying that Nick’s mistake at the content level is that the proposition doesn’t end up being complete in the way that it should have, rather than saying that the proposition expressed is something he could easily have said if he’d wanted to, but chose not to. He could have easily said, “There exist superheroes who are animals but also people,” expressing a proposition that he himself believes to be false. But he chose to make use of ‘there is’ in all likelihood because he wanted to say something different than the false proposition—he didn’t realize that he couldn’t actually express the proposition he wanted to, but this mistake fits better with the result of having a gappy proposition be the result, rather than a familiar false proposition.

4.4.3 Ontological Neutrality

Another desideratum we ought to have for any view about the meaning of existential quantifier expressions is that the view remains as neutral as possible about the matter of ontology. Just as a view about names should not force us to have to say whether Sherlock Holmes exists or not and a view about demonstratives should not force a view

about the existence of ‘that polar bear in the corner’, a view about ‘there is’ and ‘there exists’ should not say anything controversial about the existence of objects that may purport to be the subject matter of quantified claims. It’s important to remember that orthodoxy and anti-orthodoxy are views that fall in the realm of metaontology, not ontology itself.

The main ontological dispute that seems relevant here is the issue of whether our ontologies may contain nonexistent objects or not. It can be a complicated matter to separate the ontological from the metaontological question here, but it seems useful to do so. The ontological question is something like, “Are there nonexistent objects?” or, “Does the correct accounting of the objects of the world include nonexistent objects?” The metaontological question I concern myself with here is rather, “Is there a difference in meaning between the expressions ‘there is’ and ‘there exists’. and if so, what is it?” Now, it may be difficult to take up one of these questions without also thereby taking up the other, but I have attempted to do so here. And I claim that it is a virtue of a view about the metaontological question that it remains as neutral as possible on the ontological question. The dispute between the Quinian and the Meinongian over nonexistent objects seems to be a substantive one that ought not to be decided by a view about the meaning of existential quantifier expressions.

How do the two views at hand fare on this matter? I do see a slight advantage to anti-orthodoxy on this issue. The reason is that, since anti-orthodoxy heavily respects our usage, making sense of the way we attempt to quantify about nonexistent objects, it easily adapts to any kind of ontological view a bit more easily. Specifically, this view could be useful for a Meinongian who wants to make a distinction between *existence* and *being*, claiming perhaps that some objects have the property of being but lack the property of existence. I do not claim this kind of ontological view for my own, but I consider it a virtue of my anti-orthodox view of existential quantifier expressions that it does not obviously favor one view over another. Even though my ontological leanings are more in line with Quine’s than the Meinongian’s, I am not so confident in them that I would be comfortable with the orthodox view of existential quantifier expressions.

The orthodox view is not entirely unfriendly to alternative ontologies, but it is not exactly conducive to them either. It would be much more difficult to distinguish between two properties of *existence* and *being* while also maintaining that ‘there exists’ and ‘there

is' are entirely synonymous. It may be more directly important for this issue where we stand on the meaning of the predicates 'x exists' and 'x has being', an issue which I am not taking up here, but these predicates and the properties they denote will clearly be related in some way to the associated quantifier expressions.

On other matters of ontology aside from the issue of Meinongianism, the orthodox view seems unproblematic, or at least on equal footing with my anti-orthodox view. Claiming that 'there exists' means the same as 'there is' does not have any specific consequences about whether, say, numbers exist, or other similar controversial issues in ontology. In fact, when it comes to most ontological discussions, my anti-orthodox view will end up saying virtually the same thing as the orthodox view, since it is part of my view that, since 'there is' is context-sensitive, if we are in a context where we're explicitly talking about existence, 'there is' takes on the same meaning as 'there exists'.¹⁵

4.4.4 Semantic Simplicity

The final desideratum I'll discuss here on its own is the notion of semantic simplicity, which is just the idea that it is a virtue of a theory of meaning, even of a specific small group of words, that the theory avoid needless complication and attempt to give the simplest explanation possible. Of course there is no desire to over-simplify a theory—it is understandable that semantics may sometimes be complex and difficult. But if we're comparing two theories and one of them seems overly complex compared to the other, we may have a reason to prefer the simpler theory.

By this standard, the orthodox view appears to have an edge over anti-orthodoxy. An invariance view is always semantically simpler than a context-sensitive view—the word just means one thing across all contexts, without the extra semantic machinery required for context-sensitivity. There is no need to try make sense of where the contextual index might go, or what features of the context are important for fixing the content of the expression. There is also no need to worry about whether the expression operates more like an indexical, with a fixed character, or a demonstrative that requires some kind of (physical or mental) pointing.

¹⁵ This would make it difficult, though not impossible, for the Meinongians to have a conversation about which objects exist and which just have being—but they themselves would likely recognize this difficulty. It's hard for anyone to talk about nonexistent objects, once it has become explicit that you're intentionally talking about nonexistent objects and their ontological status.

If our only desideratum were semantic simplicity, I would be a proponent of the orthodox view of existential quantifier expressions. But it is important to look at more than just this feature. There's no good reason to expect semantics to be simple—simplicity really only gives us a reason to support one view over another when all else is equal, or at least very close to being equal. I won't claim that my reasons for believing that 'there is' is context-sensitive are on equal footing with the reasons for believing, say, that 'I' is context-sensitive rather than invariant. But it is good to notice that we do accept a great deal of complexity in semantics, when that complexity seems to be moving us towards the correct view.

4.4.5 Summing up the Desiderata

The above sections have detailed how I see the comparison between the orthodox view of existential quantifier expressions and my own anti-orthodoxy. On balance, I think that the only desideratum that really favors the orthodox view is that of semantic simplicity, and that is simply not enough, in the face of other features that are important to such a theory. The orthodox view does not give a good enough explanation of our usage of these expressions, and it does not give a good accounting of what our intentions are when we use them. Furthermore, it is not quite as ontologically neutral as the anti-orthodox view I've presented.

The orthodox view has remained orthodox largely, I believe, because of an absence of acceptable competitors. The primary competitors to orthodoxy generally come from philosophers who are committed to a Meinongian ontology, such as Richard Routley, Terence Parsons, Edward Zalta, and Graham Priest. Their ontological views appear to be primary, and the denial of the orthodox view of existential quantifier expressions then sounds like a patch to make their ontology more palatable. My own view is motivated solely by the desire to understand the meaning of the words 'there is' and 'there exists' better.

It is also worth pointing out that the orthodox view has only been orthodox (at least in the form discussed above) since Quine's work reinvigorated metaphysics. Graham Priest [35] points out that historically, it has been quite common to distinguish terminologically between different existential quantifier expressions, some with ontological

thickness and some without. He points out that Aristotle used what is generally translated as ‘some’ rather than anything with explicit existential reference when making claims that we would use an existential quantifier for. And this, combined with the fact that Aristotle was comfortable with reference to non-existent objects,¹⁶ seems to show at least that Aristotle did not explicitly associate existence with cases of existential quantification.

Priest points out other historical examples from Medieval logicians such as Buridan and Paul of Venice,¹⁷ and then he points out how some 19th century logicians such as Keynes and Venn brought the word ‘existence’ explicitly into logical discussions but explicitly removed ontological significance from it, using it just as a shorthand way to discuss what’s in a domain of discourse.¹⁸ Priest then blames Peirce and Frege for introducing (independently of each other) the existentially loaded quantifier into logic,¹⁹ and he gives an extensive discussion of how Russell went about perpetuating and defending (rather poorly, in Priest’s view) this existential loading.²⁰

The background discussed by Priest gives us a picture of how these expressions have been treated in the area of logic, by logicians, but it is useful for establishing that there is nothing particularly sacred about the meaning of existential quantifier expressions. The intentional ‘unloading’ of ontological significance of the word ‘existence’ by Keynes and Venn that Priest mentions appears to be a clearly technical usage of the word. When it suits our purposes, we may stipulate whatever kinds of meanings we like for these expressions—if it helps draw attention to domains for a discussion in logic, there’s nothing wrong with unloading ‘existence’ in that setting. But it’s good to be clear that this is a technical stipulation, diverging from the meaning of the word.

Although the historical comments of Priest’s are interesting, and they keep us in perspective when thinking about how orthodox the orthodox view is, it is also important to point out that my anti-orthodox view is not intended as an historically-centered view. I’m not relying on etymology or historical usage to make my case—this is a look at what

¹⁶ Priest points out that in the *Posterior Analytic*, Aristotle says, “. . . one can signify even things that are not” (*An. Post* 92b 29-30.)

¹⁷ Priest [35], p. 45. See also Priest [34], section 3.7.1, for more discussion of Medieval logicians and nonexistence.

¹⁸ Priest [35], pp. 46–47

¹⁹ Priest [35], p. 47

²⁰ Priest [35], pp. 48–51

these words mean now, regardless of their evolution. I do think it's important to see that as long as we've had fiction and make-believe and ontological confusion we've needed to have devices in our language for talking about them and purporting to quantify over them. Those devices may change over time, but it seems clear that right now, we have the flexible expression 'there is', which we can use for quantified claims whether we care about the existence of the objects we're trying to quantify over or not. And we reserve 'there exists' for those cases where we're comfortable with the ontological commitment that comes along with it.

Chapter 5

Concluding and Tying Up Loose Ends

5.1 Introduction

In this final chapter, I will conclude by tying up a few loose ends. One point that I have not discussed outside of a few brief remarks is the context-sensitivity of the word ‘knows’. It will be helpful if I make a few remarks here about some connections we might see between that contextualist view and my own—this will also extend my argument for my view as well, since these comments will show that some objections to contextualism about ‘knows’ do not apply in the case of ‘there is’. Another loose end that deserves some remark is the issue of other kinds of unorthodox views about the meaning of existential quantifier expressions. I have mentioned some in passing, but I have not discussed them thoroughly. There are ambiguity views such as Thomas Hofweber’s and Timothy Williamson’s, and there is Priest’s anti-synonymy view. I’ll discuss these in enough detail here to see where they fit in and what we can say about them, in light of the previous four chapters.

5.2 Contrast With Epistemic Contextualism

Contextualism about ‘knows’, or epistemic contextualism, is a plausible view and an elegant approach to some puzzles about knowledge and our knowledge ascriptions. It

has plenty of proponents and detractors, and I'm not going to take a stance on it one way or another. Rather, I'll discuss epistemic contextualism just enough to explore connections we might see between it and my own contextualist view about 'there is'.

The three primary models of contextualism about 'knows' appear to be the following: David Lewis's view, found in his "Elusive Knowledge" [31]; Stewart Cohen's treatment in such papers as "Knowledge and Context" [9] and "Knowledge, Context, and Social Standards" [10]; and Keith DeRose's view, collected recently in his *A Case for Contextualism* [16].¹ I'll focus on Cohen's view, since that is the one I'm most familiar with.

The basic idea of Cohen's contextualism about 'knows' is that the truth of our knowledge attributions—sentences like "*S* knows that *p*"—depends on the epistemic standards of the speaker, which include such things as the information that is currently salient to the speaker at the time of utterance. Cohen gives some details of the view, illustrating it with an example in the following passage:

A common thread in contextualist theories is that the salience of error possibilities raises the standards for how strong one's epistemic position has to be in order for one to know. Contextualism resolves the sceptical paradox by construing our apparently inconsistent intuitions as resulting from contextual shifts in these standards. So, for example, our intuition that the sentence 'Smith knows his car is parked in lot 2' is true is explained by the fact that at ordinary contexts, the statement is in fact true. Our apparently conflicting intuition that the sentence 'Smith does not know his car has not been stolen' is false results from our shifting to a stricter context when we consider the possibility that Smith's car has been stolen. At this new context, the sentence is indeed false. But deductive closure for knowledge is preserved relative to a context. So at everyday contexts, both 'Smith knows his car is parked in lot 2' and 'Smith knows his car has not been stolen' are true. At stricter contexts, neither sentence is true.²

¹ There is much more to the field and the history of Contextualism than these three figures, but they are quite central. Among other things, DeRose's book contains a nice history of Contextualism and survey of the different views.

² Cohen [14], p. 200.

As we see here, one of the features of a contextualist view about ‘knows’ is that it can give us a way to approach, and hopefully solve, some paradoxical cases. How can Smith know where his car is parked when, if we asked him if he knows his car has not been stolen, he would say, “No”? Answer: by bringing up the possibility of cars being stolen, we’re making that possibility salient, and this changes the epistemic standards, thereby changing the truth-value of knowledge attributions.

The truth-value changes from context to context on this view because the knowledge relation that appears at the level of content is sensitive to the epistemic standards of the speaker. John Hawthorne gives a brief and relatively neutral characterization of the semantics of epistemic contextualism here:

[I]t is helpful to think of the extension of ‘know’ on some particular occasion of use as a set of subject–proposition–time triples. According to the contextualist, then, the extension of some tokening of ‘know’ may include some particular triple that does not belong to the extension of some other tokening.³

So with Cohen’s example above, we might say that the proposition expressed by ‘Smith knows his car is parked in lot 2’ will have Smith himself in it, the proposition *that Smith’s car is parked in lot 2*, and a set of triples that is the content of ‘knows’. If that set of triples contains the Smith–*p*–*t*, where ‘*p*’ is the proposition *that Smith’s car is parked in lot 2* and ‘*t*’ is the time of utterance of the entire sentence, is a member of that set of triples, then the sentence is true. But if the triple is not there, it’s false. What decides whether that triple is in the set or not will be the epistemic standards in the context, such as what is salient to the speaker.

Unlike epistemic contextualism, on my view, the content of the expression is dependent on the intentions of the speaker (the intended domain), not what is salient to the speaker. This difference is important because it allows me to avoid one problem that a contextualist view like Cohen’s faces, *viz.* the problem of switching contexts. Since a crucial factor in determining epistemic standards, on Cohen’s view, is what is salient to the speaker, there is a problem insofar as it is quite difficult, once something has become salient, to make it un-salient. That is, once a speaker begins to worry about whether

³ Hawthorne [17], p. 58.

her car has been stolen, she cannot just easily put that worry to the side. She may relax and not feel over-anxious about it, but it will still be salient in a significant sense. This makes it very difficult to speak in one context about what we say in other contexts without doing a good deal of quoting and being careful about use-mention issues.

On my view, talking about different ways we may have meant our expression or our assertion is just a matter of making clear what our intentions were, or might be. Even if the context has been altered from one where we intend to quantify over nonexistent objects to one where we are explicitly concerned with ontology, or what exists, we can still shift back to ordinary contexts with relative ease. We can say things like, “Okay, enough ontology, let’s go back to talking about comic books like normal people.” This makes our shift in intentions clear, and our content follows along obediently. This would not be so easy if the content was indexed to what is salient to us, since it is reasonable to think that the ontological discussion and the worries about existence would still be salient in some sense after just completing a discussion about them.

Another issue that is raised as a problem for epistemic contextualism that will be worth discussing here is that of the model of context-sensitivity chosen. Cohen’s view is that the indexical model is the correct one for ‘knows’,⁴ and this appears to be the prevailing notion. My own contextualist view about ‘there is’ chooses the demonstrative model rather than the indexical one. What difference does this make? It may make several differences, but one that I will focus on briefly here is the fact that both kinds of contextualism will have error-theoretic components, and the model chosen may make a difference with respect to this.

I’ve discussed the error-theoretic aspects of my own view in previous chapters, noting that I am committed to saying that speakers make an error if they think they are asserting something true when they attempt to quantify over nonexistent objects. On Cohen’s view, the error a speaker is making is of a different kind. His view is designed to show how speakers may be speaking truly from context to context, even if they assert p in one context and $\text{not-}p$ in another. That is, I may truly say, “Smith knows where his car is parked,” in a low-standard context and then, while worrying about car theft a few minutes later, may truly say, “Smith does not know where his car is

⁴ In Cohen [11], for example, he says, “[T]he theory I wish to defend construes ‘knowledge’ as an indexical,” p. 97.

parked.” The error that epistemic contextualism is committed to, then, is not in the truth-value of our claims. Rather, it is an error that comes about when we look at these two different claims, uttered in different contexts, and think that they are inconsistent claims. They are not *really* inconsistent, since the first was true in the context in which it was uttered, and the second was true in its own context. Cohen’s treatment of this is just to point out that the invariantist about ‘knows’ is on no better footing, since it is just as problematic—or perhaps moreso—to deny the consistency of these claims and have to say that the speaker errs by thinking that these claims are both true.⁵

Cohen’s response seems just fine, so I am not going to say there’s a difference in that his view can’t handle the error-theoretic consequence and mine can. The only point I want to make with this comparison is that we might think that a demonstrative view lends itself a bit more to the *expectation* of an error than an indexical view. The prime difference between indexicals and demonstratives is that indexicals are governed by a fixed character, a semantic rule that leads from the context of utterance to a content, while demonstratives are governed by a pointing in the context (a mental pointing, on my view). When it come to speakers making errors, then, we might note that it is generally more difficult to become confused about indexicals than about demonstratives. If I’m a competent speaker of English, I’ll understand the rule that governs words like ‘I’ and ‘today’ in a way that prevents me from making errors about the content of them. But when pointing is involved, whether it’s physical or mental, I may understand how ‘that’ works perfectly well and still mess up when I make or interpret a demonstration.

Cohen’s treatment of the error theoretic aspect of his view handles this issue, since the error he accepts is not one at the level of content, but rather at the level of noticing the consistency of two different claims, uttered in different contexts. So this error is more like someone thinking that the claims, “Today is cold,” and, “Today is not cold,” are inconsistent until it is pointed out that they were uttered on two different days. My point here, though, is that with a demonstrative model, an error theoretic aspect of the view should be more comfortable, more expected even, then with an indexical model. Indexicals have more semantic information, in a sense, that competent speakers are expected to know, so if something goes wrong, the error may seem a bit more serious.

⁵ See Cohen [13], p. 70.

5.3 Some Other Unorthodox Views

Now I'd like to switch topics to another loose end that is worth tying up—the discussion of a few other anti-orthodox views about the meaning of existential quantifier expressions. Recall that there were two primary tenets of the orthodox view:

Orthodox View of Existential Quantifier Expressions:

- (O1) There is only one kind of existential quantification.
- (O2) Existential quantifier expressions are synonymous and interchangeable.

And my own denial of the orthodox view had two primary tenets of its own:

Anti-Orthodox View of Existential Quantifier Expressions:

- (A1) 'There is' is a context-sensitive expression.
- (A2) 'There exists' is invariant in meaning.

In effect, I deny both (O1) and (O2), since I see 'there is' and 'there exists' as having significantly different meanings, the former being context-sensitive and the latter being invariant.

Other views, proposed by philosophers like Thomas Hofweber and Graham Priest, deny the tenets of the orthodox view in different ways than my own denial. Hofweber's view is that existential quantifier expressions are ambiguous, denying (O1), while not taking issue with (O2), since he does not propose a distinction between different existential quantifier expressions. Priest takes a different route, denying (O2) by distinguishing between different existential quantifier expressions in their meaning, though it's unclear whether he should be taken as denying (O1) or not. I'll focus on these two views, since they seem interesting and relevant, though they are certainly not the only anti-orthodox views out there.

5.3.1 Hofweber's Ambiguity View

Thomas Hofweber's view is that existential quantifier expressions are ambiguous between two readings. The 'domain-conditions' reading is the one I've discussed, giving the traditional semantics for it in Chapter 2, while the 'inferential role' reading is more like substitutional quantification. Hofweber explains this in the following passage:

Beside the domain-conditions reading, quantifiers also have an inferential role reading, or also internal reading. In this reading sentences with quantifiers are inferentially related to sentences without such quantifiers (in the simplest case). For example, the quantifier ‘something’ in this reading has the inferential role that ‘Something is F’ follows from each instance of ‘F(t)’, for any term ‘t’ in the language. I have argued in [Hofweber [19]] that we do have a need in ordinary communication to have an expression that plays exactly this role. This need has nothing to do with metaphysics but rather with a need to communicate information in a situation of partial ignorance. But in a language like ours, where not every term denotes an object in the domain of discourse, and not every object is denoted by a term, no single contribution to the truth conditions can yield both a domain condition and an inferential role. In the limit, when every term denotes an object and every object is denoted by a term, these two can coincide, but in our language they do not.⁶

Hofweber is motivated by similar concerns as mine—we clearly need a device in our language for at least purporting to quantify over nonexistent objects. He just takes a different route, ending up by saying that sometimes when we quantify we mean it in the normal, ontologically-committing way, but other times we mean it just in a substitutional way. The substitutional way is non-ontologically-committing, since it is just focused on the inferential role of the quantifier, not what is in the domain.

Now, the issue of the sensibility of substitutional quantification is interesting in its own right. It may be difficult to separate off the domain from the quantifier, say that it’s unimportant on these uses, and still make sense of the claim being a *quantified* claim.⁷ My own primary reason for not finding an ambiguity view acceptable is the difficulty we have with using a quantifier expression in two different ways within a single sentence. We cannot say, for example, “There are superheroes who can fly, but there aren’t any.” If ‘there is’ were ambiguous, we should be able to say this in a way that we could, once we’re paying attention, hear as an acceptable sentence, making use of two

⁶ Hofweber [21], p. 28

⁷ See, for example, Peter van Inwagen [45] for a discussion of his worries about substitutional quantification.

different senses of the quantifier expression. I don't think that works, and I also think that my own contextualist view captures the phenomena that motivate both Hofweber and I in a better way than his ambiguity view.

Hofweber also makes some other comments that lend support to my overall approach as well as to his, while discussing the role of 'there exists' as opposed to other quantifier expressions. He points out that the Meinongian position treats 'exists' as a contextual restriction on other quantifiers, in the way that 'who owns a scooter' restricts 'everyone' in the claim, "Everyone who owns a scooter likes to watch *Battlestar Galactica*." Hofweber thinks the intuition behind looking at 'exists' as a restrictor like this is a strong intuition, and a problem for the Quinian. He takes himself to have a non-Meinongian explanation of this intuition (as I take myself to have as well). Since he thinks that there are two readings of quantifier expressions, a claim with an expression like 'something' is sometimes underspecified, on his view. If I just say, "Something makes me happy," that on its own is underspecified, whether to read 'something' in the internal (inferential role) or the external (domain conditions) way. It could be that Thor makes me happy, and Thor doesn't exist, making it more appropriate to read my claim in the inferential role way. With this in mind, Hofweber claims that 'exists' is a specifier, i.e., an indicator as to which reading of the quantifier we should use. He says, "The difference between 'something' and 'something, which exists' is that the latter forces that the quantifier is used in its external [domain conditions] use."⁸

I don't know if I would put this quite as Hofweber does, saying that the intuition that 'exists' is a restrictor is a strong one. But I do agree that something along these lines is correct, and the Meinongian kind of intuitions do need to be accounted for. I think my own view offers quite a bit better of an explanation of this phenomenon than Hofweber's, since he is in a position where he has to say that the word 'exists' pushes us towards one of the two the meanings of existential quantifier expressions, even though he's not actually making a distinction in meaning between the different existential quantifier expressions 'there is' and 'there exists'. On my view, which argues for a difference in meaning between them, it makes much more sense why the word 'exists' would have an effect on our conversation. First of all, it is associated invariantly with the domain of existing objects. Secondly, it is the kind of word that, when brought into a conversation,

⁸ Hofweber [19], p. 271

has the power to shift the conversation from a context where existence was not important to one in which it now is. If we were using ‘there is’ before, talking about superheroes to our kids, once someone brings up existence, the conversation can easily shift in an important way. And if we have a feeling that there’s been a restriction of the domain, it’s because we were intending to quantify over nonexistent objects before, and now our intentions have shifted to just the domain of existing objects.⁹

5.3.2 Graham Priest’s Denial of Synonymy

A different kind of anti-orthodox view comes from Graham Priest, a member of the lineage of neo-Meinongians. Priest [35] holds what appear to be a primary and a secondary view about the existential load that the existential quantifier—what he calls the ‘particular’ quantifier—expressions carry with them. The primary view is that ‘there exists’ is an existentially loaded expression, but ‘some’ is not. The secondary view is basically an acceptance that ‘there is’ means the same thing (or carries the same load) as ‘there exists’. I call this a ‘secondary’ view because he does not seem terribly committed to it, but he definitely leans toward it. His main point, though, is that at least one existential quantifier expression (‘some’) is not existentially loaded, which is a challenge to the received Quinian view. Priest gives both of his views in the following passage:

Meinongians of various kinds, including noneists such as myself, hold that one can quantify over something without taking it to exist. More specifically, what is most naturally called the particular quantifier (being the dual of the universal quantifier) should not be read as ‘there exists’—or even ‘there is’, there being no real difference between being and existence; it should simply be read as some, leaving it open whether the some in question exists or not.¹⁰

⁹ Another non-Meinongian ambiguity appears in Timothy Williamson’s [48], where he proposes a distinction between a ‘logical’ and a ‘commonsensical’ reading of the existential quantifier, in order to handle a puzzle about necessary existents. I don’t have the space to discuss this view further here, but it may be one more interesting view (among others that I’m sure I’m unaware of) to add in to the ambiguity mix.

¹⁰ Priest [35], p. 42

Here he clearly states both the primary and secondary views I mentioned above. The language of existential ‘loading’ does not come here, but he does use it elsewhere.¹¹ He thinks that ‘some’ is a legitimate existential quantifier expression, but it does not commit us to the existence of what we quantify over. What makes the secondary view secondary is something he says here:

Distinguishing between ‘there is’ and ‘there exists’ invites the view that there is a second-class notion of existence, ‘subsistence’—a view that has caused much confusion and so is best avoided. Honestly, there is no difference between existence and being. For all that, it is possible, of course, to use ‘there is’ simply to mean ‘some’ if one wants.¹²

This claim, especially the last sentence, makes it clear that he does hold the view that being and existence are the same, but he’s not all that interested in arguing for it, and he doesn’t seem to think it would make much difference if we went the other way on this point.¹³

Priest’s discussion is focused on Quine, since his “On What There Is” was so hugely influential as the root of the orthodox view. In discussing Quine, Priest says the following:

Quine argues that the use of names and predicates is not existentially committing, but there is absolutely no argument given as to why quantification is existentially committing. Quine simply assumes that the domain of quantification comprises existent objects—or what comes to the same thing, that the particular quantifier is to be read as ‘there is’. No argument is given for this: it is stated simply as a matter of dogma. [...] At any rate, if Russell used bad arguments for the view, Quine uses none at all.¹⁴

¹¹ Priest [35], p. 42

¹² Priest [35], p. 43, n. 6

¹³ Though there are different consistent ways one could hold a view like Priest’s, he takes ‘there exists’ to involve an existence predicate. In the formal language, we could represent this with the particular quantifier (he uses a ‘ \mathfrak{E} ’ symbol) quantifying into the existence predicate, or we could just abbreviate this as ‘ \exists ’ if we want to continue to use that as a symbol of an existentially loaded quantifier. See Priest [34] section 1.4, for a discussion of this issue.

¹⁴ Priest [35], p. 52

Priest is wrong to say there is no argument at all in “On What There Is” for this view of Quine’s, though not *tragically* wrong. I discussed Quine’s argument (about ratios) in the previous chapter, including why I found it lacking as an argument for the orthodox view. Now, Priest doesn’t claim that Quine *never* gives an argument for this view. He doesn’t discuss any others, but he does say that he’s, “not suggesting that there are no other arguments for taking the particular quantifier to be existentially loaded, possibly even in Quine’s later corpus.”¹⁵ This could be a reference to the argument in Quine [39] that I also discussed in the previous chapter. But, since Priest’s aim here was mostly historical, he focused only on “On What There Is” because it was so influential of an essay. Priest sees this essay as having solidified the view that the quantifier is existentially loaded, and he’s further pointing out the problem of this view having been solidified by a paper with no (good) argument for the view.

Overall, I’m friendly to Priest’s approach, though I do think he misses the mark in drawing the distinction between ‘some’ on the one hand and ‘there is’/‘there exists’ on the other. He does not give a substantial argument for this particular split, and I think the reasons he does provide should push him (like me) to seeing ‘there exists’ on one side and ‘there is’/‘some’ on the other. I have not discussed ‘some’ explicitly in this dissertation, since space requires me to stay quite narrowly focused on this topic, but I offer (without argument) that I maintain ‘some’ to mean essentially the same thing as ‘there is’, context-sensitive in the same way.

Also, since Priest’s view is mostly given without semantic details, I don’t have cause to agree with or criticize it on that level. He does appear to be motivated by his ontology in making this distinction, though, so reading his work as a non-Meinongian, one does not exactly have a strong reason to agree with him. It appears as if he takes the quantifier referred to by ‘some’ to be the normal existential quantifier we study in logic, with the existential loading coming in as either an ‘existence’ predicate or as a restriction of the domain of the quantifier. Either way, these are methods available only to the Meinongian who allow, in some way or another, for there to be objects that do not exist.

¹⁵ Priest [35], p. 53

5.4 Conclusion

In this chapter, my goal has been primarily to tie up a few loose ends that seemed worthy of discussion, but not until now. Both the issue of epistemic contextualism and that of other anti-orthodox views would have been a bit of a distraction in the previous chapters, but they should enrich the discussion here at the end, adding a bit to the arguments from the previous chapters.

To sum up the entire project, my primary goal has been to present and argue for an anti-orthodox view of meaning of the existential quantifier expressions ‘there is’ and ‘there exists’. In Chapter 1, I presented the notion of ‘ontological robustness’ in order to motivate my view, discussing how ‘there exists’ claims appear to always be ontologically robust, while ‘there is’ claims appear to be sometimes ontologically robust, sometimes not. In Chapter 2, I presented the details of my contextualist view of ‘there is’, including also some discussion of why I take ‘there exists’ to be invariant. In Chapter 3, I extended the discussion of domains and nonexistent objects to fill out my view and address some potential objections. In Chapter 4, I discussed the details of the orthodox view and presented my holistic argument against it, tying together the pieces from the previous chapters and adding some new points in favor of my view. And here in Chapter 5, I’ve added a discussion of epistemic contextualism and two other anti-orthodox views about existential quantifier expressions.

References

- [1] Joseph Almog, John Perry, and Howard Wettstein, editors. *Themes from Kaplan*. Oxford University Press, New York, 1989.
- [2] J. L. Austin. *Sense and Sensibilia*. Oxford University Press, Oxford, 1962.
- [3] Jody Azzouni. *Talking About Nothing: Numbers, Hallucinations, and Fictions*. Oxford University Press, Oxford, 2010.
- [4] Jon Barwise and Robin Cooper. Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4:159–219, 1981.
- [5] George S. Boolos, John P. Burgess, and Richard C. Jeffrey. *Computability and Logic*. Cambridge University Press, Cambridge, 5th edition, 2007.
- [6] David Braun. Empty names. *Noûs*, 27:449–69, 1993.
- [7] David Braun. Empty names, fictional names, mythical names. *Noûs*, 39:596–631, 2005.
- [8] David Chalmers, David Manley, and Ryan Wasserman, editors. *Metametaphysics*. Oxford University Press, Oxford, 2009.
- [9] Stewart Cohen. Knowledge and context. *Journal of Philosophy*, 83:574–83, 1986.
- [10] Stewart Cohen. Knowledge, context, and social standards. *Synthese*, 73:3–26, 1987.
- [11] Stewart Cohen. How to be a fallibilist. *Philosophical Perspectives*, 2:91–123, 1988.
- [12] Stewart Cohen. Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13:57–89, 1999.

- [13] Stewart Cohen. Contextualism defended some more. In Matthias Steup and Ernest Sosa, editors, *Contemporary Debates in Epistemology*, pages 67–71. Blackwell, Oxford, 2005.
- [14] Stewart Cohen. Knowledge, speaker and subject. *The Philosophical Quarterly*, 55:199–212, 2005.
- [15] Max Cresswell. Formal semantics. In Michael Devitt and Richard Hanley, editors, *The Blackwell Guide to the Philosophy of Language*, pages 131–46. Blackwell, Oxford, 2006.
- [16] Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press, Oxford, 2009.
- [17] John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, Oxford, 2004.
- [18] Eli Hirsch. Ontology and alternative languages. In Chalmers et al. [8], pages 231–59.
- [19] Thomas Hofweber. Quantification and non-existent objects. In Anthony Everett and Thomas Hofweber, editors, *Empty Names, Fiction and the Puzzles of Non-Existence*, pages 249–73. CSLI Press, Stanford, 2000.
- [20] Thomas Hofweber. A puzzle about ontology. *Noûs*, 39:256–83, 2005.
- [21] Thomas Hofweber. Innocent statements and their metaphysically loaded counterparts. *Philosopher's Imprint*, 7:1–33, 2007.
- [22] David Kaplan. Afterthoughts. In Almog et al. [1].
- [23] David Kaplan. Demonstratives. In Almog et al. [1].
- [24] David Kaplan. The meaning of ouch and oops: Explorations in the theory of *Meaning as Use*. UC Berkeley Graduate Council Lectures, 2004.
- [25] Edward L. Keenan. Natural language, sortal reducibility and generalized quantifiers. *The Journal of Symbolic Logic*, 58:314–325, 1993.

- [26] Edward L. Keenan. The semantics of determiners. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 41–63. Blackwell, Oxford, 1996.
- [27] Edward L. Keenan and Jonathan Stavi. A semantic characterization of natural language determiners. *Linguistics and Philosophy*, 9:253–326, 1986.
- [28] Jeffrey C. King. *Complex Demonstratives: A Quantificational Account*. MIT Press, 2001.
- [29] David Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61:343–77, 1983.
- [30] David Lewis. Noneism or allism. *Mind*, 99:23–31, 1990.
- [31] David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–67, 1996.
- [32] David Lewis. How many lives has schrödinger’s cat. In Frank Jackson and Graham Priest, editors, *Lewisian Themes*, pages 4–23. Oxford University Press, Oxford, 2004.
- [33] Stephen Nelson. Fictional worlds: An anti-realist view of fictional characters. unpublished manuscript, 2011.
- [34] Graham Priest. *Towards Non-Being: The Logic and Metaphysics of Intentionality*. Oxford University Press, Oxford, 2005.
- [35] Graham Priest. The closing of the mind: How the particular quantifier became existentially loaded behind our backs. *The Review of Symbolic Logic*, 1, 2008.
- [36] W. V. O. Quine. On what there is. *Review of Metaphysics*, 2:21–38, 1948.
- [37] W. V. O. Quine. Two dogmas of empiricism. *Philosophical Review*, 60:20–43, 1951.
- [38] W. V. O. Quine. *Word and Object*. MIT Press, Cambridge, MA, 1960.
- [39] W. V. O. Quine. Existence and quantification. *L’Age de la Science*, 1:151–164, 1968.

- [40] Agustín Rayo and Gabriel Uzquiano, editors. *Absolute Generality*. Oxford University Press, Oxford, 2007.
- [41] Nathan Salmon. Nonexistence. *Noûs*, 32:277–319, 1998.
- [42] Theodore Sider. *Four-Dimensionalism*. Clarendon, Oxford, 2001.
- [43] Theodore Sider. Ontological realism. In Chalmers et al. [8], pages 384–423.
- [44] Jason Stanley and Zoltan Gendler Szabó. On quantifier domain restriction. *Mind and Language*, 15:219–61, 2000.
- [45] Peter van Inwagen. Why i don’t understand substitutional quantification. *Philosophical Studies*, 39:281–85, 1981.
- [46] Peter van Inwagen. *Material Beings*. Cornell University Press, Ithaca, NY, 1990.
- [47] Peter van Inwagen. Meta-ontology. *Erkenntnis*, 48:233–50, 1998.
- [48] Timothy Williamson. Necessary existents. In A. O’Hear, editor, *Logic, Thought and Language*, pages 233–51. Cambridge University Press, Cambridge, 2002.