Context and Unrestricted Quantification*

Michael Glanzberg

University of California, Davis

Quantification is haunted by the specter of paradoxes. Since Russell, it has been a persistent idea that the paradoxes show what might have appeared to be absolutely unrestricted quantification to be somehow restricted. In the contemporary literature, this theme is taken up by Dummett (1973, 1993) and Parsons (1974a,b). Parsons, in particular, argues that both the Liar and Russell's paradoxes are to be resolved by construing apparently absolutely unrestricted quantifiers as appropriately restricted.

Building on Parsons' work, I have advocated a *contextualist* version of the view that there is no absolutely unrestricted quantification (Glanzberg 2001, 2004a,b). I have argued that all quantifiers must be construed as ranging over contextually provided domains, and that for any context, there is a distinct context which provides a wider domain of quantification. Hence, there is no absolutely unrestricted quantification. Instead, quantification displays a contextual version of what Dummett calls 'indefinite extensibility'. With Parsons, I have argued that this helps us to resolve the Liar as well as Russell's paradoxes.

There remain a great number of issues surrounding the sort of view Parsons and I advocate. Just how to understand the argument from paradox against absolutely unrestricted quantification

remains a delicate matter. Questions about how our view might be coherently stated, and whether it is compatible with certain ideas in metaphysics, are often raised. (Many such questions are raised forcefully in Williamson (2004).) I take these sorts of issues seriously (and have tried to address some of them in my 2004*b*). In this essay, however, I shall put them aside, in favor of developing important positive aspects of the contextualist proposal.

The rejection of absolutely unrestricted quantification is no doubt an unexpected, if not unwelcome, conclusion. In part, I believe, we have to recognize that any genuine solution to the paradoxes will force some unwelcome conclusion upon us. (We would hardly see a genuine paradox otherwise.) But I have also suggested that the contextualist view mitigates the unwelcome effects, and so offers a better-motivated, less ad hoc approach. It does so by drawing a parallel between the sorts shifts in quantifier domains required by the paradoxes and the very familiar phenomenon of *contextual quantifier domain restriction*. The more close the parallel, the more we can see what might have looked like an unexpected and unwelcome restriction on what we can say as merely an unusual manifestation of a familiar and wide-spread natural-language phenomenon.

In this essay, I shall investigate how close this parallel really is. I shall argue for a limited, but still substantial, conclusion. The kind of quantifier behavior we see with the paradoxes is not exactly the same as the more ordinary kind we see in everyday discourse. Even so, I shall argue it is importantly similar. Like ordinary quantifier domain restriction, we can understand it as the setting of a context-dependent parameter in a sentence. The parameters involved in paradoxical and ordinary cases are distinct, but the pragmatic *processes* which set them are fundamentally related. Thus, I shall argue, examining the pragmatics and semantics of quantifier domain restriction does provide us with important insights which may be applied to develop the contextualist response to the paradoxes in detail.

I should make clear at the outset that I shall not be offering a new argument against unrestricted quantification, nor am I claiming that we can resolve the paradoxes simply by observing commonplace linguistic phenomena. Rather, I shall be developing the contextualist view, by supplementing familiar arguments derived from the paradoxes by examination of ordinary

contextual domain restriction in natural language. Providing these details and establishing where and how they relate to well-established phenomena in natural language will, I hope, show the contextualist view to be plausible and well-motivated. This will offer indirect support for the view.

This essay proceeds as follow. The kind of paradoxical reasoning that the contextualist proposal takes as its starting point is reviewed in section 1, while the basic form of the contextualist response is outlined in section 2. Section 3 discusses some ideas about the semantics and pragmatics of quantifier domain restriction, focusing on ordinary cases of contextual domain restriction. Sections 4–6 apply these ideas towards a systematic development of the contextualist proposal. Some concluding remarks are offered in section 7.

1 The Paradox Reviewed

To set the stage, let us rehearse the kind of paradoxical reasoning that will be our main concern. A number of paradoxes will make the point, including forms of the Liar and Russell's paradoxes. I shall present a very elegant and highly general version of Russell's paradox due to Williamson (2004), which will serve as a good illustration of the phenomenon at issue.

Williamson invites us to consider the task of building an interpretation I for some language. As he notes, we need to say little about the formal properties of I, except that for a given predicate term 'P' and collection F, we may build an interpretation $I(F)$ which makes 'P' hold of all and only the Fs. We can remain quite neutral on the nature of the Fs as well. They need not form a set or a class, and it appears we can simply appeal to Boolos-style plural constructions (e.g. Boolos 1984) to describe them.

The crucial observation is that though interpretations need not be sets or classes, it appears we can talk about them. As Williamson notes, we do so naturally when we investigate logical consequence, for instance. In talking about interpretations, we recognize them as objects of some kind. Once we have this, we have the basics of a Russell-like construction. Let the Rs be all and only the objects o such that o is not an interpretation under which 'P' applies to o. Then there is an interpretation $I(R)$. But the object $I(R)$ itself cannot be in the domain of the quantifier *all and*

only the objects just used (nor can it be in the domain of any quantifier as interpreted by $I(R)$). If it were, we would have in the case of $o = I(R)$: $I(R)$ is an interpretation under which 'P' applies to o iff $I(R)$ is not an interpretation under which 'P' applies to o. This is a contradiction, just as in the usual version of Russell's paradox.

The response to this and other paradoxes I favor simply says that the object o must not have been in the range of the apparently unrestricted quantifier *all objects*. Indeed, it *cannot* be, on pain of logical contradiction. Thus, we must see even quantifiers like *all objects* as in important ways restricted. Of course, Williamson himself disagrees. He holds that we cannot really recognize $I(R)$ as an object, as it is fundamentally predicative or second-order. But as I mentioned above, my goal here is not to engage in this dispute directly, but rather to investigate some of the details of my favored response.

One of the very nice features of Williamson's version of the Russell argument is that it shows the object we find outside of the domain of an occurrence of *all objects* need not be a set, or class. The argument makes extremely minimal assumptions about the kind of object in question. But if we do help ourselves to some of the basic features of set theory, we can state the problem we face more simply. Relying on some simple set-theoretic reasoning, we can conclude from Williamson's construction that the interpretation of a language in a given stretch of discourse cannot itself be in the domain of any quantifier as used in that stretch of discourse. Some further set-theoretic reasoning allows to reduce this to the fact that the domain of the widest quantifier in a stretch of discourse—the domain of *all objects*—cannot itself be an object in the domain of that quantifier. As usual, we turn the Russell paradox into a proof that there is no universal set.[1]

Though Williamson's argument shows them not to be crucial, I shall make these set-theoretic assumptions for simplicity's sake in what follows. With these assumptions, we can work with the more familiar 'no universal set' form of Russell's paradox. We start with a quantifier *all objects*, recognize its domain as itself an object, and conclude on pain of contradiction that this object

---

[1]Of course, the axiom of foundation already tells us no set contains itself. There is a responses to the Liar developed in non-well-founded set theory by Barwise and Etchemendy (1987), but this is already a *contextualist* response. (I discuss the relation of my preferred view to theirs in my 2004a.)

cannot be in the domain of the quantifier. We thus find, I maintain, that our apparently absolutely unrestricted quantifier *all objects* is somehow restricted after all.

2   A Contextualist Proposal

So far, we have observed that given any quantifier domain, it is possible to build an object which does not fall under that domain, via directions provided by familiar paradoxes. With some set-theoretic assumptions, we may say that the quantifier domain itself is an object that cannot be among its own members.

Call this the *argument from paradox*. What does this argument really show? Though the issue is contentious, my starting point for this essay is that the argument shows no quantifier can range over 'absolutely everything'. To fix some terminology, say that an *absolutely unrestricted quantifier* is one that ranges over a fixed domain of 'absolutely everything'. An *absolutist* holds that there are absolutely unrestricted quantifiers. I shall assume the argument from paradox shows that absolutism is untenable.[2]

Rejecting absolutism has appeared implausible to a number of authors (e.g. Cartwright 1994; McGee 2000; Williamson 2004). To many, some form of contextualism has seemed to be the best way to address this worry. The basic contextualist idea is to see the reasoning in the argument from paradox as showing that even quantifiers like *all objects* range over contextually extensible domains. We start with a domain for a quantifier like *all objects*, and via the paradoxes identify an object not in that domain. This causes the context to change, to a new context in which *all objects* ranges over a strictly wider domain including the new object we discovered.

The contextualist holds that this conclusion is not so implausible as it might seem, because in fact contextual restrictions on quantifier domains are the norm in natural language. For instance, consider:

(1)   a. Most people came to the party.

---

[2]I hasten to repeat that the assessment of the argument from the paradox is a delicate matter. I have discussed it at greater length in my (2004*b*), and it is pursued with great subtlety by Williamson (2004), and in the papers by Fine (this volume) and Parsons (this volume).

b. Every bottle is empty.

In each, we interpret the quantifier as contextually restricted. Precisely how will depend on the context. Roughly, (1a) says that most people in the contextually salient domain came to the party, e.g. most people among my friends and colleagues. (1b) does not say that every bottle in the world is empty; rather, it is understood as saying that every bottle in the contextually salient domain is empty, e.g. the bottles near the door waiting to be taken outside.

(1a) and (1b) uses the already restricted *most people* and *every bottle*, and provide further contextual restrictions. We see the same thing with *syntactically unrestricted* quantifier: quantifiers like *everything* or *nothing* which bear no non-trivial overt restricting predicate. For instance:

(2)   a. I took everything with me.

      b. Nothing outlasts the energizer.

(2a) says that I took everything in the contextually salient domain, e.g. everything I had brought with me; while (2b) says that nothing of the contextually relevant kind outlasts the energizer, e.g. batteries.

The contextualist seeks to bolster the anti-absolutist position by noting that we already have good reason to see most quantifiers in natural language as ranging over contextually restricted domains, whether they appear overtly restricted or not. The contextualist then argues it is a small step to conclude that all quantifiers are so contextually restricted. There is no way to step outside the normal contextual restrictions on our quantifiers, the contextualist holds, and assuming we can leads to paradox.

It is the task of this essay to see how far this idea can be pressed. But some initial points about how best to formulate the contextualist proposal should be addressed before we jump into the details of contextually restricted quantifiers. The first is to clarify what we mean by talking about restricted or contextually restricted quantifiers. To regiment some terminology, let us reserve the term *restricted quantifier* for one which contains a *syntactic restrictor*: a predicate either pronounced or unpronounced but present in the underlying syntax, which restricts the domain of a quantifier. It will be convenient to further require that the restrictor be non-vacuous. Restricted

quantifiers are thus syntactically restricted, and the restrictor position is filled with a non-vacuous predicate. Each quantifier in (1) is clearly restricted.

By this definition, the quantifiers in (2) are not restricted, as predicates like *thing* are semantically vacuous. But as they occur in their intended contexts, they range over subdomains of objects we can talk about in their contexts, and we can assume that speakers will intend them to range over such restricted domains. Let us say a quantifier is *contextually restricted* if its contextually fixed domain is a subdomain of the objects available for quantification in a given context. All the examples in both (1) and (2) are of contextually restricted quantifiers, though the ones in (2) are unrestricted but contextually restricted, whereas the ones in (1) are both restricted and contextually restricted.

The quantifiers that figure into the argument from paradox, like *all objects* or *everything*, are clearly unrestricted according to this terminology. It is also highly plausible that they are contextually unrestricted. As speakers use them at the beginning of the argument, they do not intend them to range over any proper subdomain of objects they can talk about. When the speaker says *all objects* in such a context, she does not mean 'all objects that are F', or 'all objects except those in Y', or anything else that would indicate contextual restriction.[3]

To make clear how there might still be room for a contextualist response to the argument from paradox, one more definition is needed. Let us say the *background domain of a context* is the widest domain of quantification available in a given context. This will be the domain of all objects, according to a given context. It is thus the domain over which unrestricted and contextually unrestricted quantifiers range. As far as a single context is concerned, the background domain is simply 'everything', and will not be a proper subset of any other quantifier domain or predicate extension available in that context.

The contextualist response to the paradox is the view that there is contextual relativity of

---

[3]Some interesting idea of Rayo (2003) might be used to explain how some unrestricted quantifiers wind up contextually *un*restricted. Of course, by my lights, Rayo's ideas will have to re-cast in terms of setting quantifiers to range over the background domain of a context, rather than 'absolutely everything', but I think they could serve very nicely in that role.

background domains. Whereas the absolutist holds there is one fixed background domain, which is simply 'absolutely everything', the contextualist holds that different contexts can have distinct background domains. The argument from paradox shows us, given a plausible background domain, how to identify an object not in the domain. The contextualist hold that this leads us to a new context with a strictly wider background domain.

The contextualist position on the quantifiers in the argument from paradox is thus that they are both unrestricted and contextually unrestricted. This is just to say that they range over the background domain of a given context. But they are still *contextually relative* unrestricted quantifiers, in that they range over the background domain of a context, and that is a context-relative domain.[4]

I am distinguishing contextual domain restriction, which carves out a subdomain within a background domain, from the contextual relativity of background domains themselves. The contextualist takes the argument from paradox to show that background domains can *expand* in certain changes in context. Clearly this cannot be understood as relaxing some contextual restriction. Indeed, what the paradox does is show us one specific object that was left out of a background domain, and the shift in context must expand the domain to take it in.[5] To understand this, we need to understand how context can re-adjust to take a new object into a background domain, and how this can affect our uses of unrestricted quantifiers. It is to these matters that we now turn.

---

[4]Fine calls these *relatively unrestricted quantifiers*. I am changing the terminology slightly to emphasize the role of context in rejecting absolutism. It is fair enough to say that contextual relativity is a *kind* of contextual restriction, but it would be cheating for the contextualist to let the terminology conflate cases like (1) and (2) with the ones we are confronted with by the paradox. Thus, I shall apply 'contextually restricted' to the former, and 'contextually relative' to the latter.

[5]The contextualism I am proposing here assumes what Fine (this volume) calls *expansionism*, as opposed to *restrictionism*. Fine argues directly against restrictionism, while I am merely noting that the argument from paradox, and the basic contextualist reply, seem to be naturally expansionist and not restrictionist.

3 The Semantics and Pragmatics of Restricted Quantifiers

In this section, I shall investigate how quantifier domains are context-dependent. Much of the focus here will be on the ordinary sort of contextual domain restriction we see in cases like (1) and (2). For cases like these, we may rely on some independently motivated ideas from philosophy of language and linguistics to explain how context interacts with the semantics of quantifiers. With this in hand, I shall identify how context can interact with unrestricted and contextually unrestricted quantifiers as well. I shall also very tentatively explore the pragmatic mechanism by which context sets restricted quantifier domains. Once we have these ideas in hand, I shall go on to apply them to the case of background domains in subsequent sections.

3.1 The Semantics of Quantifier Domains

In this section, I shall begin by reviewing some fairly standard points about how ordinary contextually restricted quantifiers like we see in (1) and (2) work. I shall then turn to how the semantics of quantifiers can make room for background domain relativity.

For reference's sake, I shall adopt a standard generalized-quantifier treatment of the semantics of quantification. This is the semantics of *determiners*: expressions like *every, some, most, few*, etc. The standard theory treats these as *relations* between two sets, representing the contents of nominals and verb phrases. Fix a domain M for an interpretation. Assign nominals like *bottle* and verb phrases like *is empty* subsets of M as their semantic values.[6] For a term α, let $[\![\alpha]\!]^c$ be its semantic value in context c, so $[\![\text{bottle}]\!]^c, [\![\text{is empty}]\!]^c \subseteq M$. Then *Every bottle is empty* (1b) expresses the relation $[\![\text{bottle}]\!]^c \subseteq [\![\text{is empty}]\!]^c$. The semantic value of the determiner *every* is thus the generalized quantifier $\mathbf{every}_M$:[7]

(3)   For every $A, B \subseteq M$, $\mathbf{every}_M(A, B) \longleftrightarrow A \subseteq B$

---

[6]I am assuming an *extensional* semantic framework. As issues of intensionality are not relevant to our concerns here, this can be seen as a mere simplifying assumption.

[7]I am denoting the generalized quantifier which interprets an expression like *every* by the corresponding boldface expression $\mathbf{every}_M$.

Semantically non-trivial nominals, like *bottle*, give us restricted quantifiers.

One of the virtues of generalized quantifier theory is that it provides definitions for many other quantifiers as well. For instance, we can define a value for *most* by:

(4)   For every $A, B \subseteq M$, $\mathbf{most}_M(A, B) \longleftrightarrow |A \cap B| > |A \setminus B|$

The theory of generalized quantifiers, and their application to natural-language semantics, is well developed. For surveys, see Keenan and Westerståhl (1997) and Westerståhl (1989).

Definitions like (3) are relative to a fixed M, and give what are called *local* generalized quantifiers. In logic, of course, we can ask about what would happen if we varied M—varied the universe of discourse. To do this, we need what are called *global generalized quantifiers*. These are simply functions from domains M to local generalized quantifiers on M. So we could set:

(5)   **Every** is the function from M to $\mathbf{every}_M$.

Global generalized quantifiers capture the most general meanings of determiners.

It is tempting to explain the context-dependence of quantifiers we see in (1) and (2) as the result of context affecting M—the background domain—in a global generalized quantifier. As was argued by Westerståhl (1985*a*), this is not right. Westerståhl offers two important principles:

WP1: Background domains are large. Contextually restricted domains can be small.

WP2 Background domains are (relatively) stable across stretches of discourse. Contextually restricted domains are not.

(Westerståhl makes a stronger claim in place of WP2, but the weaker version will be the relevant one for our discussion here.)

To see how WP1 works, consider:

(6)   At the department meeting today, everyone complained about the Governor.

Everyone ranges over members of the department, and excludes the Governor, even though we have to have the Governor in our background domain. Hence, the contextually specified domain of *everyone* is not only quite small, but clearly smaller than the background domain.

To see how WP2 works, consider an example attributed to Peter Ludlow from Stanley and Williamson (1995):

(7)   Nobody cared that nobody came.

Here, on many natural readings, we talk about two distinct domains for distinct occurrences of the same quantifier. Hence, neither can be the background domain.

The moral of these arguments is that ordinary quantifier context-dependence is not the result of M being a context-dependent parameter. Rather, we need an additional contextual restriction on quantifier domains, within whatever background domain we have set. There are a number of different ways to do this, and the details of exactly how will not matter here. For argument's sake, I shall adopt the proposal of Stanley (2000) and Stanley and Szabó (2000), which holds that there is a contextual parameter in the *nominal* of a quantifier. Simplifying somewhat, we make *Every bottle is empty* look like:

(8) $\mathbf{every}_M(D^c \cap \llbracket \text{bottle} \rrbracket^c, \llbracket \text{is empty} \rrbracket^c)$

D^c is a contextually fixed set of elements of M, which restricts the quantifier domain by intersection.[8]

Making the semantics of ordinary quantifier domain restriction explicit makes clear that it will not directly explain the context-relativity of *background domains* which the paradoxes seem to show us. It cannot, as it is a mechanism for restriction within background domains. Even so, the semantics of quantification does show us a way to make room for the context-relativity we need. The semantics of each (local) generalized quantifier already depends on M, which is playing the

---

[8]I am suppressing some further complications in Stanley and Szabó's view. (Incidentally, the specific example (1b) is drawn from their presentation.) Stanley and Szabó (2000) are advancing a particular claim about logical form: that the contextual parameter appears in the nominal position. There are other options. For instance, von Fintel (1994) and Westerståhl (1985a) place the parameter on the determiner, though the resulting semantics is still basically that of (8). A more significantly different option is to deny there is any such parameter in logical form at all, and insist that a purely pragmatic process produces quantifier domain restriction. Representatives of this view include Bach (1994) and Carston (2004). Another option is to deny that there is any context-dependence. This route is taken by Cappelen and Lepore (2002).

role of the *background* domain. If, as the argument from paradox seems to show, M can shift with context, then we may see M as introducing context-relativity into the semantic values of determiners as given in definitions like (3) and (4).

To make this explicit, consider an occurrence of *everything is F*, in which *everything* is not contextually restricted. In such a context, we will have $D^c = [\![\text{thing}]\!]^c = M$. Tracing through definition (3), we find the sentence is true iff $M \subseteq [\![\text{F}]\!]^c$. If M is context-relative, so is the interpretation of the contextually unrestricted *everything*.

As the considerations of section 2 already suggested, this sort of contextual relativity is very different from what we see in ordinary cases of contextual domain restriction. First of all, it is a very different sort of mechanism that introduces dependency upon context. In ordinary domain restriction, it is an independent parameter D^c in the underlying logical form of a sentence. In the case of background domain relativity, it is a feature of the semantics of determiners itself that triggers context-dependence. I shall call this M-dependence, but it must be stressed that M in (3) and (4) works very differently than D^c in (8). D^c is a parameter which gets an independent value, and then composes with other semantic values of a sentence, particularly the semantic value of the nominal. M enters into determining the semantic value of a determiner directly.[9]

This is to say, in effect, that determiners function like *indexicals*. As the class of determiners is rather large, one might object that this conclusion posits massive or open-ended indexicality. It certainly does posit indexicality, but the objection is over-stated. First of all, the class of simplex or 'lexical' determiners is not really so large or open-ended, compared to classes like nouns or verbs. (The determiners form what linguists call a 'closed class', whereas nouns and verbs form 'open classes'.) Perhaps more importantly, it is known from work of Keenan and Stavi (1986) that the semantic values of possible human language determiners can be built up from a very small class of basic determiner values, together with some operations not specific to determiners. Indeed, we can built them inductively from $\textbf{every}_M$ and $\textbf{some}_M$. So it may well be that we only have to posit

---

[9]A global generalized quantifier thus gives the *character* of a determiner, in the sense of Kaplan (1989). The context-relativity of background domains implies that this character is non-constant.

indexicality in a very limited class of expressions to get the results we need.[10]

So far, my proposal is that there are two distinct sources of context-dependence in quantifiers. One—D^c dependence—is familiar and commonplace, and responsible for ordinary contextual domain restriction. The other—M dependence—is responsible for the context-relativity of unrestricted and contextually unrestricted quantifiers. If this is right, then the pressing question becomes how context can fix a background domain: how context fixes M. To try to shed some light on this, I shall begin by looking at how context works to set quantifier domains in the more ordinary cases. In doing so, I shall isolate a few general principles. Along with Westerståhl's principles about background domains, these will help us understand what context must do in the case of background domains.

3.2 The Pragmatics of Quantifier Domains

There are a few general principles about how context sets quantifier domains which we can identify by looking at discourse. One general caveat needs to be mentioned first. There is a significant step between asking how the truth conditions of a sentence vary with context, or what features of a sentence makes it context-dependent, and asking how context itself fixes some context-dependent parameter. By the lights of Stanley and Szabó (2000), the latter is a matter of *foundational* rather than *descriptive* semantics (or pragmatics). Regardless of classification, understanding how context affects content is an important part of our understanding of context-dependence. But foundational issues in pragmatics do tend to get extremely messy, and run into some very hard problems in cognitive science. Explaining what speakers will take to be salient or relevant, for instance, might well involve far-reaching theories of cognition. (To borrow a phrase from Peter Ludlow, foundational problems in pragmatics tend to be 'AI-complete'.)

The best we can do, absent such far-reaching theories, is to stay as close to descriptive matters as we can. Wherever possible, I shall try to isolate relatively clear, well-motivated descriptive

---

[10]There are some more ordinary cases where determiners display indexicality. One is the case of *many* discussed by Westerståhl (1985*b*), which appears to have a value sensitive to a contextual input of 'normal frequency' in much the way that I am positing sensitivity to background domain.

constraints, and then try to apply these constraints to shed some light on the foundational issue of what pragmatic processes are at work in domain restriction. The conclusions I shall come to will be limited, but they will be enough, I hope, to further the comparison between ordinary quantifier domain context-dependence, and background domain context-relativity.

3.2.1 Quantifier Domains by Anaphora on Predicates

One way context works to set quantifier domains, in the ordinary case of contextual domain restriction, is by anaphora-like processes. The domain-restricting parameter can be mapped to some previous material in a discourse, much like a pronoun can (cf. Geurts and van der Sandt 1999; Roberts 1995). We see this in:

(9) There were some passengers <u>on the airplane</u>. Most passengers$_{D^c}$ were killed in the crash.

The domain of *most* is contextually restricted by the predicate *on the airplane*.[11]

 Quantifier domain restrictors prefer to find antecedents in predicative material. For instance, we see a contrast in:

(10) John came to the party and Sarah came to the party.

  a. They had fun.

  b. Everyone had fun.

They in (10a) picks up the aggregate of John and Sarah, while *everyone* in (10b) ranges over people who came to the party. It picks up its domain from *came to the party*.

 Furthermore, it appears that context can construct complex defining properties out of stretches of discourse. For instance:

(11) Susan found most books which Bill needs, but few were important.

---

[11]The example is modified from one of Geurts and van der Sandt (1999), who develop this idea in a DRT-based framework. Gawron (1996) and Roberts (1995) pursue related ideas in the framework of dynamic logic.

The domain of *few* appears to be books which Susan found and Bill needs. Note, this combines relations and terms to form a defining condition.[12]

So far, we have seen that contextually restricted quantifier domains can be set by finding appropriate predicative material in previous discourse. When this happens, it is via predicative material, but that material can be complex, and built out of multiple predicates and singular terms.

3.2.2 Accommodation

In many cases, previous discourse does not provide the needed predicates to restrict a quantifier domain. This is the way we interpreted the examples in (1) and (2), for instance. Even when we do have some predicates available in previous discourse, there is no guarantee we will not need more information from context to tell us how to further restrict the domain.

In these cases, speakers will seek to recover from the context enough information to define the right domain restriction. These days, this is often glossed as a process of *accommodation*, in that it makes the continuing discourse as if the new information had been explicitly uttered.[13] There are some very rough-and-ready rules for accommodation we can state for the case of quantifier domain restriction: add information restricting a quantifier domain to make the current utterance coherent and informative, relative to what is common knowledge in the discourse at the point of utterance.[14]

---

[12]This is derived from an example of Kamp and Reyle (1993), who use it in introducing their abstraction operator for plural anaphora (cf. Sæbø 1999).

[13]The notion of accommodation stems from Lewis (1979). It is common to think of accommodation as a kind of *conversational repair strategy*. Something happens in a discourse which would cause it to break down unless some information where present, and we repair the discourse to make it present.

There are a number of distinct ways of understanding how such a process works. Though Lewis proposed that there are distinct rules of accommodation, it can be thought of as flowing from general pragmatic principles, perhaps along Gricean lines. This is the view taken by Stalnaker (1998), and is in keeping with the discussion in Heim (1983) and Roberts (1995). But in theories like that of van der Sandt (1992), accommodation is understood as a particular kind of operation within a representational theory of discourse like DRT. However, I do not think any of the claims I make here are affected by these distinctions.

[14]These ideas derive from Grice (1975), Stalnaker (1978), and van der Sandt (1992). I have discussed the issue of coherence at length in my (2002).

These rules are rough, and leave the task of accommodation drastically under-described. I shall not be able to elaborate them much more, but there is one point that will be important. Whatever determines a quantifier domain in accommodation cases—what makes the quantifier domain sustain informative and coherent discourse—is not simply a matter of what objects and properties are salient in the immediate environment. We do not accommodate *merely* by checking what is around us. For instance, consider the familiar:

(12) Everything is packed.

Make the context one in which you are about to step out the door to go on a trip. You have your suitcase, and a bunch of things in your pockets or in hand that you will want for the ride. As we normally interpret it in a case like this, the utterance of (12) is true, in spite of the fact that the things in your pocket are not packed. Accommodation has not set the domain of *everything* to be all the items that are salient in the environment, or even all the salient items that belong to you. Rather, it has made the domain all the salient items *of the kind appropriate for packing for a trip.*

What makes something appropriate in this way is part of what happens on trips, or what normally happens on trips. It is part of what cognitive scientists sometimes describe in terms of frames or scripts, or more often these days, plans. The plan or script or frame of trips is needed to fix the domain of *everything.* Hence, in the accommodation process that gives us the natural reading of (12), we will have to look to some such plan or frame or script. I shall not worry about exactly which of these notions is right for describing accommodation; rather, I shall highlight one general point. A plan or script or frame provides highly situation-specific or activity-specific information, often very complex information. What we need to accommodate in (12) is access to specific and detailed information about what kind of activity a trip is, and what happens on one. This tells us what belongs in a suitcase, which is what we use to set the domain of *everything.*[15]

---

[15]A specific proposal on where something script-like fits into domain restriction comes from the analysis of telescoping of Poesio and Zucchi (1992) (they appropriate the term 'script'). Clearly, I am using the term 'plan' very loosely—more loosely than serious work in AI—though I believe this sort of case highlights the 'AI-complete' nature of the problem we face. Some comments relating accommodation to plan recognition more properly construed may be found in Thomason (1990). This paper appears in the collection edited by Cohen *et al.* (1990), which contains

3.2.3 Domains Include Topics

I shall make one more observation about the way contextually restricted quantifier domains are set. This one relies on a general feature of discourse. Discourses have *topics*, roughly, what is under discussion at a given point in a discourse. We can think of discourse topics as given by questions: the topic of a discourse at a given moment is the question that is under discussion. If the question is something like *What did John do?*, we can talk about John being the (current) topic of the discourse. Discourse topics, and how they evolve as discourse progresses, are closely related to the messy issues we encountered when looking at accommodation. Though the facts are sometimes murky here, some helpful generalizations can be made.[16]

Quantifier domains and discourse topics interact in a number of ways.[17] One will be important for our concerns here: Generally, if something is a topic at a given moment in a discourse, we will expect contextually set quantifier domains to include it. Return to the context of (12), and consider:

(13)   a. That's a nice watch you are wearing. Tell me about it.

    b.   ? Everything is packed.

At best, the domain of *everything* can no longer be taken to exclude the watch, as it did in (12). Worse, I am not sure if the pragmatic process of setting a domain even succeeds here, as the second utterance sounds marginal to my ear.

In contrast, we can readily exclude non-topical elements from a quantifier domain by contextual restriction. We saw this with (12). To give one more example, consider:

(14)   a. John decided to ship all his belongings to England.

a number of other papers discussing the role of plans in discourse.

[16]The literature on topic is quite large. Some of it is surveyed in my (2002). Important recent developments of the idea of discourse topic include Büring (2003), Roberts (1996), van Kuppevelt (1995), and von Fintel (1994).

[17]The idea that topic and quantifier domain interact has been investigated at length for the case of adverbial quantifiers, by Partee (1991), Roberts (1995), and von Fintel (1994), among others. A somewhat programmatic suggestion along the lines I am indicating here is also found in Beaver (1994).

b. Everything is small.

Suppose this discourse is taking place as the movers are loading John's belongings into a giant shipping container. The domain of *everything* still does not include the container, roughly, as it is not what we are talking about, even though it is salient in the environment, and figures into the 'plan' for moving.

If we did bring the container into the discourse as a topic, we would get a different result. Consider:

(15) a. John decided to ship all his belongings to England. So, he went out and started investigating shipping containers. He found some that were about the right size.

b. Everything is small.

Now *everything* definitely contains the shipping containers.

I shall rely on the principle that contextually set quantifier domains include topical items. The principle I need, and think is reasonably well-illustrated by the examples we have seen, is that if an item is a topic, it must be in contextually restricted quantifier domains. I do not think it is generally true that such domains include all and only topical elements.[18]

3.3 Setting Quantifier Domains

We now have a few observations about how contextual quantifier domain restriction works. We have assumed that there is a contextual parameter in the nominal of a quantified noun phrase, which accounts for the 'ordinary' cases of domain restriction. Furthermore, we have isolated some principles which govern how this parameter is set:

i. When possible, quantifier domains are built out of predicative material. Both predicates and terms appearing in previous discourse can be used to construct complex defining predicates

---

[18]There are a number of complications to this principle. For instance, the dynamics of topics and subtopics, and interactions with the semantics of nominals, can create apparent violations of the principle. Space limits preclude investigating this, and I shall have to simply assert that attention to details can show these violations to be merely apparent.

for domains.

ii. When appropriate predicative material is not available in discourse, a process of accommodation is triggered.

iii. In constructing domains either anaphorically or by accommodation, domains are constrained to include all topical material.

iv. Accommodation often makes reference to situation-specific information.

These rules hardly tell us everything we might want to know about setting quantifier domains, but they do tell us something.

A couple of morals for application of these principles to the harder case of background domains are worth highlighting. When possible, contextual domain restriction is a pragmatic way to reproduce what could be done semantically by predicates in restricted quantifiers. We see this in principle i. But this is not always the way domain restriction works. When the process of accommodation of principle ii is triggered, it relies on many other factors than simply finding predicates. Principle iii is a clear example of this. No predicate is needed to introduce a topical item into a domain. But the same may be said of the kinds of situation-specific information invoked in principle iv. The complex information encapsulated in a plan or script need not correspond to anything that speakers can express in the context in question with the language they speak; except insofar as they can use a quantifier whose domain is restricted by that information.

4   Reflective Contexts

In that last section, we enumerated some important features of how contextually restricted quantifier domains are set by context. These applied primarily to ordinary contextual domain restriction, which we identified as setting the parameter D^c in the nominal of a quantified noun phrase. We also saw in the last section that if there is to be contextual relativity of contextually unrestricted quantifiers, it must be from a different source. It flows from the context-dependence

of the determiners themselves, which in turn flows from the context-relativity of the background domain M.

My goal now is to apply the lessons of section 3 to the case of background domains, to help us to understand how, in some extreme cases like we see in the paradoxes, background domains vary with context. Indeed, we can now see a little more clearly why focus on the paradox is so important. As we are reminded by principle iv, any account of the effects of context on quantifier domains will be highly situation-specific. We need to see with what situation the paradoxical reasoning of section 1 confronts us.

4.1 Artifacts of Discourse

The argument from paradox of section 1 leads us to identify an object not in the domain of an unrestricted quantifier. To understand how this marks a shift in context, we should begin by asking what the objects we are led to recognize are like.

An important point that is made vivid by the Williamson version of Russell's paradox of section 1 is that the objects in question are semantic in nature. In that version, the object we identify is the *interpretation* $o = I(R)$; in more typical versions, it is the background *domain of quantification* itself. The objects in question are semantic values, or more complicated objects built from semantic values, like an interpretation of an entire language relative to some context. Other forms of the paradoxes can lead us to other related objects, including truth predicates, propositions, contexts, etc. To give this category a label, let us call them *artifacts of discourse*. Semantic values of expressions relative to contexts, including quantifier domains, will be the main artifacts of discourse for our discussion here, but the category is somewhat wider. The characterization of artifacts of discourse is admittedly rough, but at least the main examples are familiar. Rather than refine the definition, I shall move on to examine how artifacts of discourse interact with quantifier domains.

The answer is that usually, the do not. Artifacts of current discourse—the quantifier domains of the context of the current point in the discourse, the interpretation of the language in the context, etc.—are *usually not* part of any contextually restricted quantifier domain given by the context of the current point in the discourse. Here is one example to illustrate the point. At the start of a set

theory class, the professor says:

(16) Everything today is finite.

She means, roughly, that everything relevant to the day's class is a finite set. But within this category, the semantic values of her own words are excluded. They are excluded even though they do, in an entirely natural sense, count as relevant to the class. (Apprehending them is crucial to understanding the class's content.) They are also excluded even on the assumption that semantic values are sets. What the professor says remains *true* if the semantic values of her words turn out to be infinite sets.

Though most speakers do not tend to care about semantics (including most set theory professors), that is not the issue here. Artifacts of current discourse tend to be excluded from quantifier domains even if we are generally talking about semantic values, or other artifacts of some discourse or another. Suppose a semantics professor says:

(17) Every semantic value relevant to today's class/in sight/at issue is an individual or a set of individuals.

The domain of the quantifier again does not include the values of her own words, and the claim is not made false by the value of *every* being a relation between sets. To take it as if it did include these objects would not be to stretch the limits of what counts as relevant; rather, it would be to perversely disregard the normal rules of discourse.

The moral is that it is extremely hard to incorporate artifacts of current discourse into contextually restricted quantifier domains. Not impossible (at least, not in all cases), but strongly discouraged by the normal rules governing discourse. It follows from this, together with principle iii of section 3.3, that artifacts of current discourse are extremely difficult to make *topics*. This is not simply to say we often do not make them topics, but that it is unusual, and hard, to succeed. At the very least, topicalizing an artifact of current discourse—making it a topic—amounts to a violation of the rules of well-organized and coherent discourse.[19]

---

[19]I am here appealing to what appears to be a fact about natural language, for which (16) and (17) provide a little bit of evidence. I have discussed the relation of topic to discourse coherence extensively in my (2002). Note

Artifacts of current discourse thus have a peculiar status. They are usually non-topical, and cannot even be coherently topicalized in ordinary discourse. But nonetheless, they are clearly active in discourse in another way. They are the semantic values of the very words we are speaking, the context in which we speak, etc. We cannot understand a discourse unless we apprehend them. To fix some more terminology, let us say that artifacts of current discourse are *implicit* in the discourse. They are not explicit, in that they are not normally available to be topics, but they are clearly an important part of the discourse.

4.2 Reflection

I have pointed out that it is very difficult to topicalize implicit objects like artifacts of current discourse. The rules of discourse ordinarily tell us not to do it. But in fact, we can do it, if we really want to. At least, for a given point in a discourse, we can step back and start talking about the semantic properties of the discourse as it stood at that point. This introduces what were the current artifacts of discourse as topics. This will be jarring, and outside the normal rules of discourse, but we have the ability to do it.[20]

Doing this will change the context. Generally, what is topical at a given point in the discourse is part of the context. This again is a fact about the way natural language works. Again, I hope it is intuitively clear. It deserves more argument, but I shall leave that to the literature on context.[21] Forcing something implicit like an artifact of current discourse to be a topic is, if anything, a highly marked change of context. It is not a natural evolution of the context as a discourse progresses, but a discontinuous jump in context caused by a change in topic violating the normal rules of discourse.

that in many discussions of the syntax of topic-marking, the term 'topicalization' is used for a particular syntactic construction. I mean making something a discourse topic. It is shorter to say 'topicalize' than 'incorporate as a discourse topic'.

[20]I am taking it for granted that we have the intellectual resources to topicalize artifacts of discourse, even if it is only in unusual contexts where we do so. I suspect this may be the crux of my disagreement with Williamson (2004). Though linguistically speaking we can, it appears to have a nominalizing effect, and so is something I believe he would have to reject.

[21]This is a common theme in most of the pragmatically oriented work on discourse topic, such as those cited in footnote 16.

With this in mind, let us turn to the situation which confronts us in the argument from paradox. The key step in the paradoxical reasoning we rehearsed in section 1 is precisely a step of topicalizing an artifact of current discourse. In the usual forms of the paradox, we topicalize the background domain of the current context (the domain over which the quantifier *all objects* ranges). In the Williamson version, we topicalize an interpretation of the language, as it is used in some context. We do this when we take the domain or the interpretation and start talking about it, particularly, asking what properties it has. The previous discussion shows that this step in the argument from paradox will amount to a change of context, and indeed, marked and unusual one.

Paradox also enters the fray here, and tells us something specific about what happens in this change in context. It introduces an object as topical which could not be within the background domain of the initial context of the argument from paradox, upon pain of contradiction. Now, our examination of ordinary cases already gave us some reason to expect this. Artifacts of current discourse are implicit, and so will not normally appear in any restricted quantifier domain, as we saw in section 4.1. Why not? The most natural reason would be that they are not even in the background domain. Considering ordinary discourse can at best lead us to find this natural, but the paradox shows us that it has to be right!

Not every artifact of current discourse leads to paradox. Let us call the ones that do *crucial artifacts of discourse.* The crucial artifacts of discourse include the background domain given by the context of the current point in a discourse, or even the entire interpretation of the language relative to that context. They also include any objects which contain or encode these, or would allow us to extract them by some process of accommodation. With the sorts of set-theoretic assumptions I alluded to in section 1, we can assume that any crucial artifact of discourse will make the background domain of the context available for topicalization. It is a common idea that what makes an artifact of discourse crucial is that it is extremely large or comprehensive relative to a given context, as are the background domain of the context, or the interpretation of the entire language as it is used in that context. A more full account of what makes certain artifacts of discourse crucial would be very useful, but I shall for the moment rest with the usual list of objects which lead to paradox.

To summarize what we have seen, fix a context c_0 for a discourse. Say a *reflective context* for c_0 is one in which we have topicalized any crucial artifact of the discourse at c_0. I am assuming that this will at least implicitly introduce the background domain M_0 of context c_0 as a topic. Call the reflective context for c_0 by c_0^R. A reflective context has the feature of taking something that was only *implicit* in c_0 and making it topical—making it explicit—in c_0^R.[22]

We have noted that we have solid reason, coming from observations about ordinary discourse, to expect that c_0^R is really a new context, and at least some reason to expect that it should have a strictly wider background domain than c_0. Furthermore, we have seen that the argument from paradox shows that c_0^R *must* have a strictly wider background domain, and so must certainly be a distinct context. In particular, the background domain M_0^R of c_0^R must contain M_0. Indeed, as M_0 is a topic in c_0^R, we may expect even contextually restricted domains in c_0^R to contain M_0. The paradox hardens our expectation that c_0^R is a new context with a strictly wider background domain into a requirement of logic.

We now have one major piece of the contextualist response to the paradoxes in place. There is, I have argued, a context shift in the course of the argument from paradox. It is of a particular sort: a step to a reflective context, which makes topical something which was merely implicit in the initial context.

5   Domains for Reflective Contexts

If we take the step to a reflective context—the step from c_0 to c_0^R—we face a problem. We have topicalized an item which cannot be in the background domain of our initial context c_0, so a new, strictly wider background domain M_0^R for c_0^R is needed.

So far, I have argued that we should have expected this, and that the paradox shows it must be the case. But that just presses the issue of what M_0^R looks like. If there is a change in context from c_0 to c_0^R, it is not enough to simply say that M_0^R must be strictly bigger than M_0. We want

[22]The idea of reflection making what was implicit to be explicit appears in other forms as well. In a proof-theoretic setting, a similar idea is discussed in Kreisel (1970). I tried to apply this to the Liar in my (2004c). A related point is made in Parsons (1974a).

to know how the change in context expands the domain.

It is here that our investigation of quantifier domains in section 3.2 will pay off. Building a new background domain is not the same as setting the contextual parameter D^c for restricted domains, but it is still setting a domain of quantification. If so, it should be governed by the principles of section 3.3. In this section, I shall show how these principles, supplemented by some of the observations about why background domains are different from section 3.1, help us to see how M_0^R might be constructed. They will tell us how context can set a new background domain as we move to a reflective context.

5.1 Triggering Accommodation

My first observation is that principles i–iii from section 3.3 trigger an accommodation process. Of course, the paradox already forces c_0^R to have a strictly wider background domain. But once we take the step to the reflective context c_0^R, this is not merely a recondite fact of logic. Once we topicalize the domain M_0, principle iii requires this object to appear in *restricted* quantifier domains relative to c_0^R. We will have to talk about it and quantify over it. Thus, we will really have to make use of the expanded domain M_0^R. So, the pragmatics of domain-setting will require us to work out what the new domain is.

If the same domain-setting processes at work in ordinary cases are at work here, then we would expect the construction of c_0^R to first try to recover a new domain M_0^R from previous discourse. But of course, this is impossible. By principle i, it would have to do so by finding appropriate predicative material. But predicative material from previous discourse will not be able to describe any domain containing M_0 as an object. At best, it will describe the objects in M_0. Hence, by principle ii, we must accommodate.

5.2 Towards a 'Plan' for Accommodation

How should this accommodation proceed? As we are reminded by principle iv, we will need to make use of highly specific information about the particular situation involved: the situation of shifting to a reflective context. This information should provide us with guidelines for building the new

domain—something like a plan or script or frame from which we can extract an accommodation process. I shall thus loosely talk about a 'plan' for accommodation. (I shall usually put 'plan' in scare quotes, to remind us how loose the talk of plans is.)

They key feature of a reflective context is the topicalizing of a crucial artifact of the discourse as it stood at c_0. To build any kind of coherent plan in this situation, we should have at our disposal vocabulary for describing such artifacts and their basic properties. So as a first step, we should add this if it is not already present:

Step 1: Add vocabulary that describes in c_0^R the semantics of the language as it was used in c_0. (Even if the relevant vocabulary was present, we may need to adjust extensions for the new context.)

As with any case of accommodation, speakers may have trouble finding words to make explicit exactly what information they have accommodated. So, when saying we should add 'vocabulary', we should say more fully to add information that could appear in a fully articulated discourse even if speakers only tacitly grasp this information.

Step 1 is really only a set-up move. It puts us in a situation to make some sort of coherent plan appropriate for a reflective context. So far, we have from the prior context the elements of M_0, from the transition to c_0^R we have M_0 as an object itself, and we have vocabulary for describing the semantic properties of the language as it was used in c_0 (for shorthand, let us say the *semantics of* c_0).

To add the next step to the 'plan', we need to remember that we are not simply trying to build any old set containing M_0. We need to construct a viable background domain. One of the crucial roles of background domains is to be the source of restricted quantifier domains. As principle i of section 3.3 tells us, the basic way restricted quantifier domains subsequent to c_0^R will be built is by forming possibly complex predicates out of material available in the context. These predicates then define restricted domains as subsets of the background domain. A viable background domain must be rich enough to allow for the formation of restricted domains by this process. In particular, once we have taken step 1, we need to build a domain which will include any restricted quantifier

domain we can define using the extended vocabulary provided by step 1.

Assuming that the extended vocabulary provided by step 1 is rich enough, this will require including subsets of M_0 as elements of our new domain M_0^R. For instance, if we have predicates like *x is a semantic value of a verb phrase in* c_0, we will have to make sure M_0^R contains each subset of M_0 definable in extended vocabulary. Thus, we should at least take as our next step:

Step 2: Close under definable subsets of M_0 in the extended vocabulary.

This begins to get us a suitable background domain for providing contextually restricted domains of quantification in c_0^R and subsequent contexts.

Actually, we need to a little more than step 2. As stated, step 2 only provides subsets of M_0. But if we are thinking of having vocabulary for the semantics of something like a natural language, we will have higher-type objects as well. Determiners will have relations between sets as their values, for instance. At the very least, we will have to iterate step 2 several times. We should re-formulate step 2 as:

Step 2′ Close under definable subsets of M_0 in the extended vocabulary. Iterate as many times as needed by the semantics of c_0.

As we iterate, the extended vocabulary will get used more and more, to describe the behavior of more and more complicated artifacts of discourse from c_0.

5.3 Westerståhl's Principles Revisited

So far, we have begun to describe a process of accommodating a new background domain M_0^R for the reflective context c_0^R. Step 2 tries to take into account the fact that we are not just constructing any old quantifier domain; rather, we are constructing a *background domain*. Once we have the extended vocabulary we need to accommodate the new domain, it ensures that we can also provide for additional contextually restricted domains that we could define using that vocabulary.

But step 2 does not go far enough. It does not, I suggest, precisely because it does not yet pay attention to the two principles we discussed while considering Westerståhl's argument in section 3.1.

Westerståhl's principles show that background domains behave very differently than contextually restricted quantifier domains. To build an acceptable background domain, we need to write into our 'plan' rules that ensure we satisfy these principles.

We satisfy both if we do as much as we can to satisfy WP1, which reminds us that background domains are large. The result of step 2 appears unduly small. It includes a few new elements from M_0, but then just stops. Using step 2′ instead will iterated this process a few times, but then it still just stops. We can certainly iterate much further than is required by step 2′. To build a large domain, as WP1 requires, we should do just that. We should iterate step 2 as far as we can.

The more we iterate step 2, the more we satisfy WP2 as well. This principle tells us that background domains should be stable. Westerståhl had in mind that they never change. I have rejected that, but I have argued that shifting to a reflective context is a fairly unusual step. We should still build a domain which minimizes the need for such transitions. The larger the domain, the less need we will ever have to expand it.

For instance, what happens if we succeed in topicalizing *the domain of semantic values from* c_0? This object will not be in the domain resulting from step 2, or step 2′. Would topicalizing it amount to a step to a new reflective context? I am inclined to think not. Topicalizing it is introducing a metasemantic object—a meta-artifact of discourse—as a topic. But the object is one we get simply by collecting together the semantic values from the original context c_0. It results from the metasemantics of the semantics of c_0. Topicalizing it does not seem to be a new instance of reflection, but rather continuing the reflection we had already started. A new reflective context would be one which reflects on the semantics of c_0^R itself. To avoid counting this kind of continued reflection as leading to a new reflective context, we need to make sure objects like *the domain of semantic values from* c_0 are already in the background domain M_0^R provided by c_0^R. To do so, we need to iterate step 2 past what step 2′ itself requires.

The more we iterate step 2, the more we satisfy both Westerståhl's principles. The more we iterate, the bigger the background domain of c_0^R is. Likewise, the more we iterate, the less we see occasions for shifting to a new reflective context with a bigger background domain. The more we

iterate, the more stable background domains are. Thus, it appears that to satisfy Westerståhl's principles for background domains, we need one more step in our 'plan':

Step 3: Iterate step 2 as far as possible.

What is as far as possible? Presumably until some appropriate closure condition is reached, or until the resources for iteration are exhausted. I shall discuss this further in section 6.

5.4 Tarski and Kripke

It is worth pausing to note that the issue of how far to iterate, and where we see new reflective contexts, essentially reprises the contrast between Tarskian and Kripkean approaches to the Liar.

If we were to stop at step $2'$, we would have an essentially Tarskian view (Tarski 1935). Though we have replaced talk of languages with talk of what is expressed in contexts, we would essentially have in c_0^R resources for describing the semantics of a language as used in c_0, but nothing more. Any further metasemantic reflection would require ascending to a new reflective context, which we can count as for all intents and purposes ascending one level in a hierarchy.

On the other hand, just as in the Liar case, iteration helps to minimize the number of distinct levels. Just as with Kripkean iteration of the Tarskian truth predicate (Kripke 1975), iterating step 2 allows for some modest amounts of metasemantic discourse within c_0^R itself.

I have suggested that Westerståhl's principles give us reason to pursue this more Kripkean strategy. Even so, it is still possible to reflect on the semantics of c_0^R itself, and that will induce a new reflective context. If $c_0^R = c_1$, then we can always move to c_1^R by topicalizing a crucial artifact of the discourse at c_1.[23] Though the proposal I am making is more Kripkean than Tarskian, it is still in effect a hierarchical proposal.[24]

---

[23]Kripke notes something like this in talking about "the ghost of the Tarski hierarchy" (1975:80).

[24]It may be that from c_1^R, the next reflective context up, we might be able to find a predicate which defines M_0^R as a subdomain of M_1^R. Indeed, the 'plan', and the considerations of section 6 to follow, make this seem likely. But the process of building M_0^R as part of the step to a reflective context is not that process at all. Rather, it is the process of from the 'bottom up', or as I shall suggest in a moment, inductively, generating a new domain which can serve as a background domain. Hence, the basic outlook is still, in the terms of Fine (this volume), expansionist.

6 Iteration

So far, we have taken a number of steps towards articulating the contextualist response to the paradox. Semantically, we have identified two different sources of context-dependence in section 3.1. Contextual domain restriction sets a parameter D^c, which combines with the semantic values of nominals. The context-relativity of unrestricted quantifiers flows from the context-relativity of the background domain M, which affects the semantic values of determiners. In effect, the context-relativity of M gives determiners an indexical character.

Pragmatically, we saw in section 4 how the step to a reflective context counts as a genuine change in context, and how it induces a change in the background domain. Unlike cases of contextual domain restriction, these context shifts are highly unusual, and violate some general guidelines for keeping discourse orderly and coherent.

When it comes to setting domains of quantification, we isolated some general principles in section 3.2–3.3, by studying ordinary contextual restriction. We applied these, together with Westerståhl's principles from section 3.1, to the case of background domains in section 5. The plan we developed in section 5 differed from the process of setting restricted quantifier domains in some important ways. But where it does, it is because of the specific features of reflective contexts and background domains. The rules for setting quantifier domains already require us to take into account such specific features of contexts. Thus, it is fair to say that the pragmatic processes that set background domains and contextually restricted domains are in fundamental respects the same.

The 'plan' I sketched in section 5 instructed us in step 3 to iterate step 2. How far such iteration will go remains an open question. In this section, I shall examine this question, by bringing to bear some tools from mathematical logic. This will also enable me to offer a somewhat idealized mathematical model of what the background domains described in section 5 might look like. My discussion here will, by necessity, be somewhat more technical than what has come before. For those readers wishing to skip the technicalities, the principal claim of this section can be summarized as follows: logic provides some plausible stopping points for the iteration required by step 3.

6.1 Reflective Contexts and Constructible Sets

The 'plan' I sketched in section 5 will have a very familiar ring to logicians. It is essentially the instructions for building levels of the constructible hierarchy with urelements. In this section, I shall spell out the basic idea of this connection. I shall give a more technically precise account in section 6.2.

The 'plan' of section 5 told us to take our initial background domain M_0 and start a process of adding elements to it. The elements we are to add are those subsets that are definable in an appropriately extended vocabulary. Adding definable subsets as members is just the process of building the constructible sets. Starting with M_0 in effect would be to build the constructible sets with urelements from M_0. Thus, the plan tells us to build the constructible sets with urelements up to some appropriate level in the constructible hierarchy.

Why just the constructible sets? Why not build up all the sets, up to some appropriate level in the rank hierarchy? The quick answer is because the plan does not tell us to. The plan has us expand our initial domain M_0 for specific reasons: to make an acceptable background domain for our reflective context c_0^R. None of those reasons indicated going beyond the constructible sets, as they do not ask us to include anything beyond the definable. This plan came from considering what happens in the step to a reflective context, and asking how the general guidelines for fixing quantifier domains should be applied to it. Thus, what the plan tells us to do really does seem to be all that we should do. As the plan tells us only to add constructible sets, that is all that should be added.

There is also good reason to keep the process of building a new context as constrained as we can. Establishing a new context is something that speakers do. In passing to a reflective context c_0^R, speakers will at least implicitly have to carry out the task of building the new background domain M_0^R. The general principle 'do no more than required' seems to be a good one to invoke for what is already a massive task. Thus, though speakers may be able to understand what it would be to build a larger domain than the plan calls for, it is not required, and so they will not build it.

What of the rest of set theory? The most natural proposal, I think, is that an initial background

31

domain M_0 had better include the usual objects of mathematics and science. Shifting to talk about mathematics or science is not shifting to a reflective context, and we have no particular reason to think it requires a shift of background domain of any kind. Hence, if we are concerned about sets, we should think of them as included in M_0.[25]

If this is right, the question of how far to iterate in step 3 of the plan comes down to what level of the constructible sets with urelements to stop at. The considerations we just raised point to a strategy for answering this question as well. The process of building a new domain for c_0^R must be something speakers can at least implicitly make sense of. Thus, the iteration should proceed as far as speakers can likewise make sense of.[26]

How far is that? The plan describes a kind of inductive process for generating a new domain by adding elements in stages. Thus, it is natural to suggest that the right stopping place is the limit of the lengths of inductive processes that are available relative to the expressive resources of c_0. I shall spell this out a little more formally in the next section.

6.2 A Little More Formally

Let us now try to spell out the basic idea of section 6.1. To begin, suppose we describe c_0 by fixing a language $\mathcal{L}$ and a structure $\mathfrak{M}_0$ for $\mathcal{L}$. $\mathfrak{M}_0$ provides the semantics of c_0, and so interprets $\mathcal{L}$ as it is being used in context c_0. To facilitate doing a little logic, let us suppose $\mathcal{L}$ is a first-order language. This is certainly an idealization, but it will prove useful. (I shall use M_0 for the universe of $\mathfrak{M}_0$, which preserves its role as the background domain of c_0.)

---

[25]Kit Fine suggested to me that the sorts of domain completion found often in mathematics, like passing from the real numbers to the complex numbers, might count as cases of background domain expansion. The proposal I am floating here would be to treat them as a cases of passing from one restricted domain to another, all within M_0. This would leave the cases of genuine background domain expansion limited to passing to reflective contexts. Both options are consistent with the general outlook of contextualism. If we accept Fine's option, we no longer need to assume that all, or even most, of the objects of mathematics are in a given initial background domain.

[26]If we really do assume that the initial background domain M_0 contains all the sets, we will encounter some technical complications I shall ignore for current purposes. Rather than constructible sets, for instance, we should perhaps be talking about iterated predicative classes. For thinking about the matter more technically, I shall assume that M_0 is a set.

As c_0 is supposed to be an initial context, where we have not engaged in any reflection, we could either suppose that $\mathcal{L}$ contains no semantic vocabulary, or that all the semantic vocabulary has empty interpretations in $\mathfrak{M}_0$. It will simplify matters slightly, and allow us to invoke some standard notation, to take the former route. When we move from c_0 to c_0^R, step 1 of the 'plan' of section 5 instructs us to add all the vocabulary we might use to describe the semantic properties of the language as used in c_0. We thus extend $\mathcal{L}$ to a language $\mathcal{L}^*$ having the needed semantic vocabulary.

I shall continue in my practice of treating semantic values as sets, and generally, of identifying semantic talk with set-theoretic talk. Relative to this simplifying assumption, we should suppose $\mathcal{L}^*$ supplements $\mathcal{L}$ by adding set-theoretic vocabulary: a membership relation $\in$, and quantifiers and variables over sets. $\mathcal{L}^*$ is a two-sorted language, with one sort of variables ranging over elements of M_0, the other ranging over sets. We will enforce this strictly, by having set variables not range over urelements.

We are thinking of $\mathcal{L}^*$ as being used in context c_0^R. Relative to this context, it must be interpreted by an appropriate structure. A structure for $\mathcal{L}^*$ is a structure $\langle \mathfrak{M}, A, E \rangle$, where A is a universe of sets with urelements drawn from M, and E interprets $\in$. Conventionally, we insist that A contains only sets and no urelements, to keep the two sorts of quantification in $\mathcal{L}^*$ separate.

As step 1 tells us to move from $\mathcal{L}$ to $\mathcal{L}^*$, step 2 tell us to add to our domain the subsets of M_0 definable in $\mathcal{L}^*$. Step 3 tells us to iterate this process. Iterating up to an ordinal α is essentially building $L(M_0, \alpha)$: the constructible sets with urelements from M_0 up to level α. We can then use $L(M_0, \alpha)$ to form the $\mathcal{L}^*$-structure $L(\alpha)_{\mathfrak{M}_0} = \langle \mathfrak{M}_0, L(M_0, \alpha) \cap V_{M_0}, \in \rangle$.[27]

Carrying out the 'plan' for building a background domain sketched in section 5 can thus be modeled by the process of constructing $L(\alpha)_{\mathfrak{M}_0}$ for an appropriate α. As step 3 of the plan leaves

---

[27]V_{M_0} is the universe of *sets* with urelements from M_0. The intersection with V_{M_0} is needed to make sure the second universe is strictly a universe of sets, and not of both sets and urelements. Clearly, the details of working with $\mathcal{L}^*$ and urelements here can be somewhat fussy. For a full exposition, see Barwise (1975), who's notation and conventions I am following. The basic idea is still just that of building the constructible sets as is done in any set theory textbook, or more closely, building $L_\alpha[M_0]$ while taking definability in $\mathcal{L}$ into account.

open how far to iterate step 2, we have so far left open which ordinal α will be appropriate. But setting that aside, $L(\alpha)_{\mathfrak{M}_0}$ provides an interpretation of $\mathcal{L}^*$, the enriched language being used in the reflective context c_0^R. This in turn gives us an expanded background domain for c_0^R, which we get by combining M_0 with the domain A of sets we have added. In fact, this simply gives us $L(M_0, \alpha)$.

Step 3 tells us to iterate 'as far as possible'. How far is that? A moment ago I suggested that we think about 'as far as possible' in terms of the limits of processes speakers could at least implicitly make sense, given the expressive resources they start with in c_0. The slightly picky definitions we have just given will allow us to flesh this out formally. As we are assuming speakers can pass to a reflective context, they are able to reason about the semantics of $\mathcal{L}$ as used in c_0, which is given by $\mathfrak{M}_0$. Moreover, in following the 'plan' of section 5, they must be able to make sense of a kind of 'bootstrapping process' which uses $\mathfrak{M}_0$ to build up new domain in stages. Formally, this is an inductive definitions on $\mathfrak{M}_0$. Insofar as they understand these inductive processes, they should be able to make sense of iterating 'as far as such processes go'. This suggests the limit of the lengths of inductive definitions on $\mathfrak{M}_0$ as a way to capture the amount of iteration needed in the step from c_0 to c_0^R.

The formal model of iterating to exactly this limit is the well-studied structure known as $HYP_{\mathfrak{M}_0}$. In the case where $\mathfrak{M}_0 = \mathbb{N}$, the natural numbers, $HYP_{\mathfrak{M}_0}$ is simply $L(\omega_1^{CK})_{\mathbb{N}}$, where ω_1^{CK} is the first non-recursive ordinal. For $\mathbb{N}$, we build $HYP_{\mathbb{N}}$ by iterating through all the recursive ordinals. More generally, for structures $\mathfrak{M}_0$ which share enough properties with the natural numbers, $HYP_{\mathfrak{M}_0} = L(\alpha)_{\mathfrak{M}_0}$ for α the limit of the closure ordinals of (first-order positive) inductive operators on $\mathfrak{M}_0$.[28]

One of the very nice features of the HYP construction is that it is iterable. The problem of building background domains is not restricted to c_0^R. As we observed in section 5.4, the process of

---

[28]These results are surveyed in Barwise (1975), among places. As I mentioned above, we might have to think of our starting structure $\mathfrak{M}_0$ as looking different from the natural numbers. We might encounter contexts in which our background domain already looks like a model of set theory, for instance. On the other hand, thinking about the issue from the perspective of linguistics or cognitive scientist, we might see things very differently. From those perspectives, a natural working hypothesis might be that the background domain we start with is *finite*.

stepping to new reflective contexts continues. From $c_0^R = c_1$, we can pass to c_1^R, which will need a new background domain. HYP allows for a uniform account of these steps. From $HYP_{\mathfrak{M}_0}$ we can form $HYP(HYP_{\mathfrak{M}_0})$, which can provide the background domain for c_1^R.[29]

So, one appealing option for how far to iterate is given by $HYP_{\mathfrak{M}_0}$. There are other options as well. For instance, we might take the idea of iterating 'as far as possible' to be that of iterating until you get nothing new. Insofar as this might mean nothing *recognizably* new, this might mean iterating until $L(\alpha)_{\mathfrak{M}_0}$ is an elementary substructure of $L(\alpha + 1)_{\mathfrak{M}_0}$. An ever stronger option, taking into account the role of inductive definitions I just alluded to, would be to iterate until $L(\alpha)_{\mathfrak{M}_0} \prec_1 HYP(L(\alpha)_{\mathfrak{M}_0})$. Ordinals satisfying these properties can be found.[30]

We thus have a number of viable options for how far iteration must go to build an appropriate background domain for c_0^R. As a working hypothesis, I am inclined to opt for $HYP_{\mathfrak{M}_0}$. It provides a workable picture of the semantics of $\mathcal{L}^*$ in c_0^R, and a plausible picture of what background domain this context might give us. Furthermore, it is fairly well-understood mathematically. All the same, the little exercise in mathematical modeling we have just been through does not conclusively tell us which option to choose (and as an exercise in mathematical modeling, it builds in some incidental features of the mathematics used, as well as the features we are trying to capture with the model). In considering multiple options, I do not want to suggest that there is nothing to distinguish among them, which would induce rampant indeterminacy in the notion of background domain. Rather, I think the moral to be drawn is that we do not yet know enough to be certain just how far iteration really does go.

Let me close this section by briefly mentioning one important feature of the formal model. It respects the distinction between artifacts of the discourse and other objects in a thoroughgoing way. Non-artifacts of the discourse correspond to urelements. These do not change as we move

---

[29]There is also a nice match-up between $HYP_{\mathfrak{M}}$ and Kripke constructions on $\mathfrak{M}$, as has been thoroughly investigated by McGee (1991), and by my (2004a).

[30]The former is what is known as an $\alpha + 1$-stable ordinal, the latter an α^+-stable ordinal (where α^+ is the next admissible ordinal greater than α). These were extensively investigated by Richter and Aczel (1974), who show that they are analogs in generalized recursion theory of indescribable cardinals. They also show that the least α^+-stable ordinal is the limit of the closure ordinals of non-monotone Π_1^1-inductive operators.

to reflective contexts. Indeed, the predicates *x is an urelement* and *x is a set* are Δ, and so are *absolute* in the logician's sense of not changing as models expand. Furthermore, the same holds for any 'ordinary' non-semantic expression. In the model, these will be expressions of $\mathcal{L}$, and their interpretations all remain constant. For instance, for a predicate P of $\mathcal{L}$, the set $\{m \mid P(m)\}$ does not change at all as we move from c_0 to c_0^R, etc. P in $\mathfrak{M}$ is the same as P in $HYP_\mathfrak{M}$ is the same as P in $HYP(HYP_\mathfrak{M})$. The absoluteness of the set/urelement distinction ensures this, mathematically trivially.

This holds for all formulas of $\mathcal{L}$. Quantification in $\mathcal{L}$ is not a problem, as quantification over urelements is bounded quantification in $HYP_{\mathfrak{M}_0}$. Nor is negation a problem, in the formal treatment using a two-sorted language I have offered here. However, the two-sorted language may appear somewhat artificial. If instead, for instance, we think of an expression like P as having its extension in $M_0 \cup A$, the combined domain urelements and sets, we have to a little more careful. $\neg P$ would then change its extension as A expands. But even so, the absoluteness of the set/urelement distinction will allow us to form an absolute predicate $U(x) \wedge \neg P(x)$, where $U(x)$ is *x is an urelement*. The extension of this predicate will not change as the domain of sets expands. (As usual, this observation really applies to Σ-formulas of $\mathcal{L}^*$.)

The moral of this observation is that changing to a reflective context does change the background domain of quantification, and the interpretation of the language of artifacts of discourse, but it leaves the interpretations of other aspects of the language as we speak it in c_0 unchanged.[31]

7 Conclusion

I have now presented the core features of the contextualist approach to the paradoxes, and so made my basic case. I have argued that quantifiers which are both unrestricted and contextually unrestricted show relativity to a contextually determined background domain. They do so because

---

[31]In a longer version of this paper, I used this moral to counter an objection from Williamson (2004) that non-absolutist views of quantification, including the contextualist view I have been pursuing here, cannot even capture *restricted* quantification properly. However, a full discussion of this interesting challenge will have to wait for another occasion.

of the way the background domain M figures into the semantics of determiners. I noted that this is a distinct phenomenon from the usual quantifier context-dependence, as it is relativity to the background domain M, over and above the context-dependence of the parameter D^c in the nominal of a noun phrase. But I also argued that the pragmatic mechanisms that set M are fundamentally related to those that set contextually restricted quantifier domains. First of all, I argued that there really is a context shift in the argument from paradox, as it involves a step to a reflective context. Furthermore, I showed how the same general principles which govern the setting of contextually restricted domains govern the setting of a new expanded background domain for a reflective context. Where setting M for a reflective context differs from setting a contextually restricted domain, it is because setting M relies on the specific nature of reflective contexts. Yet it is a general feature of domain-setting to rely on such specific features. Finally, to flesh out my proposal, I offered some formal models of what an expanded background domain might look like.

To close, let me mention some issues that remain open (over and above the larger issues I put aside at the beginning of this essay). Two seem to me most important. First, this essay has concentrated on how domains expand, particularly, how they expand in the step to a reflective context, which I see as the crucial step in the argument from paradox. Very little has been said here about how an *initial* background domain is set, and understanding this is important for a full account of the context-relativity of background domains. Second, a great deal of weight has been placed on the distinction between artifacts of discourse and other objects. I gave a rough-and-ready version of this distinction, and more needs to be done to understand it in its full generality.

I hope that developing the positive aspects of the contextualist approach systematically, as I have tried to do here, will serve to dispel some the air of mystery which seems to attach to it. Answering these questions would do so all the more.

References

Bach, K. (1994). 'Conversational Impliciture', *Mind and Language*, 9: 124–162.

Barwise, J. (1975). *Admissible Sets and Structures* (Berlin: Springer-Verlag).

Barwise, J. and Etchemendy, J. (1987). *The Liar* (Oxford: Oxford University Press).

Beaver, D. (1994). 'Accommodating Topics', in P. Bosch (ed.), *Proceedings of the IBM/Journal of Semantics Conference on Focus*, vol. 3 (Heidelberg: IBM), 439–448.

Boolos, G. (1984). 'To Be Is to Be the Value of a Variable (or to Be Some Values of Some Variables)', *Journal of Philosophy*, 81: 430–449.

Büring, D. (2003). 'On D-Trees, Beans, and B-Accents', *Linguistics and Philosophy*, 26: 511–545.

Cappelen, H. and Lepore, E. (2002). 'Insensitive Quantifiers', in J. Keim Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 197–213.

Carston, R. (2004). 'Explicature and Semantics', in S. Davis and B. Gillon (eds.), *Semantics: A Reader* (Oxford: Oxford University Press), 817–845.

Cartwright, R. L. (1994). 'Speaking of Everything', *Nous*, 28: 1–20.

Cohen, P. R., Morgan, J., and Pollack, M. E. (eds.) (1990). *Intentions in Communication* (Cambridge: MIT Press).

Davis, S. (1991). *Pragmatics* (Oxford: Oxford University Press).

Dummett, M. (1973). *Frege: Philosophy of Language* (Cambridge: Harvard University Press).

——— (1993). 'What Is Mathematics About?', in *The Seas of Language* (Oxford: Oxford University Press), 429–445.

Fine, K. (this volume). 'Relatively Unrestricted Quantification', in A. Rayo and G. Uzquiano (eds.), *Absolute Generality* (Oxford: Oxford University Press).

Gawron, J. M. (1996). 'Quantificational Domains', in S. Lappin (ed.), *Handbook of Contemporary Semantic Theory* (Oxford: Blackwell), 247–267.

Geurts, B. and van der Sandt, R. (1999). 'Domain Restriction', in P. Bosch and R. van der Sandt (eds.), *Focus: Linguistic, Cognitive, and Computational Perspectives* (Cambridge: Cambridge University Press), 268–292.

Glanzberg, M. (2001). 'The Liar in Context', *Philosophical Studies*, 103: 217–251.

——— (2002). 'Context and Discourse', *Mind and Language*, 17: 333–375.

——— (2004*a*). 'A Contextual-Hierarchical Approach to Truth and the Liar Paradox', *Journal of Philosophical Logic*, 33: 27–88.

——— (2004*b*). 'Quantification and Realism', *Philosophy and Phenomenological Research*, 69: 541–572.

——— (2004*c*). 'Truth, Reflection, and Hierarchies', *Synthese*, 142: 289–315.

Grice, P. (1975). 'Logic and Conversation', in P. Cole and J. L. Morgan (eds.), *Speech Acts*, vol. 3 of *Syntax and Semantics* (New York: Academic Press), 41–58. Reprinted in Grice (1989).

——— (1989). *Studies in the Way of Words* (Cambridge: Harvard University Press).

Heim, I. (1983). 'On the Projection Problem for Presupposition', *Proceedings of the West Coast Conference on Formal Linguistics*, 2: 114–125. Reprinted in Davis (1991).

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic* (Dordrecht: Kluwer).

Kaplan, D. (1989). 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes From Kaplan* (Oxford: Oxford University Press), 481–563. First publication of a widely circulated manuscript dated 1977.

Keenan, E. L. and Stavi, J. (1986). 'A Semantic Characterization of Natural Language Determiners', *Linguistics and Philosophy*, 9: 253–326. Versions of this paper were circulated in the early 1980s.

Keenan, E. L. and Westerståhl, D. (1997). 'Generalized Quantifiers in Linguistics and Logic', in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language* (Cambridge: MIT Press), 837–893.

Kreisel, G. (1970). 'Principles of Proof and Ordinals Implicit in Given Concepts', in A. Kino, J. Myhill, and R. E. Vesley (eds.), *Intuitionism and Proof Theory* (Amsterdam: North-Holland), 489–516.

Kripke, S. (1975). 'Outline of a Theory of Truth', *Journal of Philosophy*, 72: 690–716. Reprinted in Martin (1984).

Lewis, D. (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8: 339–359. Reprinted in Lewis (1983).

——— (1983). *Philosophical Papers*, vol. 1 (Oxford: Oxford University Press).

Martin, R. L. (ed.) (1984). *Recent Essays on Truth and the Liar Paradox* (Oxford: Oxford University Press).

McGee, V. (1991). *Truth, Vagueness, and Paradox* (Indianapolis: Hackett).

——— (2000). 'Everything', in G. Sher and R. Tieszen (eds.), *Between Logic and Intuition* (Cambridge: Cambridge University Press), 54–78.

Parsons, C. (1974*a*). 'The Liar Paradox', *Journal of Philosophical Logic*, 3: 381–412. Reprinted in Parsons (1983).

——— (1974*b*). 'Sets and Classes', *Nous*, 8: 1–12. Reprinted in Parsons (1983).

——— (1983). *Mathematics in Philosophy* (Ithaca: Cornell University Press).

——— (this volume). 'The Problem of Absolute Generality', in A. Rayo and G. Uzquiano (eds.), *Absolute Generality* (Oxford: Oxford University Press).

Partee, B. H. (1991). 'Topic, Focus and Quantification', *Semantics and Linguistic Theory*, 1: 257–280.

Poesio, M. and Zucchi, A. (1992). 'On Telescoping', *Semantics and Linguistic Theory*, 2: 347–366.

Rayo, A. (2003). 'When Does 'Everything' Mean Everything?', *Analysis*, 63: 100–106.

Richter, W. and Aczel, P. (1974). 'Inductive Definitions and Reflecting Properties of Admissible Ordinals', in J. E. Fenstad and P. G. Hinman (eds.), *Generalized Recursion Theory* (Amsterdam: North-Holland), 301–381.

Roberts, C. (1995). 'Domain Restriction in Dynamic Semantics', in E. Bach, E. Jelinek, A. Kratzer, and B. H. Partee (eds.), *Quantification in Natural Languages* (Dordrecht: Kluwer), 661–700.

——— (1996). 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics', *Ohio State University Working Papers in Linguistics*, 49: 91–136.

Sæbø, K. J. (1999). 'Discourse Linking and Discourse Subordination', in P. Bosch and R. van der Sandt (eds.), *Focus: Linguistic, Cognitive, and Computational Perspectives* (Cambridge: Cambridge University Press), 322–335.

Stalnaker, R. C. (1978). 'Assertion', in P. Cole (ed.), *Pragmatics*, vol. 9 of *Syntax and Semantics* (New York: Academic Press), 315–322. Reprinted in Stalnaker (1999).

——— (1998). 'On the Representation of Context', *Journal of Logic, Language, and Information*, 7: 3–19. Reprinted in Stalnaker (1999).

——— (1999). *Context and Content* (Oxford: Oxford University Press).

Stanley, J. (2000). 'Context and Logical Form', *Linguistics and Philosophy*, 23: 391–434.

Stanley, J. and Szabó, Z. G. (2000). 'On Quantifier Domain Restriction', *Mind and Language*, 15: 219–261.

Stanley, J. and Williamson, T. (1995). 'Quantifiers and Context-Dependence', *Analysis*, 55: 291–295.

Tarski, A. (1935). 'Der Wahrheitsbegriff in den formalizierten Sprachen', *Studia Philosophica*, 1: 261–405. References are to the translation by J. H. Woodger as "The Concept of Truth in Formalized Languages" Tarski (1983).

——— (1983). *Logic, Semantics, Metamathematics* (Indianapolis: Hackett), 2nd edn. Edited by J. Corcoran with translations by J. H. Woodger.

Thomason, R. H. (1990). 'Accommodation, Meaning, and Implicature', in P. R. Cohen, J. Morgan, and M. E. Pollack (eds.), *Intentions in Communication* (Cambridge: MIT Press), 325–363.

van der Sandt, R. A. (1992). 'Presupposition Projection as Anaphora Resolution', *Journal of Semantics*, 9: 333–377.

van Kuppevelt, J. (1995). 'Discourse Structure, Topicality and Questioning', *Journal of Linguistics*, 31: 109–147.

von Fintel, K. (1994). *Restrictions on Quantifier Domains*, Ph.D. dissertation, University of Massachusetts at Amherst.

Westerståhl, D. (1985*a*). 'Determiners and Context Sets', in J. van Benthem and A. ter Meulen (eds.), *Generalized Quantifiers in Natural Language* (Dordrecht: Foris), 45–71.

——— (1985*b*). 'Logical Constants in Quantifier Languages', *Linguistics and Philosophy*, 8: 387–413.

——— (1989). 'Quantifiers in Formal and Natural Languages', in D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. IV (Dordrecht: Kluwer), 1–131.

Williamson, T. (2004). 'Everything', *Philosophical Perspectives*, 17: 415–465.