
Quantification and Non-Existent Objects

THOMAS HOFWEBER

1.1 Non-existent Objects

Whether or not there are non-existent objects seems to be one of the more mysterious and speculative issues in ontology.¹ To affirm that there are non-existent objects is to affirm that reality consists of two kinds of things, the existing and the non-existing. The existing contains all of what is in our space-time world, plus all abstract objects, if there are any. Most people, it seems fair to say, would think that this is all there is. For them the only real question in ontology can be what kinds of existing things there are. However, followers of Meinong maintain that this isn't all there is. There is also another kind of things, those that do not exist. And to say this, the Meinongians continue, is to accept that reality is divided into two basic kinds of things, the existing and the non-existing. Whether or not reality contains two basic categories of things, existing and non-existing, or only one, existing, is what the debate about non-existent objects is all about. And as such it seems to be the most speculative of the debates in ontology. How could we human beings possibly decide it? One might think that to find out whether or not there are abstract objects is hard to decide, since they are not in space and time, causally inaccessible, unobservable, etc.. But whatever difficulty there might be to answer the question whether or not there are abstract objects, it has to be even harder to decide whether or not there are non-existent objects. Abstract objects, if there are any, at least exist. non-existent objects at best seem to fill out the space of what there is.

¹Thanks to Johan van Benthem, Sol Feferman, John Perry and Ed Zalta for comments on earlier drafts.

To affirm that there are such things seems to engage one in ontological speculation of the highest kind.

Reasonable as this might be, it on the other hand seems to be quite trivial to argue that there are non-existent objects. To accept that there are non-existent objects doesn't seem to come down to much more as to accept such trivialities as that, say, Santa Claus doesn't exist. To illustrate this, let's first consider the question whether or not there are non-cooperative objects. This, it seems, is easily answered by the fact that Fred does not cooperate. That Fred doesn't cooperate can be said in many different and more or less direct ways. Consider:

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

To say that Fred doesn't cooperate is, modulo subtleties, the same as to say that Fred is a non-cooperative person, (or thing, or object). These are just different ways of saying the same thing. So, since Fred is a non-cooperative object, doesn't that answer the question whether or not there are non-cooperative objects? And similarly for the case of Santa's non-existence. That Santa doesn't exist can also be said in many different ways:

- (4) Santa does not exist.
- (5) Santa is non-existent.
- (6) Santa is a non-existent person, (or thing, or object).

So, isn't it trivial to decide whether or not there are non-existent objects? After all, Santa doesn't exist, and to say that Santa is a non-existent object seems to be no more than a fancy way of saying that Santa doesn't exist.

To be sure, one can believe that the acceptance of non-existent objects is implicit in accepting that Santa doesn't exist. But this might not be taken to settle the issue about non-existent objects. One can simply claim that all this shows is that Santa's non-existence is to be taken to be an equally controversial and difficult issue as there being non-existent objects. This is certainly an option, but not one that seems to be very attractive to take. After all, a philosopher's claiming that the existence or non-existence of Santa Claus is an open and substantial philosophical problem isn't usually taken to shed a favorable light onto contemporary

²Of course, in ordinary discourse "object" is often contrasted with "person". However, I will use the word "object" to be more general than "person", as it is commonly used in debates about ontology.

philosophy. Usually, but not always, the participants in the debate are happy to concede that Santa doesn't exist, but they don't take this to answer the question whether or not there are non-existent objects. This question isn't answered by this because it doesn't answer the question whether or not there is a Santa Claus. Only if we knew that

(7) There is a Santa Claus.

and in addition that

(4) Santa doesn't exist.

would we have an answer to the question whether or not there are non-existent objects. This reasoning thus might grant that Santa doesn't exist, but transfers the real issue of deciding whether or not there are non-existent objects into the issue of deciding whether or not there is a Santa Claus. Thus the issue gets transformed into an issue about the truth of statements involving quantifiers. And to get the truth value of such quantified statements as (7) right is the really tricky case, or so the common line.

This, again, seems perfectly reasonable, but also quite problematic. Can the move to quantified statements really be that central? Consider again the question whether or not there are non-cooperative objects. Suppose we agree that

(1) Fred is doesn't cooperate.

and thus that

(3) Fred is a non-cooperative person, (or thing, or object)

Could it really be that we might reasonably have a substantial disagreement about whether or not

(8) There are non-cooperative objects.

is true? It might seem reasonable to say that this last issue has been resolved by example: of course there are non-cooperative objects, and Fred is one of them. Similarly for the case of non-existent objects. If we agree that

(4) Santa does not exist.

and thus that

(6) Santa is a non-existent person, (or thing, or object).

how can it still be a substantial question whether or not there are non-existent objects? Again, this question seems to have been answered by example. Thus it seems to be a trivial inference to conclude from

(6) Santa is a non-existent person, (or thing, or object).

that

- (9) There is a non-existent object, namely Santa.

If one would want to deny that inference one would have to deny that there is a tension or inconsistency between

- (10) Santa is a non-existent object.

and

- (11) Nothing is a non-existent object.

But these sure seem to be in conflict with each other, and it will need quite a bit of philosophical sophistry to explain that away.

So, we are in a dilemma, a dilemma with the following two horns:

- a) On the one hand, it seems reasonable that in order for us to decide the ontological dispute about non-existent objects we have to look at the truth value of statements with quantifiers that range over non-existent objects. To do this is to engage in a substantial and possibly quite speculative philosophical project.
- b) On the other hand, it seems that it follows trivially from the uncontroversial facts like that Santa doesn't exist that there are non-existent objects. This inference can go as follows:
- (4) Santa Claus does not exist.
- (10) Santa is a non-existent object.
- (9) There is a non-existent object, namely Santa.

To understand this dilemma better is to gain a better understanding of the debate about non-existent objects, and about ontology in general. This paper is supposed to shed some light on how this dilemma is to be resolved. To do so we have to have a closer look at quantification.

1.2 Quantification

The standard debate about whether or not we should accept an ontology of non-existent objects is very closely connected to the debate whether or not we should or accept quantification over non-existent objects. In fact, it is a classic "one person's modus ponens is another person's modus tollens" debate. It is commonly agreed upon among the members of that debate that

- (I) If there are statements that are literally true and that contain quantifiers that range over non-existent objects then reality consists of two kinds of things: the existing and the non-existing.

However, it is controversial whether or not we should accept such quantified statements. On the one side (the Meinongean, or modus ponens,

side) are people who point out that we do in practice use such quantified statements and that therefore we should just accept that reality contains more than we naively thought. On the other side (the anti-Meinongean, or *modus tollens*, side) are people who think that such an ontology is absurd and that therefore we should not quantify over non-existent objects, unless, of course, we are not trying to make a literally true statement.³

1.2.1 Quinean and non-Quinean quantifiers

It seem that there is very good reason to believe that (I) is true and that therefore one either has to reject all quantification over non-existent objects, or accept a Meinongean ontology. To see if this is indeed so we shall have a look at how (I) is motivated.

First, an observation that leads to some useful terminology. Those who accept quantification over non-existent objects basically accept that quantifiers come in two kinds. There are ordinary quantifiers, as they occur in ordinary utterances of

(12) Someone ate my sandwich.

and there are the one's that apparently require an ontology of non-existent objects, as in

(13) Someone is smarter than any real detective, namely Sherlock Holmes, but unfortunately he doesn't really exist.

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

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

is more than odd. When quantifiers are used in the way in which they apparently range over Sherlock and the like, such an explicit modification without change of truth conditions does not seem possible. (13) modified this way seems clearly false.

Let's call the occurrence of a quantifier in an utterance that is such that we can explicitly modified with "which/who exists" without change of truth conditions of that utterance a **Quinean quantifier**. Let's call those occurrences where such a modification is not possible without change of truth conditions a **non-Quinean quantifier**. I'd like to stress

³One further option is to accept quantification over "non-existent" objects like Santa, but deny that they are properly non-existent. For example, according to van Inwagen quantification over Santa Claus is to be accepted, and in some sense it is true that Santa doesn't exist, but still, Santa isn't a non-existent object. See (van Inwagen 2000). The present considerations are most relevant for this line of reasoning, too, even though the difference between it and Meinongianism isn't discussed in any detail here.

here that the distinction between Quinean and non-Quinean quantifiers is one that applies at the level of individual utterances of quantifiers. It is not a distinction at the level of language. So, even if this distinction is legitimate and not empty, it is a further question whether or not our language has two kinds of quantifiers in it, or whether or not to account for the difference between the two kinds of occurrences of quantifiers in another way.

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

- i) Are there any legitimate uses of non-Quinean quantifiers? Is there a need for us to take recourse to them when we try to make a literally true statement? Is there a need for them outside of metaphysics, in ordinary communication? Should we accept such utterances as true?
- ii) If yes, how should we understand the function of the non-Quinean quantifiers in these uses?

1.2.2 Meinongians and non-Quinean quantifiers

Meinongians believe that the question “Are there non-existent objects?”, understood as a substantial ontological question, has an affirmative answer. The main, or at least one of the main arguments in favor of this view, is that we do in fact use non-Quinean quantifiers in apparently true sentences. If we look at what we accept to be true we see that it contains statements involving non-Quinean quantifiers. And the way to understand them is, according to Meinongians, as follows:⁴

- An expression like “something” can be used both as a Quinean and as a non-Quinean quantifier, as in
 - (15) Something is eating my cheese, probably a mouse.
 - and in
 - (16) Something is keeping me awake at night, namely the monster I dream about.
- Thus there has to be some difference in the particular occasions of the utterance which makes it that one can be modified with “which exists” without change of truth conditions, whereas the other one can’t.
- The way to understand this is simply the following: Quinean quantifiers are a case of a well known way in which the context of utterance of a sentence with a quantifier contributes to the truth conditions, or what is said with the utterance, namely contextual

⁴See, for example, (Parsons 1980).

restriction of the quantifier. Quinean quantifiers are implicitly restricted to what exists. Non-Quinean quantifiers don't have this restriction. This phenomenon is just like contextual restrictions of quantifiers in standard examples of utterances like:

(17) Everyone has to die.

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

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

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

In addition and on top of that, it seems that one can't reasonably draw the line which non-existent objects one takes to be part of reality, if one accepts any at all. It seems that one can't reasonably say that all the non-existent objects there are are the ones that we happen to talk about, like Santa Claus and the like. If non-existent objects are part of reality at all then it seems that there isn't any good reason to assume that just the ones we happen to talk about are part of reality. We might as well have come up with other myths and stories, and everything else would have remained the same. It would be very surprising if we got so lucky that the non-existent things we actually talk about are part of reality, but the ones we might as well have talked about aren't. Thus it seems that if one accepts non-existent objects at all then one has to accept a plenitude of non-existent objects: every conceivable one has to be just as good as any other one. Thus if there are non-existent objects at all then there have to be all conceivable non-existent objects. The only way how it might be otherwise is that what non-existent objects there are somehow *depends* on our talking about them. But this option seems hard to defend. It is already very difficult to make sense of how the existence of something depends on our talking about it. But how can the *non-existence* of something depend on our talking about it?

This reasoning, and the dilemma that arises from the acceptance of (I), defines much of the debate about non-existent objects. I suspect that it is the apparent implausibility of there being a plenitude of non-existent objects, that reality contains already anything we might conceivably come up with and start to talk about, together with the acceptance

of (I), that makes people vehemently deny that there can be any true statements containing non-Quinean quantifiers. On the other hand, there seem to be a number of very plausible cases where we seem to have to accept quantification over non-existent objects. These cases, together with the acceptance of (I), seems to imply such an ontology of non-existent objects.

I think the above reasoning that leads to the acceptance of (I) is mistaken, and in this paper I will try to spell out what the mistake is on which it is based. In particular, I will try to show that the literal truth of statements with non-Quinean quantifiers does not imply the substantial metaphysical pictures. I will defend this by motivating that in ordinary communication quantifiers have more than one function. In the next section we will talk about the functions of quantifiers in ordinary communication. After that we will return to talk about non-existent objects.

1.3 The function of (some) non-Quinean quantifiers

In this section I will argue that quantifiers like “something” really have two different, but related, roles in communication. They have different functions, and they can differ in what contribution they make to the truth-conditions. Once spelled out what these are we will see that this is most relevant to our discussion about non-existent objects.

1.3.1 Subtle contextual contributions to content

The main step in the Meinongian’s argument for (I) is to view the contextual difference between Quinean and non-Quinean quantifiers as one of contextual restriction. The Meinongian’s motivation that non-Quinean quantifiers are restricted quantifiers is closely related to a certain view about what the role of the context of an utterance is in determining its content. One view is that the context provides

- the values of the indexicals and demonstratives that occur on the sentence uttered, and
- contextually restricts the quantifiers that occur in the sentence uttered.

If this is all the context does, and given that there is a contextual difference between Quinean and non-Quinean quantifiers, then it seems obvious that the difference between these two quantifiers is one of contextual restriction.

Undoubtedly, the context makes at least the above two contributions to the content. The question we have to address here is whether or not this is all that the context does. If the context could contribute in

other, maybe more subtle, ways to what the content is then we would have to ask ourselves if the argument that the contextual difference between Quinean and non-Quinean quantifiers indeed is one of restricted quantification. I will now point to several examples that show that there are other, more subtle contributions that context can make. I will go over three cases of this phenomenon, one of them is related to quantifiers. Then we will return to the discussion of non-Quinean quantifiers.

Genitive

Consider a standard use of genitive, like

(19) Joe's car has a flat.

It seems that what the genitive "s" does in (19) is that it contributes to what is said overall that Joe owns a certain car. It is the car that Joe owns that is said to have a flat. Thus it seems that what the genitive does is contribute the relation of ownership that holds between Joe and a certain car to the content of what is said.

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

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

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

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

Joe's book is full of mistakes.

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

that a book that has that relation to Joe is F.

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

Plural

Another, very widely discussed, example of a not immediately obvious kind of contextual contribution to content is plural.⁵ Consider the following example:

(22) Four philosophers carried three pianos.

The plural phrases “four philosophers” and “three pianos” can each make at least two different contributions to the truth conditions of the utterance. First, they could be read as being about individual philosophers, or individual pianos. Secondly, they could be read as being about a group or collections of philosophers or pianos. Thus an utterance of (22) can have at least the following truth conditions:

- (23) a. Four philosophers together carried three pianos each (one after the other).
 b. Four philosophers each carried three pianos each (one after the other).
 c. Four philosophers together carried three pianos together (all of them at once).
 d. Four philosophers each carried three pianos together (all at once).

Now, given what we all know about pianos and the strength of philosophers, the most natural way to understand an utterance of (22) is of course (23a). But an utterance of (22) can have the truth conditions as spelled out in (23d). That it usually doesn't comes from the fact that we talk about pianos, and philosophers, and what we know about them. If we talk about other things then this won't be the default reading, as in

(24) Four philosophers carried three books.

⁵A survey of a number of issues related to plural and quantifiers can be found in (Lønning 1997). See also (van der Does 1995) on different attempt to locate the source of the readings.

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

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

If a sentence contains a semantically underspecified item in it then an utterance of that sentence will have more than one reading. It will be possible to utter it with at least two different truth conditions. This is like ambiguity, but different in certain important respects from (standard cases of) ambiguity. In the case of ambiguity, too, one can utter a phonetically identical sentence with different truth conditions. But in the case of semantic underspecification we are dealing neither with lexical ambiguity, nor with structural ambiguity. It isn't the case that any one of the words in (22) has two different meanings, nor that the sentence can have two different structures. Rather, some of the items in the sentence are not specified completely. The context of the utterance will have to fill in the details that the language left out.

Reciprocals

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

(25) The Smiths like each other.

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

- (26) The people in the room were no further than one yard from each other.
- (27) The exits on the Santa Monica freeway are no further than one mile from each other.

In the case of (26) can be required for this utterance to be true that everyone in the room could touch everyone else by extending their arm. But in the case of (27) it would still be true even if the first and the last exit are 30 miles apart. All that is required here is that there is another exit every mile.⁶

1.3.2 Back to quantifiers

What these examples show is that the role of context in determining content is much more subtle than simply to contribute the value of indexicals and demonstratives or to contextually restrict quantifiers. Context can make a contribution to content even when there are no overtly context sensitive elements in the sentence uttered. The simplest way to describe this is that certain words and phrases are semantically underspecified: the contribution that the shared language makes to the content of the utterance is only part of the contribution that the utterance of the word or the phrase makes to the content. The rest is supplied by the context. Certain features of the situation of the utterance make it such that one of the other options that the contribution from language allows is picked. To be sure, this is only a rather simplified way of describing a complex phenomenon, but I think it is good enough for now. It allows us to see that the role that context plays in the determination of content is more complex than what the above simple picture seems to suggest. In particular, if the role that context plays in the determination of truth conditions is more complex than to fill in the values of demonstratives and indexicals, and to contextually restrict the domain of quantification, maybe the difference between Quinean and non-Quinean quantifiers, or at least the one's we looked at above, isn't in the ballpark of contextual restriction of quantifiers. In fact, this is exactly what I want to argue for next.

In the remainder of this section I will motivate that quantifiers, too, are semantically underspecified. They do play two different but closely related roles in communication. On different occasions one and the same quantifier can make different contributions to the truth conditions. And only on one of these occasions is quantification closely related to ontology.

I think we can see what the difference between Quinean quantifiers

⁶See (Dalryple et al. 1998) for a discussion of such cases.

and (some) non-Quinean quantifiers really is when we look at what use we have for non-Quinean quantifiers in ordinary everyday communication. It is particularly helpful to look at these situations of communication, and not primarily at metaphysical debates and the role of quantifiers there. If we can find a use for non-Quinean quantifiers in ordinary everyday communication, and if we can find out what the function of these quantifiers is in these uses then this should shed light on the use of non-Quinean quantifiers in ontology and metaphysics, too.

In the remainder of this paper I will propose an account of what the function is of some non-Quinean quantifiers (I can't claim to have an account of all uses of quantifiers that are not Quinean quantifiers). According to this account such quantifiers have a real function in ordinary communication. However, if we look at what this function is, and if we have the issues about context from above in mind, we can see that this is more plausibly accounted for in a non-Meinongian framework. This is not to say that I will argue that Meinongianism is incoherent or absurd. Not at all. But Meinongianism is motivated in part by a mistaken view about the function of quantifiers. In particular, the motivation for (I) essentially relies on this view about quantification.

1.3.3 Communicative functions of quantifiers

To see whether or not we should accept (I) we have to understand better what we do with quantifiers. What functions do they have in communication? What are we doing when we take recourse to them?

There is one thing we do with them which I take to be quite uncontroversial. In this use we have for quantifiers we make statements whose truth depends on what things there are out there in the world. Among statements of this kind I take to be ordinary utterances of:

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

(29) Everybody is hungry, let's take a break.

These are the uses of quantifiers where the truth of these statements depends on what things there are in our domain of discourse, or what things make up reality. Statements with quantifiers in these uses in them impose a condition on the domain of discourse for them to be true. These statements are only true if the domain of discourse satisfies that condition (like containing a thing which eats my cheese). Because of that I will call this reading of quantifiers the **domain conditions reading** or also the **external reading**.

There is however also another, different use we have for quantifiers. I think the believers in non-existent objects make the mistake of collapsing this use of quantifiers into the domain conditions use. Let me explain.

A situation where we have to take recourse to strong expressive power that our language offers us is when we communicate information that is lacking in a certain respect. This is well known from the discussion about the function of a truth predicate. Among the needs we have for a truth predicate is the one to communicate information that is “incomplete” in a certain sense.⁷ To mention the standard example, suppose I believe everything the Pope said. If I know what the Pope said then I can communicate this without using a truth predicate, as with

- (30) The Pope said that p and I believe that p , and the Pope said that q , and I believe that q , etc. (and that’s all the Pope said).

But if you do not know what the Pope said you can’t put it that way. You will have to say something like

- (31) Everything the Pope says is true.

Thus having a truth predicate gives one increased expressive power, and one might think that this is why we have one in the first place.

A similar situation occurs with quantifiers. To do this, let’s consider an example where we have “incomplete” information, but we know exactly in which respect the information we have is lacking: the case of forgetting, but remembering that one forgot.

Suppose you have to write a psychological profile of Fred. One day you learn the most valuable information that Fred is a big admirer of Clinton. This is most useful to you since now you know a lot more about what kind of person Fred admires, what character traits he values and so on. You note this to yourself:

- (32) Fred admires Clinton very much.

The next day, however, you can’t recall who it was that Fred admires so much. You do remember that you knew yesterday, and that it was most useful information to you. But now you just can’t recall who that was. But you didn’t forget everything you knew yesterday. You still know that whoever that was, he is also admired by many Democrats. This, again, is still very useful information to you, since it allows you to connect Fred in a certain way to Democrats. And you can still express and communicate the information you now have. You can say:

- (33) There is someone Fred admires very much, and that person is also admired by many Democrats. I just can’t recall any more who that is. . .

⁷The information is incomplete in an epistemic sense, a sense where the person who is communicating it finds it lacking in a certain way and would like to be more specific, but for some reason or other isn’t able to.

This will communicate the information you still have.

Now, this situation is completely general. Nothing hangs on that who was admired is Clinton. What you might have learned about Fred, and what might have been just as important and useful to you, is that

(34) Fred admires Sherlock Holmes very much.

Again, this allows you to conclude all kinds of things about Fred, about the character traits he values and the like. And again, the next day you might have forgotten who it was that Fred admires so much, but you still remember that whoever it was is also admired by many detectives. This, again, is still useful information, even though not as good as what you knew before. Now it is lacking a certain part, namely who it is that Fred admires so much. However, you can still communicate the “incomplete” information you have by saying something like

(35) There is someone Fred admires very much and who is also admired by many detectives. I just can’t remember who that is any more. . .

We need quantifiers in these situations, when we want to communicate information that is missing certain parts. And in situations like the above, where the speaker doesn’t know what the part is that is missing, but only that there is a missing part, one has to use a quantifier to communicate the information one still has.

This situation is one where quantification is necessary to communicate the information we want to communicate, but the only instances of the quantifier might be things that don’t exist, like Sherlock, in the above case. In this sense it can be said to be quantification ‘over’ non-existent objects. This is an example where such quantification is necessary to communicate the information we want to communicate.

One might think that this plays in the hands of the Meinongean, since it only shows that after all we do have to accept quantification over non-existent objects, and therefore the Meinongean ontology. But that would be too fast. Whether or not the truth of such quantified statements brings with it the Meinongean ontology is just what we have to look at here. And to do so we have to see what the communicative function of such quantifiers is. What are they supposed to do on these occasions?

Let’s first have a brief look at what we use the quantifier for in the above situation. In the above situation we lost a certain part of the information we had before. The rest that we still remember is most useful and we want to communicate it. But it isn’t clear what we can do with the part that was lost. We can’t just say

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

That just isn't grammatical. We have to put something in the place of the forgotten part. And what we have to put in place of the forgotten part has to be neutral with respect to what information the original, forgotten, part contributed. And this is exactly what the quantifier does for us on these occasions. It replaces the forgotten part, and together with the pronoun "who" it makes sure that the truth conditions contain that whoever is admired by Fred is the same as who is admired by many detectives. Whether or not that thing is real isn't what matters here.

If we look at what feature the quantifier must have so that it can do this we can see that it has to have a certain inferential role. The information we originally had, and represented with the use of a singular term, has to imply the representation of the 'incomplete' information that we represent with a quantifier. In other words, $(\dots t \dots)$ has to imply $(\dots \text{something} \dots)$, and it has to do this whether or not 't' stands for something real. As we saw above this is not what matters in these situations. Whether or not Fred admires someone real or unreal isn't what is of issue here.

When we use quantifiers to communicate incomplete information we use them for their inferential role. Let's call this reading of quantifiers their **internal reading** or **inferential role reading**. How this reading relates to the external reading will be what we will have to look at next.

1.3.4 Inferential role and domain conditions

Whether or not expressions that have as their semantic function to impose conditions on the domain of discourse have a certain inferential role in a certain language depends on a number of factors. In the simplest case, when every object in the domain of discourse is such that some term in the language stands for it, and if in addition every term in the language stands for some object in the domain of discourse then inferential role and domain conditions do coincide. A phrase that has a certain inferential role will make the same contribution to the truth conditions as one that imposes a certain domain condition.⁸ If every term in our language stands for some object in the domain of discourse then imposing the domain conditions will still get one the inferential role.

If we grant that we have a need for both inferential role and domain conditions then we have at least two ways of accounting for their relationship. One is to say that in the uses where we use quantifiers for their inferential role they do not make a different contribution to the truth conditions than when we use them for their domain conditions, and the other option is to say that they do. I will discuss both of these in this

⁸I am assuming for now that inferential role is had because of contributions to the truth conditions.

order.

Imposing domain conditions would get a phrase a certain inferential role if our language were a certain way. Roughly, if it wouldn't exhibit partiality (and contain intentional verbs, or the like). By partiality I simply mean that some of the terms or names in our language stand for nothing whatsoever. If our language exhibits partiality then inferential role and domain conditions go apart. The inference from $(\dots t \dots)$ to $(\dots \text{something} \dots)$ wouldn't be valid any more if the quantifier were read externally. Thus a quantifier can only do both, impose domain conditions and occupy a certain inferential role, if partiality is not true about our language.

Meinongeanism can be seen as defending the claim that partiality is not true of our language because the domain of quantification is larger than we thought, it is so large that it contains an object for every one of our terms. And the motivation for this, as we saw above, was to first point to uses of quantifiers that apparently range over non-existing objects, and secondly to provide a restricted quantification analysis of the difference between Quinean and non-Quinean quantifiers. We have seen that we need quantifiers to range over non-existent objects in ordinary everyday situations of communication, at least on the sense of ranging over non-existent objects in which the instances of such quantified statements involve terms that do not stand for any existing objects (whether they stand for nothing at all, or non-existing objects). The Meinongean can account for the inferential role by claiming that it is really a case of imposing domain conditions on a larger domain, assuming that partiality isn't true of our language. But this does not seem to be so plausible. It seems that there are several *prima facie* good reasons for claiming that partiality is true, i.e. that there are names or other terms in our language that stand for nothing whatsoever. The two most prominent are:

- Whatever the mechanism of reference is, however our words manage to stand for objects, they can break down. For one, we are only fallible creatures, and however we manage it that our words stand for things out there in the world, we can fail in cases. So, even in cases where we try to talk about regular existing and concrete things there might be some errors involved that make the referential connection break down, and we end up talking about nothing.
- Sometimes we just make up stories and we do not even try to talk about anything in reality. To be sure, we might do something that we aren't even trying to do, but it sure enough would be a miracle

if reality contained all the things already that we might make up, and that we end up talking about real things even when we aren't trying.

To be sure, this is not intended as a refutation of Meinongianism. In this paper I only want to question the motivation for Meinongianism, in particular its reliance on (I), not to refute Meinongianism. The above cases should make it plausible that there are empty names in our language, and thus that partiality is true. In addition, the reliance on empty names is in no way necessary for present purposes. We can give similar examples as the above ones using non-denoting descriptions. If we assume that no single person invented the wheel we can give an example like the above one using:

- (37) There is someone Fred admires very much, and who is also admired by many bicyclists, namely the inventor of the wheel.

So, it seems that there is good reason to believe that our language does contain empty names, and that thus domain conditions and inferential role go apart. Imposing a certain domain condition won't get one a certain inferential role. But apparently we need and use quantifiers for both. How can that go together?

1.4 A non-Meinongean proposal

We have seen that we need quantifiers for their inferential role, and for their domain conditions. We have also seen that it is plausible to assume that these two do not coincide with respect to truth conditions in a language like ours, even though they do coincide in simpler languages and are closely related. All this can be accommodated in the following simple theory:

Quantifiers, like many other natural language expressions, are semantically underspecified. They make different contributions to the truth conditions on different occasions, depending how they are used and other contextual features. On one of the ways in which they can contribute to the truth conditions they will have a certain inferential role, on the other they will impose certain domain conditions. This is completely consistent with an endorsement of partiality. The contribution that a quantifier makes in its internal, inferential role reading, does not have to reduce to the contribution it makes in its external reading. These are two different but related semantic functions that quantifiers have. As we have seen, it is a mistake to think that the only contextual difference between two uses of quantifiers is that of contextual restriction. That this is the only difference between them was suggested by the simple picture about the role of context in determining content, but we have seen that

this is a mistake. Content plays a much more complex and elusive role, and in particular, there are lots of expressions that are semantically underspecified in the sense that they make different contributions to the truth conditions on different occasions of utterance.

The view that contextual contributions come down to fixing the values of demonstratives and indexicals, and to contextually restricting quantifiers, naturally suggests itself from the background of first order logic. There both of these phenomena are prominent and can be nicely captured within the language and semantics of first order logic. However, the technical tools appropriate for natural languages will go beyond this, and the examples given above, and many others, suggest that taking first order logic as ones ideal is no ideal worth having.⁹

There is, of course, the issue what the truth conditions of the quantifier in its internal role is. In particular, how can the quantifier get the inferential role it is supposed to get by having certain truth conditions. The Meinongians have an answer to this, but as we saw the answer seems to make an implausible assumption about our language. But the Meinongian's answer isn't the only answer. Quantifiers do not have to impose domain conditions as their contribution to the truth conditions. When they are used for their inferential role it is implausible that they get their inferential role through imposing domain conditions, since this would only work if our language had features it doesn't seem to have. So, what truth conditions would give a quantifier the inferential role we want it for? I won't get into the details of this here,¹⁰ but what we have to find out here is simply what contribution to the truth conditions would give an expression the inferential role for which we want the quantifier. In our case, what contribution to the truth conditions would make the inference from (... *t* ...) to (... something ...) valid, whether or not *t* stands for anything real? There are many different contributions to the truth conditions that the quantifier could make so that it would have this inferential role. The simplest one is for it to make a contribution such that the statement with the quantifier in it is truth conditionally equivalent to the disjunction of all the statements of the form (... *t* ...) that imply it. To be sure, there are infinitely many of such statements, but if the quantifier would make such a contribution to the truth conditions then the statements in which it occurs would have this inferential role. This is the simplest and most trivial way to get a certain inferential

⁹That first order languages have their limits in capturing natural language quantification, even in cases completely unrelated to our debate here, has been widely discussed. See for example (Barwise and Cooper 1981) or the section on generalized quantifiers in (Gamut 1991) and the references given there.

¹⁰However, I do get into the details of it in (Hofweber 1999), chapter 2.

role. To say this is, of course, not to say that the underlying structure of such sentences is in any way infinitary. This relates in certain ways to substitutional quantification. A substitutional interpretation in effect gives the quantifier a certain inferential role. But the present proposal differs from substitutional quantification in a number of ways.¹¹

The present proposal can be summed up with the following points:

- Quantifiers like “something” are semantically underspecified and make different contributions to the truth condition of an utterance on different occasions of the utterance. Quantifiers are just one of many items in our language that have this feature.
- The different contributions to the truth conditions that quantifiers can make correspond to different functions they have in communication: occupying a certain inferential role and imposing certain domain conditions.
- In languages like ours inferential role and domain conditions do not coincide with respect to truth conditions.
- It shouldn’t be surprising that one and the same item has these different functions on different occasions, since they are closely related and in fact coincide in simpler languages.

1.5 Solving the dilemma

The present proposal nicely resolves the dilemma that we started out with at the beginning of this paper (see page 4). It seemed that on the one hand answering the ontological question about non-existent objects would involve figuring out the truth value of statements whose quantifiers range over non-existent objects, and this seems to be a substantial and difficult task. But on the other hand it seems to follow quite trivially from the fact that Santa doesn’t exist that such statements are true. According to the present account both sides have some truth to them.

- On the one hand it is indeed trivial to conclude that there are non-existent objects from nothing more than the premise that Santa doesn’t exist. The sense in which this is trivial is that if the quantifier is used in its inferential role reading then it trivially follows from “Santa doesn’t exist” that “There is something which doesn’t exist, namely Santa.” And from “Santa is a non-existent object” it follows trivially that “There are non-existent objects,” again assuming that the quantifier is read internally.

¹¹Again, see (Hofweber 1999) chapter 2 for the details of this.

- On the other hand, it is a substantial metaphysical issue to decide whether or not there are non-existent objects, using the quantifier externally. The speculative ontological issue is not decided by the trivial inferences. Whether or not there are non-existent objects, using the quantifier externally, is left open by the endorsement of the trivial arguments. To answer this question is indeed to engage in speculative ontology.

If what I said so far is right then (I) is wrong, and the Meinongean's motivation for an ontology of non-existent objects is too quick. There are true statements with quantifiers in them such that the only true instances of the quantifiers are with names or terms that do not stand for anything that exists. But even though this is true, it doesn't answer the question whether or not there are non-existent objects in our ontology. This would only be answered if we had a true statement with a quantifier ranging over non-existent objects in it, and that quantifier is used in its external reading. We thus avoid the modus ponens—modus tollens dilemma from page 4 by denying the conditional that gave rise to it.

The Meinongean's view that the difference between Quinean and non-Quinean quantifiers is one of contextual restriction of the quantifier is based on a too simple model of the role of context in determining truth conditions, a view that is based on assuming a too close connection between natural languages and first order languages. The difference between Quinean and some non-Quinean quantifiers is in the ballpark of semantic underspecification and the difference between inferential role and domain conditions. And that inferential role is not had because of imposing certain domain conditions is a plausible fact about our language.

1.6 Objections

Before we conclude I'd like to add a brief discussion of two objections that can be raised against the present proposal. One deals with the fact that I have not said enough about the issue of quantifier scope and how it relates to the use of quantifiers together with intentional verbs. The other deals with the claim that when we use the word 'exists' in conjunction with quantifiers we thereby explicitly restrict the quantifier, thereby giving support to the contextual restriction picture.¹²

1.6.1 Quantification and scope

One might object that I did not pay enough attention to the difference between

¹²More details are in (Hofweber 1999) chapter 2.

(38) There is someone whom Fred admires and who is also admired by many detectives.

and

(39) Fred admires someone who is also admired by many detectives.

The latter, the objection continues, is completely compatible with the claim that we never have to quantify over non-existent objects (contrary to what Meinongians and I claim), because the quantifier is within the scope of the intentional verb. Only in the case of (38) do we quantify over non-existent objects.

To be sure, I have not talked much about this difference, but there is good reason for this. Without a doubt, there is an issue of scope when quantifiers interact with intentional verbs, but I think it is a mistake to closely connect the different scope reading that such quantifiers can have with how these items are arranged in the sentences that get uttered. It seems to me that both sentences, (38) and (39), can have both scope readings, and that using one or the other is rather in the ballpark of topicalization. The difference will be apparent when we notice how we would continue a conversation after (38) or (39). In the first case we would usually continue to talk about the person, whoever it is, who is admired by Fred. In the second case we would normally continue to talk about Fred.

It is a common strategy in philosophy to claim that quantifiers that apparently range over non-existent objects can only occur within the scope of an intentional verb. I think it is a mistake to think that we only have use for quantifiers this way. But I won't be able to defend this here in sufficient detail.¹³ This strategy of arguing is, of course, not a defense of Meinongianism. The option defended in this paper is to accept certain quantified statements as true, but to give a different account of the function of these quantifiers on these occasions than a Meinongian would give.

1.6.2 The role of “exists”

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

¹³See, however, (Hofweber 1999) chapter 2.

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

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

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

(41) Four philosophers wrote a book.

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

(42) Four philosophers together wrote a book.

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

(43) Four philosophers each wrote a book.

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

1.7 Quantification and ontology

The present view about contextual differences in different uses of quantifiers has a relevance for ontology beyond the debate about non-existent objects. In fact, if what I said here is right then what is commonly called Quine's criterion for ontological commitment has to be false. We are not committed to all the things our quantifiers range over, even in our best theories. We are only committed to what the quantifiers used externally range over. The more positive alternative view that arises from all this is in certain respects similar to the view Carnap endorsed in his "Empiricism, semantics, and ontology" (Carnap 1956). There will be a difference between internal and external questions about what there is, based on whether or not the quantifier used in asking the question is used internally or externally. And in accordance with Carnap's view the internal questions (about numbers, properties and the like) will have trivial affirmative answers. The external questions, however, won't have trivial affirmative answers. Contrary to Carnap, however, I do think that these external questions are fully meaningful.¹⁴

1.8 Conclusion

We started with a dilemma about whether or not to accept that there are non-existent objects is something trivial, or metaphysically substantial. And we saw that a central claim in the debate about the ontology of non-existent objects is the acceptance of (I). Given (I) there only seem to be two options one has: accept quantification over non-existent objects and an ontology of non-existent objects, or reject both of them. We have seen that a central part of the motivation for (I) is the contextual restriction view of the difference between Quinean and non-Quinean quantifiers. I have argued that once we look at the role of context in general and at what the function of certain non-Quinean quantifiers is in ordinary communication then we can see that these are not our only options. There are uses of quantifiers where they are non-Quinean quantifiers, but they are not restricted external quantifiers. It is a mistake to think that the only contextual difference between different uses of quantifiers is one of contextual restriction. The much more broad thesis of semantic underspecification applies to quantifiers as well as many other expressions, and in particular gives us the distinction between internal and external uses of quantifiers. With this distinction in mind we can see that (I) is mistaken. And we can resolve the dilemma we started out with: it is trivial that there are non-existent objects, if the quantifier

¹⁴For more on all this see (Hofweber 1999), (Hofweber b), (Hofweber a), and (Hofweber c)

is understood internally. But it is a substantial metaphysical issue to decide whether or not there are non-existent objects, if the quantifier is understood externally. Whether or not non-existent objects are part of reality is left open by everything I have said here. But if I'm right then it will be harder to argue that there are than Meinongians assume it is.

References

- Barwise, J., and R. Cooper. 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4:159–219.
- Carnap, R. 1956. Empiricism, Semantics, and Ontology. In *Meaning and Necessity*. Chicago: University of Chicago Press.
- Dalryple, M., M. Kanazawa, Y. Kim, S. Mchombo, and S. Peters. 1998. Reciprocal Expressions and the Concept of Reciprocity. *Linguistics and Philosophy* 159–210.
- Gamut, L. T. F. 1991. *Logic, Language, and Meaning*. University of Chicago Press.
- Hofweber, T. n.d.a. Innocent Statements and their Metaphysically Loaded Counterparts. unpublished manuscript.
- Hofweber, T. n.d.b. Number Determiners, Numbers, and Arithmetic. unpublished manuscript.
- Hofweber, T. n.d.c. A Puzzle about Ontology. unpublished manuscript.
- Hofweber, T. 1999. *Ontology and Objectivity*. Doctoral dissertation, Stanford University.
- Lønning, J. T. 1997. Plurals and Collectivity. In *Handbook of Logic and Language*, ed. Johan van Benthem and Alice ter Meulen. Elsevier.
- Parsons, T. 1980. *Nonexistent Objects*. New Haven: Yale University Press.
- van der Does, J. 1995. Sums and Quantifiers. *Linguistics and Philosophy* 16:509–550.
- van Inwagen, Peter. 2000. Quantification and Fictional Discourse. In *Empty Names, Fiction, and the Puzzles of Non-Existence*, ed. A. Everett and T. Hofweber. Stanford: CSLI Press.

Index

context, 7–13

genitive, 9–10

Meinong, Alexius, 1

non-existent objects, 1–25

ontology, 1

Parsons, Terence, 6n

plural, 10–11, 23–24

quantification

 and ontological comittment, 24

quantification, 3–8, 12–25

reciprocals, 11–12