To appear in J. Hawthorne, ed., *Philosophical Perspectives* vol. on philosophy of language

Everything

Timothy Williamson

I

I am about to take a flight. Everything is packed into my carry-on baggage.

On reading the last sentence, did you interpret me as saying falsely that *everything* — everything in the entire universe — was packed into my carry-on baggage? Probably not. In ordinary language, 'everything' and other quantifiers ('something', 'nothing', 'every dog', ...) often carry a tacit restriction to a domain of contextually relevant objects, such as the things that I need to take with me on my journey. Thus a sentence of the form 'Everything Fs' is true as uttered in a context C if and only if everything that is relevant in C satisfies the predicate 'F'; 'everything' ranges just over the contextually relevant things. Such generality is restricted in a context-relative way.

Is there also absolute generality, without contextual restrictions? In that sense, absolutely everything Fs only if everything that is relevant in any context Fs. To use 'everything' to express absolute generality, we need a context in which absolutely nothing is excluded as irrelevant. Are there such contexts?

Bradley describes metaphysics as 'the effort to comprehend the universe, not simply piecemeal or by fragments, but somehow as a whole' ([AR]: 1). How could we comprehend the universe as a whole in a context in which some parts of it were irrelevant? Although Bradley does not expect us as metaphysicians to describe each part individually, each part makes a difference to the whole that we must comprehend.

At the core of metaphysics is ontology. Quine poses the ontological problem by asking 'What is there?', and answers, correctly but uninformatively. 'Everything' ([FLPV]: 1). To interpret his question and answer as restricted to a domain that excludes some contextually irrelevant things would be to misunderstand Quine by losing the total generality of the problem that he means to raise. He intends a context in which absolutely nothing is excluded as irrelevant. Everything whatsoever contributes to what there is.

Consider an example of a more specific metaphysical theory: out-and-out, no-holds-barred ontological naturalism, as in the slogan 'Everything is part of the natural world', in brief, 'Everything is natural'. To interpret such naturalists as leaving it open that there are some contextually irrelevant non-natural things would be to miss their point, by failing to appreciate the radical extent of their claim (whether it is true or false). To understand them properly, one must interpret them as generalizing without any restriction whatsoever, even if through so doing one comes to regard them as mistaken. They do not want our charity if it would interpret them as speaking truly of less than absolutely everything. They are playing for higher stakes than that.

The next two sections develop a naive conception of absolute generality. 'Naive' here does not imply 'false'; the conception is naive because it does not confront the most serious reasons for scepticism about absolute generality, reasons which seem to motivate a relativist conception of generality. Section IV sketches those reasons. Section V discusses a difficulty in stating generality-relativism coherently. Sections VI-VIII argue that such generality-relativism is far more disturbing in its implications than is often realized. The final section extends the naive conception by adumbrating a response to the argument of section IV. The upshot is a provisional argument for both the necessity and the possibility of the concept of absolutely everything.

Throughout sections II-III, the naive theorist of absolute generality is speaking.

II

I am not so naive as all that. I know that the word 'everything' can express a merely context-relative generality, however emphatically one utters it and however hard one hammers the table with one's fist. [1] Not even the addition of the word 'absolutely' forces it to range over more than comparatively few things: 'Of course I'm late —you left me to pack *ABSOLUTELY EVERYTHING!*'. Although one can try using more words ('absolutely everything in the whole universe with no exceptions whatsoever'), in some contexts even they may express only a relative generality. It makes little difference whether one says the words out loud or in one's heart. Nevertheless, I expect that a good listener with a reasonable command of English and a modicum of intelligence will interpret my use of the word 'everything' when I talk in the vein above of Bradley, Quine and naturalism in the sense that I intend. To avert any residual misunderstanding, I will further illustrate my use of 'everything' below. My remarks will constitute no attempt to define 'every' and 'thing' in terms that are somehow more basic, because I doubt very much that any terms are more basic. Definitions must come to an end somewhere. My elucidations will be ultimately circular, but not uninformative. My aim is simply to give you enough exposure to my use to enable you to understand it easily by the direct method. Although in principle a reader of gross hermeneutic obtuseness could still misunderstand me, I am confident that you will not.

If you doubt that this way of proceeding is acceptable, consider an analogy. After very little interaction with a normal human being, you can know that by '+' she means plus. Consequently, by '+' she does not mean quus, another arithmetical operation that coincides with plus over the cases that you discussed with her but diverges from plus for some larger numbers (Kripke [WRPL]). Even if the hypothesis that she meant quus predicts the very assents and dissents that you witnessed, only a dogmatic sceptic would deny that you know enough about her to know that she meant plus, not quus, even before you have tested her on the cases over which the two hypotheses diverge. Although someone could mean quus by '+', they could do so only perversely, by unnatural *ad hoc* means (such as explicit definition or extra complexity in a hard-wired algorithm) that are not required for meaning plus by '+'. The internal state and external environment of a normal human speaker in a society that uses '+' in the normal way are sufficient for meaning plus by '+'. It is not necessary to form a separate intention for each pair of natural numbers about the result of applying the operation. We find it hard to explain in detail why the ordinary conditions are sufficient, and how we know of speakers who meet them that they do mean plus by '+', but it is a good maxim in philosophy that failure to explain something obvious does not justify scepticism about it. Similarly, after very little interaction with me in a suitable context, you can know that by 'everything' I meant absolutely everything (or, more precisely, that I used the word in that context to express that content). Consequently, I did not mean absolutely everything except a few grains of sand in Australia that did not come up in our discussion. Even if the hypothesis that I meant everything except those grains of sand predicts the very assents and dissents that you witnessed, you know enough about me to know that I meant absolutely everything, not absolutely everything except those grains of sand, even before you ask me about them.

Although someone could mean absolutely everything except those grains of sand by ‘everything’, they could do so only perversely, by unnatural *ad hoc* means, that are not required for meaning absolutely everything by ‘everything’. My internal state and external environment on that occasion were sufficient for meaning that. It is not necessary to form a separate intention for each thing to make the quantifier range over it. It is hard to explain in detail why the ordinary conditions are sufficient, and how you know that by ‘everything’ I do mean absolutely everything, but remember the maxim that failure to explain something obvious does not justify scepticism about it.[2]

Now let me continue with my elucidations. Henceforth I will always use the word ‘everything’ in the absolutely general way that I am elucidating. As usual, ‘Nothing Fs’ is equivalent to ‘Everything not-Fs’, and ‘Something Fs’ to ‘Not everything not-Fs’. ‘There is an F’ will be used as a variant of ‘Something is an F’.

‘Everything Fs’ and ‘Something Fs’ can be symbolized in quantificational logic as $\forall x\ \alpha$ and $\exists x\ \alpha$, where the variable ‘x’ occurs in name position in the formula α; α may be understood as ‘x Fs’. In contemporary model theory, each model has a domain, a nonempty set; when formulas are interpreted with respect to that model, quantifiers such as $\forall x$ and $\exists x$ are read as restricted to the members of that domain. Thus the domain of the model acts as the contextual restriction. In the set theory on which contemporary model theory relies, there is no universal set, no set of which everything is a member. Consequently, the quantifiers are never interpreted as ranging over everything. So my use of the quantifiers cannot be interpreted with respect to such models.[3] Nevertheless, I can state the truth-conditions for quantified formulas, used in my way. As usual, it is done for truth under an assignment of values to variables. For any assignment A, variable x and thing d, let A[x/d] be the assignment just like A except that it assigns d to x.

[∀] $\forall x\ \alpha$ is true under A if and only if everything d is such that α is true under A[x/d].

[∃] $\exists x\ \alpha$ is true under A if and only if something d is such that α is true under A[x/d].

Naturally, 'everything' and 'something' in these clauses must be read unrestrictedly. We can easily use the unrestricted quantifiers to recover quantification over a restricted domain by introducing a monadic predicate *D* that applies to all and only the members of the domain. The formulas $\forall x\ (Dx \supset \alpha)$ and $\exists x\ (Dx\ \&\ \alpha)$ serve as the universal and existential quantifications of α restricted to the domain. If we want, we can add the axiom $\exists x\ Dx$ to express the nonemptiness of the restricted domain. On this analysis, unrestricted quantification turns out to be simpler and more basic than restricted quantification. Since the restricted quantifiers are decomposed into unrestricted quantifiers and other constituents, understanding restricted quantification involves understanding unrestricted quantification.

The same holds for contextual restrictions in natural languages. When they make a difference to what is said, they do so by contributing tacit constituents to semantic structure. When one says that everything (in my sense) Fs, one is not just generalizing over contextually relevant items in a context in which everything happens to be relevant. If relevance in my context is parasitic on relevance in your context, and everything is relevant in your context, then everything is relevant in my context, even if I am not in a position to know that everything is relevant in my context; when I utter the sentence 'Everything Fs' subject to that contextual restriction, which is in fact universally satisfied, I do not mean that everything Fs. When one says that everything Fs, one is simply speaking in a context in which no special features trigger the insertion of a restricting constituent into the semantic structure (see also McGee [E]: 77). That one failed to think of something when one spoke does not imply that it is contextually excluded; forgetting counterexamples is not avoiding them. But if one specifically intends a restriction to certain things (as I do not when I speak of everything), then perhaps that intention can succeed, and the restriction is entered into semantic structure.

By contrast, contemporary model theory treats domains like unarticulated background constraints, external to the semantic structure of formulas. But examples from natural languages indicate that no such general restriction will capture the phenomena:

(!) Every passenger packed everything into carry-on baggage.

The natural reading of (!) is roughly that every passenger packed everything that was suitably relevant to him or her into his or her carry-on baggage (for present purposes we can ignore contextual restrictions on 'every passenger'). Different passengers packed numerically different things. That is quite different from the claim that every passenger packed everything in the domain of the context into his or her carry-on baggage, which implies that the passengers all packed numerically the same things. To capture the effect, we need the contextual restriction on 'everything' to be bound by 'every passenger', and therefore to contribute to semantic structure.[4] Even in contemporary model-theoretic semantics, the restriction to the domain of the model is made explicit in the meta-language, which thereby exhibits complexity in the model-theoretic readings of $\forall x$ and $\exists x$.[5]

In some obvious ways, natural languages differ in their apparatus of generality from languages with $\forall x$ and $\exists x$. Natural languages have determiners such as 'every', 'some' and 'no', which combine with singular or plural nouns, perhaps qualified by adjectives or relative clauses (category N), to form explicitly restricted quantifiers ('every donkey', 'a brown donkey', 'no donkey that I have ever seen'). Even 'everything', 'something' and 'nothing' decompose into 'every thing', 'some thing' and 'no thing'. The restricted quantifiers then combine with predicates (category VP) to form sentences: 'Every donkey brays', 'A brown donkey ate thistles', 'No donkey that I have ever seen talks'. Sometimes those sentences can be paraphrased by means of monadic quantifiers such as $\forall x$ and $\exists x$; thus 'Every donkey brays' can be formalized as $\forall x\ (Dx \supset Bx)$. But such paraphrases are not always possible. If 'M' is read as 'most things are such that', 'Most donkeys talk' cannot be formalized as $\mathrm{M}\ x\ (Dx \supset Tx)$ (which is true because most things are not donkeys) or $\mathrm{M}x\ (Dx\ \&\ Tx)$ (which implies that most things are donkeys) or in any such way. We need a structure like $(\mathrm{M}x{:}\ Dx)\ Tx$, to separate the inputs of Dx and Tx, so that the semantics of $\mathrm{M}x$ can be sensitive to whether most of the things that satisfy the first input also satisfy the second. But the role of restricted quantification in the apparatus of generality in natural languages does not mean that they constitute less hospitable environments for absolute generality. For, as already noted, determiners can be combined with a noun such as 'thing', which in the relevant sense applies unconditionally to whatever the determiners range over. We can regard 'thing' as a logical constant, a universally applicable noun, and treat the determiners as carrying no restriction in themselves.[6] Thus 'Everything Fs' is equivalent to '(Every *x*: *x* is a thing) *x* Fs', where the determiner 'every' can range over everything and '*x* is a thing' is a trivial logical truth for each value of '*x*'. When context does impose a restriction on 'every' or 'thing', that is a phenomenon of the kind already dealt with. For most purposes we can ignore the structural ways in which natural languages slightly complicate the expression of absolute generality.[7]

In case it is not already obvious, let me emphasize that one can never defend the claim that everything Fs against a purported counterexample by admitting that it does not F but denying that it is a thing. If the word 'it' there refers at all, it refers to a thing in the relevant sense; if 'it' does not refer, nothing has been denied to be a thing. Whatever is is a thing. If there were any non-things, they too would be things: so there are no non-things. In any sense of 'exist' in which there are non-existents, they are things just as much as existents are. Any natural or unnatural kind or substance is a thing; so too is any member of the kind or sample of the substance. Whatever is abstract or concrete or neither is a thing. Whatever is basic or derived, simple or complex, is a thing. Whatever can be named is a thing; so too is whatever cannot be named. Any value of a variable is a thing, and everything is the value of a variable under at least one assignment (compare [FLPV]: 13).

You might fear that this absolute promiscuity of 'everything' makes it useless after all for the purposes of a metaphysician with a restrictive ontological thesis, such as ontological naturalism. For if 'Everything is natural' has the obviously untrue consequence 'The Platonic Form of the Good is natural', then ontological naturalism as so formalized is trivially false, and the would-be ontological naturalist must formulate the view otherwise, presumably with a restricted quantifier. But that objection to the use of the absolute 'everything' is fallacious. The naturalist should simply deny that 'Everything is natural' entails 'The Platonic Form of the Good is natural', on the grounds that 'The Platonic Form of the Good' denotes nothing whatsoever. The naturalist should not even assent to '"The Platonic Form of the Good" denotes the Platonic Form of the Good'. The argument from a universal generalization to an instance formed with an empty term is not in general valid, even if there is no restriction on which things are values of the universally quantified variable.

Whatever has a name is a thing; but it does not follow that if a sentence of the form 'Everything Fs' is true, then so is any sentence of the form 'n Fs', where 'n' is (replaced by) a name. For the name 'n' may not refer. I assent to the truism 'Everything is a thing', but I withhold assent from 'John Smith-Jones is a thing', because the name 'John Smith-Jones' does not refer (as used in this context). For the same reason, I withhold assent from 'John

Smith-Jones is non-existent' and 'John Smith-Jones is named "John Smith-Jones"'. But if a sentence of the form 'Everything Fs' is true, and the name 'n' does refer, then the sentence 'n Fs' is true too. Of course, to say that a name refers is to say that it refers to something; what else could it mean? Every candidate for being succeeds in being, but not every candidate for referring succeeds in referring. That goes for reference in thought too. Whatever one can think of is a thing, but that does not guarantee that whenever one appears to be thinking of something, one really is thinking of something. Equally, whatever one cannot think of is a thing. One's generalizations are not restricted to those things that one can single out. There are many things that I cannot think of. Inevitably, I cannot think of an example.

If there are fictional Xs, then they are things too, whether or not they are Xs, and whether or not they are real Xs (if that is a different question). But that does not commit me to the claim that if, according to a story, there is a golden mountain, then there is a golden mountain, any more than it commits me to the claim that if, according to a story, the streets of London were paved with gold, then the streets of London were paved with gold. The inference schema from 'According to the story, P' to 'P' is invalid. Of course, if I were telling the story I would say 'There is a golden mountain' or 'The streets of London were paved with gold'. But I am not currently engaged in story-telling. I am making serious assertions; I intend them to be true, not just true according to a fiction. Nevertheless, I do not intend them to be true by virtue of any restriction that would exclude a fictional golden mountain as irrelevant. In the special case of existential sentences, someone might try to bridge the gap between 'According to the story, P' and 'P' by defending the principle that if, according to a story, something is an X, then some fictional (or perhaps real) thing is an X, in which case something is an X. For many values of 'X' that conclusion looks hopeless: there just is no golden mountain in Scotland, however vividly I tell stories according to which there is a golden mountain in Scotland. More to the point, the unrestricted reading of the quantifiers in no way forces me into the principle that if, according to a story, something is an X, then some fictional (or perhaps real) thing is an X; that is an extreme step in the metaphysics of fiction and quite independent of the pure logic of unrestricted generality. The unrestricted reading of the quantifiers does not even force me into the principle that if, according to a story, something is an X, then some fictional (or perhaps real) thing is such that, according to the story, it is an X, a principle which yields the Barcan formula for fiction: if, according to the story, something is an X, then something is such that, according to the story, it is an X. Since a non-X might be such that, according to the story, it is an X, for example, an innocent man might be such that, according to the story, he is a murderer, that is a less extreme step in the metaphysics of fiction; but it is still a substantive step, and one independent of the pure logic of unrestricted generality.

If there are possible Xs, then they are things too, whether or not they are Xs, and whether or not they are actual Xs (if that is a different question). But that does not commit me to the claim that if there could have been a golden mountain, then there is a golden mountain, any more than it commits me to the claim that if I could have been a murderer, then I am a murderer. The inference schema from 'Possibly P' to 'P' is invalid. In the special case of existential sentences, someone might try to bridge the gap by defending the principle that if possibly something is an X, then some possible (perhaps actual) thing is an X, in which case something is an X. For many values of 'X' that conclusion looks hopeless: there just is no city on the moon in 2003, however easily there could have been one. More to the point, the unrestricted reading of the quantifiers in no way forces me into the principle that if possibly something is an X, then some possible (perhaps actual) thing is an X; that is an extreme step in the metaphysics of modality and quite independent of the pure logic of unrestricted generality. The unrestricted reading of the quantifiers does not even force me into the principle that if possibly something is an X, then some possible (perhaps actual) thing is such that possibly it is an X, a principle which yields the Barcan formula for possibility: if possibly something is an X, then something is such that possibly it is an X. Since a non-X might be such that possibly it is an X, for example, an innocent man could have been a murderer, that is a less extreme step in the metaphysics of modality; but it is still a substantive step, and one independent of the pure logic of unrestricted generality. If there could not have been anything other than what there actually is, then the Barcan formula for possibility is correct: but our current task is not to decide whether it is correct.[8]

If there are past or future Xs, then they are things too, whether or not they are Xs, and whether or not they are present Xs (if that is a different question). But that does not commit me to the claim that if there was or will be a golden mountain, then there is a golden mountain, any more than it commits me to the claim that if it was raining, then it is raining. The inference schema from 'It was the case that P' or 'It will be the case that P' to 'P' is invalid. In the special case of existential sentences, someone might try to bridge the gap by defending the principle that if it was or will be the case that something is an X, then some past or future (or perhaps present) thing is an X, in which case something is an X. For many values of 'X' that conclusion looks hopeless: there just is no king of England, although there was one and probably will be one again. More to the point, the unrestricted reading of the quantifiers in no way forces me into the principle that if it was or will be the case that something is an X, then some past or future (or perhaps present) thing is an X; that is an extreme step in the metaphysics of time and quite independent of the pure logic of unrestricted generality. The unrestricted reading of the quantifiers does not even force me into the principle that if it was or will be the case that something is an X, then some past or future (or perhaps present) thing is such that it was or will be an X, a principle which yields the Barcan formulas for the past and future: if it was the case that something is an X, then something is such that it was an X, and if it will be the case that something is an X, then something is such that it will be an X. Since a non-X might be such that it was or will be an X, for example, an adult was once a child, that is a less extreme step in the metaphysics of time; but it is still a substantive step, and one independent of the pure logic of unrestricted generality. If there never was and never will be anything other than what there actually is, then the Barcan formulas for the past and future are correct: but our current task is not to decide whether they are correct.

Far from prejudging issues in the metaphysics of fiction, modality or time, the use of unrestricted quantification enables one to articulate those issues. Are there fictional Xs? Are there merely possible Xs? Are there past or future Xs? The questions are worth asking only if they do not embody contextual restrictions that would exclude fictional Xs, or merely possible Xs, or past or future Xs respectively.

III

The metaphysical grandeur of absolutely universal quantification might lead one to fear that it would be inferentially isolated from our more mundane concerns. Fortunately, the logic of quantification integrates even the unrestricted variety with the rest of our thought. Indeed, if all restricted quantifiers are semantically complex combinations of unrestricted quantifiers and restricting conditions, as was proposed in the previous section, then all arguments that involve restricted quantification thereby semantically involve unrestricted quantification. But a few examples will demonstrate the logical integration of unrestricted quantification with the rest of the language, independently of the semantic analysis of restricted quantification.

For simplicity, let us consider the universal quantifier $\forall$, as interpreted by the unrestricted clause [$\forall$]. Although the usual soundness proofs for first-order logic demonstrate the validity of the standard axioms and rules of inference only with respect to the domain-relative model-theoretic semantics, the same axioms and rules remain evidently valid on the absolutely general reading. If that reading also validates further axioms and rules of inference, let us ignore them for present purposes.[9]

Since any universal generalization is refuted by a counterexample, $\sim\forall x\, Rx$ is a logical consequence of the singular sentence $\sim Rt$, and $\sim\forall x \sim Rt$ is a logical consequence of the singular sentence Rt (for simplicity, empty singular terms are ignored in this section). If you assert that this page is rectangular, you can deduce that something is rectangular; if you deny that this page is rectangular, you can deduce that something is not rectangular. By taking a stand on the singular question, you are committed to a stand on an absolutely general question. Such negative commitments are scarcely surprising: as the range of generality widens, it becomes easier to deny and harder to assert a universal generalization.

It is more interesting to consider how speakers can incur positive logical commitments to absolutely universal generalizations. Some are logical truths, and therefore logical consequences of any premises whatsoever: for example $\forall x \sim(Rx \,\&\, \sim Rx)$ and $\forall x\, (\forall y\, Ry \supset Rx)$. Nothing is both rectangular and not rectangular; anything is rectangular if everything is. Whatever one says, one is logically committed to such claims. Other universal generalizations depend on specific singular premises: for example, $\forall x\, (\sim Rx \vee Rt)$ and $\forall x\, (Rx \vee \sim Rt)$ are logical consequences of Rt and $\sim Rt$ respectively without being logical truths themselves. If you take a stand on the singular question whether this page is rectangular, you are committed either to the claim that nothing is rectangular unless this page is rectangular or to the claim that everything is rectangular unless this page is not rectangular.

Once we have a few universal generalizations, we can use them to extract further universal generalizations from singular claims. For example, given that the longer-than relation is universally transitive ($\forall x\ \forall y\ \forall z\ ((Lxy\ \&\ Lyz) \supset Lxz)$), from the premise that this is longer than that (Lab) one can deduce that everything is longer than that if longer than this ($\forall x\ (Lxa \supset Lxb)$). Even lengthless objects satisfy the generalization vacuously. No contextual restrictions are needed to secure the truth of such universal generalizations.

The logic of identity provides further examples. Since the reflexivity of identity is a logical truth, we can assert that everything is self-identical ($\forall x\ x{=}x$). Similarly, since the indiscernibility of identicals is logically valid, from the identity statement that Hesperus is Phosphorus we can deduce the universal generalization that everything orbits Phosphorus if it orbits Hesperus ($\forall x\ (Oxh \supset Oxp)$ follows logically from $h{=}p$). Again, from the singular statement that I am sitting, I can deduce that everything is such that if it is not sitting then it is not me ($\forall x\ ({\sim}Sx \supset {\sim}x{=}i)$ follows logically from Si).

Set theory yields more cases. Nothing is a member of the null set ($\forall x\ {\sim}x{\in}\varnothing$). Someone who leaves it open that some contextually irrelevant things might be members of the null set has not yet mastered elementary set theory. Similarly, on the standard view of sets, nothing whatsoever is a member of itself ($\forall x\ {\sim}x{\in}x$).

In ordinary language, natural laws can be expressed with natural kind terms complementing the determiner ‘every’: as a simple example, ‘Every whale is a mammal’. Such generalizations entail the corresponding generalizations with the absolute ‘everything’ and the conditional: ‘Everything is such that it is a whale only if it is a mammal’ ($\forall x\ (Wx \supset Mx)$).

Such examples can be multiplied indefinitely. They show that large numbers of absolutely universal generalizations are within easy deductive reach from our mundane knowledge. Prior to any distinctively metaphysical theorizing, we are committed to more than merely context-relative generalizations. Unrestricted generality is inescapable.

IV

The naive theorist of absolute generality has spoken. The naive conception has been subject to scepticism on various grounds, many of them weak. By far the strongest grounds for scepticism are associated with the paradoxes of set theory, in particular Russell's paradox of the set of all sets that are not members of themselves (is it a member of itself?). But the bearing of the paradoxes on the naive conception is not obvious.

Russell's paradox reveals an inconsistency in the Naive Comprehension principle. That principle allows us to substitute any predicate for 'F' in the schema: there is a set of which everything is a member if and only if it Fs. Substituting 'is a set that is not a member of itself' in the schema yields an inconsistent instance. The usual response is in effect to impose a limitation of size on sets, by eliminating axioms that would imply the existence of sets with "too many" members (more precisely, with as many members as there are sets). One still maintains a restricted form of Naive Comprehension, the Separation principle, which allows us to substitute any predicate for 'F' in the schema: for every set *x*, there is a set *y* of which everything is a member if and only if it both Fs and is a member of *x*. Separation is consistent with limitation of size, for the new set *y* will have at most as many members as the old set *x*. But Separation implies that there is no universal set, no set of which everything is a member, for otherwise we could take x to be the universal set, recover Naive Comprehension and derive Russell's paradox. Some theories postulate proper classes, too big to be sets, with the more liberal Class Comprehension principle, which allows us to substitute any predicate for 'F' in the schema: there is a class of which every set is a member if and only if it Fs. Class Comprehension implies that there is a proper class of which every set is a member, but it does not imply that there is a class of which everything is a member, a genuinely universal class, for it does not imply that proper classes are ever members of classes.

What has all this to do with unrestricted generality? It shows that if we generalize over everything, we are not generalizing over just the members of a set or class, as sets and classes are now standardly conceived. If a domain is a set or class, as now standardly conceived, then we are not generalizing over just the members of a domain. But the naive conception has already assimilated that point: it denies that when we generalize over everything, we generalize over just the members of a domain. Richard Cartwright ([SE]) has forcefully criticized the assumption that whenever we generalize coherently there must be something, a domain, to which everything over which we are generalizing belongs; he calls it the All-in-One Principle.[10] Moreover, the standard conceptions of sets and classes are not the only ones available. Some non-standard set theories postulate a universal set, at the cost of restricting the Separation principle.[11] So far we have no urgent threat to the naive conception of unrestricted generality.

Nevertheless, there is a problem. Sooner or later the naive theorist will want to generalize over all (legitimate) interpretations of various forms in the language. For example, the inference from $\forall x\, Px$ and $\forall x\, (Px \supset Qx)$ to $\forall x\, Qx$ is truth-preserving however one interprets the predicate letters P and Q. Such generalizations are the basis of Tarski's account of logical consequence [CLC] and its model-theoretic descendants.[12] Whatever their relation to intuitive ideas of logical consequence, generalizations about the truth-preservation of inferences under all interpretations of their non-logical constituents are interesting and important in their own right. The naive theorist wants to make such generalizations when $\forall x$ is read as unrestricted. In principle, when we apply the definition of logical consequence, it must be possible to interpret a predicate letter according to any contentful predicate, since otherwise we are not generalizing over all the contentful arguments of the right form. Thus, whatever contentful predicate we substitute for 'F', some legitimate interpretation (say, I(F)) interprets the predicate letter *P* accordingly:

(1) For everything o, I(F) is an interpretation under which *P* applies to o if and only if o Fs.

On the naive theorist's maximally liberal understanding of 'thing', even an interpretation such as I(F) counts as a thing: to claim that it is not a thing would be self-defeating. Now define a verb 'R' thus:

(2) For everything o, o Rs if and only if o is not an interpretation under which *P* applies to o.

The naive theorist is committed to treating 'R' as a contentful predicate, since it is well-formed out of materials entirely drawn from the naive theory itself. Now put 'R' for 'F' in (1) and apply (2):

(3) For everything o, I(R) is an interpretation under which *P* applies to o if and only if o is

not an interpretation under which *P* applies to o.

In particular, since ‘everything’ is unrestricted, o can be I(R) itself. Thus (3) implies:

(4) I(R) is an interpretation under which *P* applies to I(R) if and only if I(R) is not an interpretation under which *P* applies to I(R).

But (4) is an instance of an inconsistent form in propositional logic. Although the argument is obviously a variant of Russell’s Paradox, it employs no notion of set, class or domain, standard or non-standard. It does not assume that everything, or everything that Rs, belongs to a single thing of any kind. It assumes about interpretations only what is needed for generalizing over them to capture the naive theorist’s intuitive idea of generalizing over all contentful arguments of the right form in the definition of logical consequence.

Paradox does not arise as soon as we interpret a language as expressing unrestricted. generalizations. It arises once we start to form general about such interpretations. By themselves, the truth-theoretic clauses [∃] and [∀] for unrestricted quantifiers lead to no contradiction. Formally, they look exactly like the standard clauses. We might call (1)-(4) a paradox of logical truth rather than a paradox of truth. However, once we interpret a language in a certain way, we are led to reflect on the validity of forms of inference in such a language, and so to survey the variety of interpretations of that kind. When that reflective process ends in paradox, the coherence of the original interpretation is called into question.[13]

What has gone wrong? A sceptic about absolute generality will suggest that generalizations in the original object-language (such as $\forall x\ Fx$) cannot coherently be interpreted as ranging over such things as interpretations of that very language. Since the meta-linguistic schema (1) gives the application-condition of a predicate letter *P* of the original language under an interpretation I(F), it should be restricted to those things over which generalizations in the object-language do range, for *P* does not apply under an interpretation to something over which those generalizations do not range, whether it Fs or not. Therefore, the substitution of 'I(R)' for the variable 'o' in the step from (3) to (4) is illicit, and the paradox is blocked.

On the envisaged alternative to the naive conception, when we move from an object-language L to a meta-language L* in which we can generalize adequately over interpretations of L, our range of generality widens. Reflecting in L* on L, we realize that L did not express absolute generality. Whatever language we speak, we can always come to realize that it does not express absolute generality by moving to such a meta-language. Therefore, absolute generality is inexpressible. If it is inexpressible, it is also unthinkable, for if it were thinkable, it would have been expressed in section II. Only relative generality is left. For the sake of argument, let us provisionally suppose that the naive conception does collapse like that. We must explore the generality-relativist alternative.

V

The generality-relativist faces a problem of a kind that Wittgenstein sketches in the Preface to *Tractatus Logico-Philosophicus*: 'in order to be able to set a limit to thought, we should have to find both sides of the limit thinkable (i.e. we should have to be able to think what cannot be thought)' ([TLP]: 3). *What*, according to the generality-relativist, cannot be thought?

David Lewis puts the problem in a nutshell when he points out, of someone who asserts ‘that some mystical censor stops us from quantifying over everything without restriction’, ‘Lo, he violates his own stricture in the very act of proclaiming it!’ ([PC]: 68). According to the generality-relativist:

(5) It is impossible to quantify over everything.

The generality-relativist, who does not pretend to do the impossible, therefore admits:

(6) I am not quantifying over everything.

By simple quantificational logic, (6) implies:

(7) Something is not being quantified over by me.

Conceding the argument so far, the generality-relativist therefore utters (7) at a time t_0. Suppose that the generality-relativist speaks truly. Then a standard principle of semantic ascent for truth yields:

(8) ‘Something is not being quantified over by me’ is true as uttered by the generality-relativist at t_0.

Obvious semantic principles for the quantifier and predicate in (8) are these:

(9) ‘Something Fs’ is true as uttered by *s* at *t* if and only if something over which *s* is

quantifying at t satisfies 'Fs' as uttered by s at t.

(10) Something satisfies 'is not being quantified over by me' as uttered by s at t if and only if it is not being quantified over by s at t.

From (8) and (9) we have:

(11) Something over which the generality-relativist is quantifying at t_0 satisfies 'is not being quantified over by me' as uttered by the generality-relativist at t_0.

From (10) and (11) we have:

(12) Something over which the generality-relativist is quantifying at t_0 is not being quantified over by the generality-relativist at t_0.

Of course, (12) is inconsistent. Thus the generality-relativist did not speak truly in uttering (7). But since the generality-relativist was committed to (7) by (5), the claim that it is impossible to quantify over everything is self-defeating. In brief, if 'everything' in (5) is restricted then even the generality-relativist agrees that (5) is false, since restricted quantification is possible by the generality-relativist's own lights; but if 'everything' in (5) is unrestricted, then (5) is an instance of what the generality-relativist claims to be impossible.

Can the generality-relativist use the distinction between object-language and meta-language to construct a more circumspect formulation of generality-relativism? In a meta-language L* for a language L, the generality-relativist can say:

(5L) It is impossible to quantify in L over everything.

Since (5L) is a sentence of L* but not of L, it is not self-defeating in the same way as (5). What stand to (8), (9) and (10) as (5L) stands to (5) are these sentences:

(8L) 'Something is not being quantified over in L by me' is true in L* as uttered by the generality-relativist at t_0.

(9L) 'Something Fs' is true in L* as uttered by *s* at *t* if and only if something over which *s* is quantifying in L* at *t* satisfies 'Fs' in L* as uttered by *s* at *t*.

(10L) Something satisfies 'is not being quantified over in L by me' in L* as uttered by *s* at *t* if and only if it is not being quantified over in L by *s* at *t*.

But from (8L), (9L) and (10L) we reach only a conclusion that the generality-relativist can happily accept:

(12L) Something over which the generality-relativist is quantifying in L* at t_0 is not being quantified over in L by the generality-relativist at t_0.

The problem with (5L) is that it is too weak to capture generality-relativism, for it merely states an expressive limitation of a particular language. For all that (5L) says, it is possible to quantify in L* over everything. The generality-relativist wants to add:

(5L*) It is impossible to quantify in L* over everything.

But if the generality-relativist utters (5L*) as a sentence of L*, the paradox recurs, since the qualification 'in L*' can be added uniformly throughout. The generality-relativist must utter (5L*) as a sentence of a meta-meta-language L**, and so on. What is self-defeating is a claim of the form:

(5M) It is impossible to quantify in my current language over everything.

The generality-relativist can put the point less drastically in terms of change of context rather than change of language, a formulation better suited to the flexibility of natural languages. In a context C* one can defensibly say:

(5C) Not everything is quantified over in C.

But (5C) is not true as uttered in C itself, nor is the corresponding (5C*) true as uttered in C*. It is always easier to point out other people's limitations than one's own. Generality-relativism appears not to be fully expressible in any context.

The generality-relativist is tempted by the idea that no context subsumes every context, and therefore by this generalization:

(13) For any context C0, there is a context C1 such that not everything that is quantified over in C1 is quantified over in C0.

For (13) to express the intended idea, the quantifiers must at least range over every context (although not everything). But (13) entails (5C), for it entails that there is a context C1 such that not everything that is quantified over in C1 is quantified over in C, which entails that not everything is quantified over in C. Since (5C) is not true as uttered in C, (13) is also not true as uttered in C. Since C can be any context, (13) is not true in its intended sense as uttered in any context.

The generality-relativist may try to do better by semantic ascent. In the original context C, one might utter not the self-defeating (5C) but this:

(5CM) 'Not everything is quantified over in C' is true as uttered in C*.

It is understood that 'C' as uttered in C* refers to C. Of course, the price in C of accepting (5CM) while rejecting (5C) is that one must give up a homophonic account of the truth-condition of the quoted sentence; in C, one rejects this biconditional:

(14) 'Not everything is quantified over in C' is true as uttered in C* if and only if not everything is quantified over in C.

The justification for rejecting (14) would be that the word 'everything' has an indexical element. As used in C (for example, on the right-hand side (14)) it expresses a different content from that which it expresses as used in C* (for example, as described on the left-hand side of (14)).

The meta-linguistic strategy can be extended to a generalization of (5CM) over all contexts:

(15) For every context C0, there is a context C1 such that 'Not everything is quantified over in C0' is true as uttered in C1 (where 'C0' as uttered in C1 refers to C0).

How adequately does (15) express generality-relativism? If the speaker does not know what the sentence 'Not everything is quantified over in C0' expresses in another context, then the claim that it is true as uttered in that other context is not very helpful. Can the generality-relativist state in C an appropriate truth-condition for the quoted sentence as uttered in C1? As already noted, the meta-linguistic strategy depends on the failure of a homophonic statement in one context of the truth-condition of such a sentence as uttered in a wider context. Nor should the generality-relativist try this:

(14*) 'Not everything is quantified over in C' is true as uttered in C* if and only if not everything that is quantified over in C* is quantified over in C.

For (14*) and (5CM) still entail (5C): trivially, if not everything that is quantified over in C* is quantified over in C, then not everything is quantified over in C. Since the meta-linguistic strategy depends on accepting (5CM) while rejecting (5C) in C, the generality-relativist cannot accept (14*) in C. The generality-relativist should not expect to be able to state in one context even non-homophonically the truth-condition of a generalization as uttered in a wider context, because that would involve transcending the limitations of the narrower context from inside. The generality-relativist may still claim to know the standing linguistic meaning of the quoted sentence in (15), but when the speaker has so little grip on the truth-condition of that sentence (15) becomes a somewhat tenuous expression of generality-relativism. The generality-absolutist can use unrestricted generalizations to state the truth-conditions of contextually restricted generalizations in other contexts; the generality-relativist must be careful not to fall back inadvertently on the absolutist conception.

A further source of unease about the meta-linguistic strategy is that generality-relativism was motivated in section IV by the idea that meta-linguistic reflection on the interpretations of one's language forces one into a language of wider generality: semantic ascent takes one to a wider context. By contrast, the point of (5CM) and (15) is to apply semantic ascent while remaining in the narrower context. This is not yet a sharp objection, but it points towards problems for generality-relativism that will be clarified in section VII.

The generality-relativist might try to avoid meta-linguistic formulations by treating 'in a context' as an operator rather than a meta-linguistic device and replacing (5CM) by something like:

(5CO) In C*, not everything is such that, in C, it is something.

The operator strategy depends on the invalidity of the inference from (5CO) to (5O):

(5O) Not everything is such that, in C, it is something.

For (5O), just like (5C), is not true as uttered in C. The operators 'in C' and 'in C*' in (5CO) and (5O) are supposed to behave somewhat like a counterfactual operator 'in such-and-such circumstances it would have been the case that'. To capture the generality of (13), the generality-relativist might try to use 'in every context' and 'in some context' as irreducible operators analogous to 'necessarily' and 'possibly' respectively, while still permitting them to bind other operators:

(13O) In every context C0, in some context C1, not everything is such that, in C0, it is something.

In particular, to acknowledge the limitation of the current context, the generality-relativist can say:

(16) In some context C1, not everything is such that, in this context, it is something.

Structurally, (16) is analogous to the claim in modal logic that possibly not everything is actually something: there could have been something other than what there actually is. That modal claim entails the rejection of the Barcan schema for possibility, that if possibly something Fs then something possibly Fs. For if we substitute 'is not actually something' for 'Fs', the antecedent of the resulting instance of the Barcan schema is the analogue of (16), while the consequent says that something is possibly not actually something, which is inconsistent: since everything is actually something, everything is necessarily actually something (the occurrence of 'necessarily' makes no difference because 'actually' is being understood rigidly). Analogously, (16) entails the rejection of the Barcan schema for contextuality, that if, in some context, something Fs then something is such that, in some contexts, it Fs. For if we substitute 'is not such that, in this context, it is something', the antecedent of the resulting instance of the Barcan formula is equivalent to (16), while the consequent says that something is such that, in some context, it is not such that, in this context, it is something, which is inconsistent: trivially, it is true to say that everything is such that, in this context, it is something, so everything is such that, in every context, in this context, it is something (the occurrence of 'in every context' makes no difference because 'in this context' is being understood rigidly).

The trouble with the operator strategy can be explained by analogy with the modal case. Those who regard modal operators as irreducible do not give up unrestricted quantification merely by rejecting the Barcan formula for possibility and accepting the modal analogues of claims such as (5CO), (13O) and (16). They reject the modal realist comparison between other possible worlds and other places (given that all places are equally real). In asserting that there *could have been* something other than what there actually is, they do not assert outright that there *is* something other than what there actually is. That a list could have been incomplete does not imply that it is incomplete. Analogously, those who regard contextual operators as irreducible do not give up unrestricted quantification merely by rejecting the Barcan formula for contextuality and accepting claims such as (5CO), (13O) and (16). They reject the contextual realist comparison between other contexts and other places (given that all places are equally real). In asserting that, in some context, there is something other than what, in the current context, there is, they do not assert outright that there is something other than what, in the current context, there is. That a list is such that, in some context, it is incomplete does not imply that it is incomplete. On the operator treatment, what is simply true may, in another context, be false. Thus the operator strategy does not provide an adequate articulation of generality-relativism. In this respect the meta-linguistic treatment looks more promising, because (5CM) says that the quoted sentence *is* true-as-uttered-in-C*.

The operator treatment of 'in a context' is independently problematic. The sentence 'I am Tony Blair' is true as uttered by Tony Blair, but I cannot report that in English by saying that in a context in which Tony Blair is speaking, I am Tony Blair. In no context am I him. In David Kaplan's terminology, phrases of natural language such as 'in a context' do not function as *monsters* ([D]: 510-11). At least in those respects, we cannot understand the operator as having the kind of effect that it would need to have in order to provide a workable alternative to meta-linguistic devices in the description of contextual shifts, rather than a mere stipulated notational variant. We may be unable to understand 'in a context' in the way that the operator strategy requires.

Generality-relativists seem to be unable to articulate their position in any coherent way. Can they respond by arguing that generality-absolutists face the same problem? Suppose that the generality-absolutist utters the sentences of which (5) and (6) are the negations:

(~5) It is possible to quantify over everything.

(~6) I am quantifying over everything.

Naturally, the generality-relativist will deny that the generality-absolutist can use (~5) to express generality-absolutism or (~6) to express a generality-absolutist claim that entails generality-absolutism. For 'everything' in (~5) and (~6) is used in a way that the generality-relativist finds coherent only if it is used with a contextual restriction, in which case (~5) and (~6) do not express generality-absolutism. However, what matters much more is whether the generality-absolutist can regard generality-absolutism as expressible. For the problem for generality-relativism was that by the generality-relativist's own lights it seems to be inexpressible. According to the generality-absolutist, in a context in which 'everything' is used unrestrictedly, (~5) *does* express generality-absolutism. The generality-absolutist can even regard (5) in the same context as an expression of generality-relativism, albeit not one that the generality-relativist can acknowledge as such.

The generality-relativist might respond by arguing that not even the generality-absolutist should regard (~5) or (~6) as expressing generality-absolutism. For even if 'everything' had been used restrictedly, (~5) and (~6) would still have been true. They are trivially self-verifying sentences, in no way distinctive of generality-absolutism; they do not express philosophically significant claims. Even if the word 'absolutely' were added to (~5) and (~6), it too could be given a restricted reading. But that generality-relativist response is fallacious. The generality-absolutist does not claim that the sentence type (~5) or (~6) is intrinsically an expression of generality-absolutism, but only that it expresses generality-absolutism when used in a suitable context. For an analogy, consider these sentences:

(17) It is possible to refer to this place.

(18) I am referring to this place.

Whatever place the context singles out, (17) and (18) are true. In that sense, they are trivially self-verifying sentences. But it would be absurd to argue that therefore, as used in my present context (in which 'this place' happens to refer to Oxford), (17) does not express the claim that it is possible to refer to this place (Oxford), or that (18) does not express the claim that I am referring to this place. The sentences are true precisely by expressing those very truths. Of course, (17) and (18) are not very informative about this place, and (~5) and (~6) are not very informative about everything, but that does not imply that they do not express what they were supposed to express.

Although generality-relativists' intuitive aspiration cannot be fully articulated, they can still say enough to mark sharp explicit disagreements with generality-absolutists: it is not a matter of pure ineffable showing. They need not even generalize over all contexts in order to formulate a disagreement. We saw in section II that the generality-absolutist is quite happy to claim to have been quantifying over absolutely everything in some specified congenial context C. The generality-relativist can shift from C to another context C* and try to demonstrate in C* that not everything was quantified over in C, more things beings quantified over in C* than in C. For example, the generality-relativist might use the presentation of the Russellian paradox of section IV as a way of shifting to C*, claiming that not all interpretations of a predicate letter over the individuals quantified over in C are themselves quantified over in C: some emerge only in C*. In giving that argument, the generality-relativist does not self-defeatingly treat quantification in C* as absolute; another shift of the same sort yields a third context C**, in which the generality-relativist can argue that not everything was quantified over in C*, more things being quantified over in C** than in C*. The argument iterates indefinitely. Of course, the generality-absolutist might say in C* 'I see now that in C I was under an illusion; I was not quantifying over absolutely everything then, but *now* I am'. But that way leads to disaster. The form of the generality-relativist's argument is not specific to the particular context C. The generality-absolutist loses all credibility by responding to the iteration of the argument in C** 'I see now that even in C* I was under an illusion; I was not quantifying over absolutely everything then, but *now* I am', even though the claim is formally coherent.

Generality-absolutists should instead take a stand with respect to a judiciously selected original context C, and argue in C* that everything was already quantified over in C. Indeed, they should have the nerve to stake their position on C, and say in C* 'If I was not quantifying in C over everything, then I give up'. Generality-relativists should likewise have the nerve to stake their position on the shift from C to C*, and say in C* 'If you were quantifying in C over everything, then I give up'. Thus the disputants may agree to treat the shift from C to C* as a *crucial test* to decide the issue between them. If not everything quantified over in C* was already quantified over in C, then the generality-relativist wins the dialectic. If everything quantified over in C* was already quantified over in C, then the generality-absolutist wins the dialectic. That is not to say that the link between the result of the test and the underlying issue is purely logical in either direction. The generality-absolutist could in principle consistently say that the particular context C was a bad choice for accidental reasons and will turn out not to have been one of the absolute contexts after all, because something quantified over in C* was not already quantified over in C. Equally, the generality-relativist could in principle consistently say that the particular contextual shift from C to C* was a bad choice for accidental reasons and will turn out not to have been one of the expansive shifts from C after all, because everything quantified over in C* was already quantified over in C. Nevertheless, if generality-absolutists are willing to treat the context C as their best shot, and generality-relativists are willing to treat the contextual shift from C to C* as their best shot, then the dispute between them can take a definite, explicit and tractable form, even though generality-relativism itself is inexpressible and regards generality-absolutism as equally inexpressible. The distinction between the two sides does not completely evaporate into ineffability. Indeed, if it did so, then the Russellian paradox would no longer motivate generality-relativism, for the generality-absolutist could block the proof of the contradiction in the very same words that the generality-relativist uses, whatever those words are, while meaning them in an ineffably different spirit.[14]

VI

As interpreted in section I, much theorizing in metaphysics depends for its sense on the truth of generality-absolutism. If so, then generality-relativism implies that much of metaphysics is in trouble. But it is not absurd to think that much of metaphysics *is* in trouble. If we rely on generality-absolutism in using words like 'everything' only when pursuing with more or less *hubris* the maximalist ambitions of traditional metaphysics, then the practical costs of generality-relativism might not be very high. Most of our life, even most of our intellectual life, does not appear to depend on the prospects for traditional metaphysics. The generality-absolutist use of 'everything', if so distant from ordinary thought and talk, would radically contrast with more familiar uses of words to express concepts of metaphysical interest: it would not be so easy to give up the 'is' of identity or the 'is' of predication. Of course, the generality-absolutist argues that dependence on generality-absolutism pervades ordinary practice far more extensively than the generality-relativist thinks. Semantically, as argued in section II, the generality-absolutist can analyse restricted quantifiers as constructed out of unrestricted quantifiers, so that the meaningfulness of the former requires the meaningfulness of the latter. Logically, section III showed that if we have absolute generalizations in our language, then they will be inferentially integrated with the more mundane parts of the language. But those points do not show that generality-absolutism is true, and the generality-relativist will regard those ways in which ordinary practice is alleged to depend on generality-absolutism as illusory.

Is it really non-metaphysical business as usual for the generality-relativist? This section begins an argument that it is not.

Consider sentences such as 'Every electron moves at less than the speed of light' and 'No donkey talks'. They too are subject to contextual effects. Someone might utter them with an implicit restriction to particles in a particular cloud-chamber or animals in a particular circus. Intuitively, however, a physicist could utter 'Every electron moves at less than the speed of light' as a generalization (true or false) about all electrons whatsoever, with no electron excluded as irrelevant. Similarly, one could utter 'No donkey talks' as a generalization (true or false) about all donkeys whatsoever, with no donkey excluded as irrelevant. They are uncontroversially intelligible universal generalizations about absolutely every member of a comparatively limited kind. Call them *kind-generalizations*. We expect any language worth its salt to handle kind-generalizations. The generality-absolutist might formalize them as absolutely universal generalizations of conditionals ($\forall x\ (Ex \supset Lx)$, $\forall x\ (Dx \supset \sim Tx)$). But such generality is not obviously needed. To handle them, we need to quantify over absolutely every member of the relevant kind, but that does not appear to involve quantifying over absolutely everything.

For the sake of clarity, let us suppose that the generality-relativist uses a first-order formal language with universal and existential quantifiers. Each quantifier is associated with a limited domain. In order to gain flexibility for the generality-relativist, we allow a sentence to mix quantifiers over different domains. For convenience, we make the relativization to domains explicit with indices. For instance, $\forall[1]x\ \exists[2]y\ Rxy$ says in effect that everything in the first domain *R*s something in the second domain.

The generality-relativist might formalize 'No donkey talks' as $\forall[1]x\ (Dx \supset \sim Tx)$, which is true on the intended interpretation if and only if everything in the first domain does not talk if it is a donkey. If the first domain contained only animals in the circus, that interpretation would be too narrow to capture the intended kind-generalization, for some donkeys would have been excluded from the domain. The first domain must contain absolutely every donkey, but that is a far more modest requirement than that it should include absolutely everything. If the domain contains absolutely every donkey, then no donkey in the domain talks if and only if absolutely no donkey talks, so $\forall[1]x\ (Dx \supset {\sim}Tx)$ is true on its intended interpretation if and only if absolutely no donkey talks. That seems to give the generality-relativist what is needed.

There is a hitch. Suppose that I formalize 'Every creature with a kidney has blood' as $\forall[1]x\ (Kx \supset Bx)$ and stipulate that the domain is to consist of absolutely all and only creatures with a heart. Intuitively, by uttering the formula under that interpretation I do not say that absolutely every creature with a kidney has blood, because the truth of the formula under that interpretation does not exclude the existence outside the domain of a creature with a kidney but no heart or blood. Yet the domain does in fact contain absolutely every creature with a kidney, because it contains absolutely every creature with a heart and as a matter of empirical fact absolutely every creature with a kidney is a creature with a heart (we may assume). Whether the domain contains absolutely every member of the relevant kind is a purely extensional matter; what one says by uttering a generalization over a domain is an intensional matter. To say that absolutely no donkey talks, it is not sufficient to utter the formula $\forall[1]x\ (Dx \supset {\sim}Tx)$ with the intended interpretation of the predicates over a domain that does in fact contain absolutely every donkey. For example, one might stipulate that the domain is to contain all and only terrestrial animals, under the mistaken impression that there are talking donkeys on Mars. By uttering the formula with that interpretation, one would not say that (absolutely) no donkey talks, even though the domain does in fact contain absolutely every donkey, since it contains absolutely every donkey on earth, and in fact absolutely every donkey is on earth (we may assume).

To say that absolutely no donkey talks, speakers need some sort of access to the information that absolutely every donkey is in the domain. To express that information, we need a formula such as $\forall x\ (Dx \supset \exists[1]y\ x{=}y)$ ('Absolutely every donkey is identical with something in the first domain'). But in the current language the outer quantifier needs a domain index too, as in $\forall[2]x\ (Dx \supset \exists[1]y\ x{=}y)$. But to express the information that absolutely every donkey is in the first domain with that formula we must already have access to the information that absolutely every donkey is in the second domain. This looks dangerously like the start of a vicious regress: to have access to one piece of information in the series, we need prior access to a previous piece of information in the series.

Can we solve the problem by simply stipulating that absolutely every donkey is in the domain? We can hardly stipulate it unless we can say it, and that was exactly the original problem: how can generality-relativism allow us to say that absolutely every donkey is such-and-such?

Generality-absolutism faces no corresponding problem of principle, for it treats the presence or absence of contextual restrictions as a feature of the relevant semantic structures. When one omits all such restrictions, one makes an absolutely unrestricted generalization by default.

To see the problem for generality-relativism from another angle, consider the role of determiners in the corresponding sentences of natural languages. In 'Absolutely no donkey talks' and 'Absolutely every donkey brays', the word 'absolutely' modifies the determiners 'no' and 'every', not the noun 'donkey'; 'absolutely no' and 'absolutely every' are complex determiners (compare 'almost no' and 'almost every'). If the relevant use of 'absolutely' cancels contextual restrictions, then it should produce contextually unrestricted readings of the determiners 'no' and 'every'; but how could they be squared with generality-relativism?

The generality-relativist can try again. The kind-generalization that there are no talking donkeys requires the truth of the formula $\forall[i]x\ (Dx \supset \sim Tx)$ on its intended interpretation for each limited domain i. Rather than try to force the claim into a single formula with a fixed limited domain, why not treat it as a *schema* with each formula of the form $\forall[i]x\ (Dx \supset \sim Tx)$ as an instance? To assert the kind-generalization would be to set forth the schema, thereby committing oneself to each of its instances (perhaps some of those instances could be expressed only in other contexts).

The trouble with the suggestion as it stands is that it does not allow for unasserted occurrences of kind-generalizations, for instance those in the scope of negation or the antecedent of a conditional. If one negates the schema $\forall[i]x\ (Dx \supset \sim Tx)$, the result is a new schema, $\sim\forall[i]x\ (Dx \supset \sim Tx)$; its instances are the negations of the instances of the old schema. In setting forth the negated schema one would commit oneself to each of its instances, which is in effect to say of each limited domain that it contains a talking donkey. But someone can deny that there are no talking donkeys without undertaking all those commitments, for the denial is correct if at least one limited domain contains a talking donkey, even if others do not. A domain of dumb animals contains no talking donkey. Negating a schema is like trying to negate a conjunction by taking the conjunction rather than the disjunction of the negations of its conjuncts. If one tries to negate the schema as a whole, rather than negating its instances, one is in effect treating it as a universally quantified sentence of an extended language, something like $\forall i\ \forall[i]x\ (Dx \supset \sim Tx)$, where the variable '$i$' ranges over domains. That formula says in effect that there are absolutely no talking donkeys, for the original schema had the intended generality only if absolutely everything belongs to a domain. Indeed, in the extended language one can express the absolutely universal generalization of any predicate $\varphi(x)$ by $\forall i\ \forall[i]x\ \varphi(x)$. But since generality-relativists want to banish such absolute uses, they will not treat schemas as universally quantified formulas in an extended language. Thus the obstacle to embedding schemas remains.

The notion of a schema can be applied more subtly. First, say that there are no talking donkeys in the first domain:

A[1] $\forall[1]x\,(Dx \supset {\sim}Tx)$

Now use a schema to force the first domain to contain absolutely every donkey, by requiring of each domain *i* that any donkey in it is already in the first domain:

B[*i*] $\forall[i]x\,(Dx \supset \exists[1]y\; x{=}y)$

By simple reasoning from A[1] and B[*i*] we conclude that there are no talking donkeys in domain *i*:

A[*i*] $\forall[i]x\,(Dx \supset {\sim}Tx)$

We might think of schema B[*i*] as a stipulative characterization of the first domain, not a substantive claim up for assertion or denial. Thus we have no motive to embed the schema itself in more complex sentences, for example, by attempting to negate it. In the presence of schema B[*i*], A[1] has the same force as the schema A[*i*]. But the sentence A[1] can be straightforwardly embedded in more complex sentences. Thus A[1] does the work of the claim that there are absolutely no talking donkeys without employing absolute generality.[15]

If this account works for ‘There are no talking donkeys’, it will also work for ‘Every electron moves at less than the speed of light’ and other kind-generalizations. If some domain for the language contains absolutely every member of a kind, we can build that into a schema and use generality over that domain to capture generalizations about every member of the kind. But since by hypothesis no domain contains absolutely everything, it must somehow be illegitimate to make a stipulation of the corresponding schema that would force the first domain to contain absolutely everything ($\forall[i]x\ \exists[1]y\ x=y$).

For the A[*i*]/B[*i*] account to work properly, schemas must be open-ended. That is, their instances must not be limited to sentences currently expressible in the language. In a given context, only comparatively few domains will be specifiable; larger domains are semantically accessible only from other contexts. If A[*i*] had only the currently accessible instances, it would not capture the generality of the claim that there are absolutely no talking donkeys. Rather, commitment to A[*i*] must involve commitment to instances in other contexts; likewise for B[*i*], since any instance of A[*i*] is supposed to follow from the corresponding instance of B[*i*] and A[1]. Each schema is given not by a list of its instances but by a general condition for being an instance. More specifically, we need a general condition for an expression to be of the form $\forall[i]$. What sort of condition will do?

The obvious proposal is that the condition should be a semantic one: an instance of $\forall[i]$ is a universal quantifier over a domain. More formally, given an interpreted language, one might lay down something like this:

SECOND # is an instance of $\forall[i]$ if and only if for some domain D, for every formula α and variable v, $\#v\ \alpha$ is a formula and for every assignment A, $\#v\ \alpha$ is true under A if and only if α is true under A[v/d] for every member d of D.

For simplicity, we temporarily ignore variation in the context of utterance and omit reference to contexts in SECOND.

The semantic condition SECOND quantifies both over domains ('some domain D') and over the members of any domain ('every member d of D'). From the quantifier 'some domain D' and negation we can define the quantifier 'every domain D'. Consider the complex quantifier Q 'for every domain D, every member d of D', as in 'For every domain D, every member d of D is self-identical', interpreted as in the theoretical context CT in which SECOND was set forth. Suppose that in another context CT* we can truly say, quantifying over everything relevant in CT*, 'Q does not range over everything'. Thus some member of the domain of CT* is excluded as a value of the variable 'd' in SECOND. The excluded member is a semantic pariah; it might even be a talking donkey. Consequently, SECOND is inadequate to its purpose, because it does not allow for an instance of $\forall[i]$ that ranges over the domain of CT*. SECOND is adequate only if 'Q ranges over everything' remains true however one enlarges the context. But then Q is equivalent to 'absolutely everything d' as conceived by the generality-absolutist: the conceptual resources of SECOND, together with negation, enable us in effect to express absolute generality. That is not to say that absolute generality is expressible after all in an individual formula of the object-language; SECOND is stated in the meta-language. The point is rather that in order to understand a schema such as $A[i]$ or $B[i]$, one must grasp what distinguishes its instances from its non-instances; if that is a condition like SECOND, then one must grasp a condition that embeds the conceptual resources for absolute generality. The vague word 'grasp' is used deliberately here, for it is notoriously difficult to specify how competent speakers must be related to a semantic theory for their language. What the argument shows is that when speakers express an understanding of the schemas at issue according to a semantic condition like SECOND, either that expression is inadequate or they employ conceptual resources sufficient for absolutely general thought. Thus generality-relativism is in danger of instability under reflection. In explaining resources for it to handle kind-generalizations, we exceeded its limits in the meta-language.

Yet if the argument for generality-relativism in section IV is any good, it applies just as much to the meta-language as it does to the object-language.

We can reinforce the conclusion by considering non-semantic alternatives to SECOND as ways of defining schemas such as B[*i*]. More specifically, we might try to specify the condition for being an instance of the schematic quantifier ∀[*i*] in purely syntactic terms, perhaps by describing an inferential role that an expression will play if and only if it is a universal quantifier. However, we shall see that no such attempt can succeed. For we can construct a pseudo-quantifier that behaves syntactically but not semantically just like a universal quantifier. If it counts as a genuine instance of ∀[*i*], we shall be forced to deny that there are absolutely no talking donkeys.

We start with a standard one-sorted first-order language L with identity; the construction can easily be generalized to many-sorted languages with quantifiers over different domains. For simplicity, we assume that L has no functional expressions. What it is for a formula of L to be true under an interpretation is characterized as usual. We then expand L to a language L^ by adding the expression ∀^, called a *pseudo-quantifier*, countably many expressions called *pseudo-variables* and some expressions called *pseudo-constants*. The pseudo-variables and the pseudo-constants are also called the *pseudo-terms*. The atomic terms of L and the pseudo-terms count as the atomic terms of L^. Well-formedness is defined for L^ in the obvious way, except that quantifiers from L bind only variables from L and the pseudo-quantifier binds only pseudo-variables. We will define what it is for a formula of L^ to be true under an interpretation I of L; the extra vocabulary of L^ does not require extra components of interpretations. As the generality-relativist requires, I is associated with a domain, dom(I). If *t* is a term of L, whether a variable or a constant, I assigns it a denotation I(*t*) in dom(I). A pseudo-term has no denotation under I. To each *n*-place predicate letter *P*, I assigns an extension I(P), a set of n-tuples of members of dom(I) ($n \geq 0$). For atomic formulas, we lay down these rules:

If $t_1, ..., t_n$ are terms of L, then $Pt_1...t_n$ is true under I if and only if $<\mathrm{I}(t_1),...,\mathrm{I}(t_n)>\in\mathrm{I}(P)$.

If some of $t_1, ..., t_n$ are terms of L while others are pseudo-terms, then $Pt_1...t_n$ is not true under I.

If $t_1, ..., t_n$ are pseudo-terms, then $Pt_1...t_n$ is true under I.

The identity sign is interpreted similarly:

If t and u are terms of L, then $t=u$ is true under I if and only if I(t) is identical with I(u).

If t is a term of L and u is a pseudo-term or *vice versa*, then $t=u$ is not true under I.

If t and u are pseudo-terms, then $t=u$ is true under I.

The recursive clauses for the truth-functional operators of L are standard. The universal quantifier $\forall$ of L is interpreted over dom(I) as usual. The clause for the pseudo-quantifier is rather trivial:

If α is a formula and v is a pseudo-variable, then $\forall$^v α is true under I if and only if α

is true under I.

Since the semantic clauses for expressions of L are just the same in L^ as in L, a formula of L is true in L^ under an interpretation I if and only if it is true in L under I. Intuitively, expressions of L have the same meaning in L^ as in L. As for logical consequence in L^, the pseudo-quantifier behaves just like a genuine universal quantifier over a domain with exactly one member. Formally, for any interpretation I of L over a domain and any object o not in dom(I), let I[o] be the two-sorted interpretation of L^ on which both ∀ and ∀^ are genuine universal quantifiers with domains dom(I) and {o} respectively, the variables and terms of L have the same denotations as under I, the pseudo-variables and pseudo-terms are treated as genuine variables and terms denoting o, and the extension of the *n*-place predicate letter *P* under I[o] is $I(P) \cup \{o\}^n$, the result of adding the *n*-tuple of o with itself to the extension of *P* under I. Then one can show by a trivial induction on the complexity of formulas that a formula of L^ is true under I if and only if it is true under I[o]. But interpretations of the form I[o] play no role in the official semantics of L^, because they risk violating the intended interpretation of predicate letters by enlarging their extensions. The official semantics of L^ maps predicate letters to their usual extensions in L. Whatever the pseudo-quantifier means, it is at least an operator with a well-defined truth-conditional semantics.

We can expand L^ to a language L^^ by defining a second pseudo-quantifier ∀^^ that behaves with respect to logical consequence just like a universal quantifier over a domain that results from adding one object to the domain of the quantifiers of L (dom(I)∪{o}). Formally, if α is a formula of L^^ and *v* a variable of L, then ∀^^*v* α is true under I if and only if both ∀*v* α and ∀^*v*^ α^ are true under I, where *v*^ is a pseudo-variable that does not occur in α and α^ is the result of substituting *v*^ for every free occurrence of *v* in α. Thus ∀^^, like ∀^, has a well-defined truth-conditional semantics.

Let the original language L contain the monadic predicates D and T, and I be an interpretation such that dom(I) contains absolutely every donkey and the extensions of D and T under I are the set of donkeys and the set of talkers respectively. If x^ is a pseudo-variable, then the formulas Dx^ and Tx^ are automatically true under I, on the semantics above; thus Dx^ ⊃ ~Tx^ is not true under I. Consequently, ∀^x^ (Dx^ ⊃ ~Tx^) is not true under I. Therefore ∀^^x (Dx ⊃ ~Tx) is not true under I, where 'x' is a variable of L. Syntactically, it looks just as though we have enlarged the original domain to include a talking donkey. Thus, if we demarcate the instances of schema A[i] syntactically, they will include falsehoods. Similarly, the formula ∀^^x (Dx ⊃ ∃y x=y) is not true under I, so if we demarcate the instances of schema B[i] syntactically they too will include falsehoods. But it is absurd to suppose that one can refute the claim that there are absolutely no talking donkeys merely by constructing trivial devices such as ∀^^. The formula ∀^^x (Dx ⊃ ~Tx) is false on the interpretation above whether or not there are talking donkeys. Thus advocates of the schematic approach must explain why that formula is not a genuine instance of A[i]. More specifically, they must explain why ∀^^ is not an instance of ∀[i], a genuine universal quantifier.

The natural move is to revert to a semantic explanation: ∀^^ is not a genuine universal quantifier because it does not satisfy SECOND or some such condition. For no domain D is ∀^^x (Dx ⊃ ~Tx) true under an assignment A if and only if Dx ⊃ ~Tx is true under A[x/d] for every member d of D.[16] For a variable 'x' of L, Dx ⊃ ~Tx is true under A[x/d] for any object d whatsoever, since there are absolutely no talking donkeys; nevertheless, ∀^^x (Dx ⊃ ~Tx) is false under A. In effect, the explanation is that a formula ∀^^x α need not be true under A even if, for every domain D and every member d of D, α is true under A[x/d]. But, as noted above, if the schematic account is to have its intended effect when the instances of the schema are demarcated by SECOND, ‘for every domain D, every member d of D’ must function like an absolutely universal quantifier. Thus the schematic account depends on conceptual resources sufficient for absolute generality in explaining why $\forall$^^ is not a genuine universal quantifier.

Could the generality-relativist treat the explanation as itself schematic? In effect, the supposed explanatory schema would be something like ‘$\forall$^^ does not have the semantics of a universal quantifier over domain D’; different instances would correspond to different substitutions for the schematic variable ‘D’. What the generality-relativist must explain is why $\forall$^^x ($Dx \supset {\sim}Tx$) is not a genuine instance of the schema A[i]. That explanandum is not itself schematic; it is just the negation of the non-schematic claim that $\forall$^^x ($Dx \supset {\sim}Tx$) is a genuine instance of A[i]. But the non-schematic conclusion that the formula is not a genuine instance does not follow from any one instance of the explanatory schema. It follows only from all of them, taken together. Thus what we are dealing with is no mere explanation schema but rather a single explanation with the instances of the explanatory schema gathered together as its premises. To apply the schema like that, we cannot simply engage with its instances individually; we must engage with the schema itself as having generality over all domains. But that is in effect to treat it as a universal generalization over domains. Thus the explanation is implicated in the universal quantification over domains and their members that commits one to absolute generality. Furthermore, that a particular symbol ($\forall$^^) is not a genuine universal quantifier is just the kind of linguistic hypothesis that can be denied as well as asserted, and made the antecedent of a conditional: yet we saw that a purely schematic interpretation does not permit such embedded uses.

To recapitulate: a kind-generalization is true only if it is also true under arbitrary extensions of the domain; the generality-relativist cannot capture that reading in a single formula (such as A[1]) without the help of a schema (such as B[*i*]); the instances of the schema must be demarcated semantically; that demarcation involves conceptual resources that are taboo to the generality-relativist. One can generalize about absolutely every donkey without generalizing about absolutely everything, but only if one does not think too hard about what one is doing. Thought and talk about absolutely everything is not the isolated self-indulgence of metaphysicians; it is the natural outcome of reflection on thought and talk about absolutely every member of a specific kind.

Given how close kind-generalizations are to absolutely universal generalizations, how could the original language express the former without being able to express the latter? In order to express thought about absolutely every donkey, the generality-relativist used schema B[*i*] to convey that absolutely every donkey is in the first domain, because every donkey in an arbitrary domain *i* is already in the first domain. A similar schema conveys that everything in domain *i* is already in the first domain:

$B^{+}[i] \quad \forall[i]x\ \exists[1]y\ x{=}y$

For the generality-relativist, schema $B^{+}[i]$ would somehow be an illegitimate stipulation: its intention of making the first domain contain absolutely everything by stipulating it to contain everything in an arbitrary domain *i* must fail. But if stipulating schema $B^{+}[i]$ fails to make the domain contain absolutely everything, why should stipulating B[*i*] succeed in making it contain absolutely every donkey? Thus the threat of absolute generality by the back door moves down from the meta-language to the object-language.

Generality-relativists have not earned the right to employ an absolutist understanding of ‘every donkey’ while rejecting as illusory an absolutist understanding of ‘everything’. They can use such quantifier phrases to make context-bound generalizations, but they cannot even generalize over other contexts in the schematic way that they would like to say that the generalizations apply to them too. That result is disappointing for theoretical purposes. It does not permit us to understand a sentence such as ‘Absolutely every electron travels at less than the speed of light’ in the way that we should like. We cannot get beyond a contextually restricted claim about every relevant electron because we cannot understand a sentence such as ‘Absolutely every electron is relevant’ in the way that we should like. Scientists may try to formulate a contextually unrestricted claim as the most theoretically illuminating to investigate. For the generality-relativist, the best to be hoped for is that contextual restrictions do not obtrude in practice.

VII

Set the regrets of the last section aside as the last remnants of the naive conception of absolute generality. Instead, let us ask whether generality-relativism is stable even under reflection on the semantics of a language fit simply for the expression of context-bound generality, a language of the sort that one might expect to be innocuous from the perspective of generality-relativism.

Consider a first-order language L in which truth-conditions are context-sensitive and quantification is always over a limited, contextually specified domain. The universal quantifier $\forall$ has the obvious semantics:

[∀C] For every context C, $\forall x\ \alpha$ is true in C under A if and only if every member d of the domain of C is such that α is true in C under A[*x*/d].

As usual, α is a formula of L, '*x*' a variable and A an assignment.

Clause [∀C] employs quantifiers in the meta-language, including 'every member'; are they contextually restricted? The semantic theorist lays down [∀C] in a theoretical context CT. Given the semantic clause [∀C] of the meta-language as used in CT, one can say in a meta-meta-language: the truth-condition for $\forall x\ \alpha$ in C under A is in effect that α is true in C under A[*x*/d] for every member d of the domain of C *over which the meta-language quantifier 'every member' as used in CT ranges.* If CT is associated with a domain of quantification, then [∀C] as used in CT makes the domain of quantification for $\forall x\ \alpha$ as used in C the intersection of the domains of C and CT. Consequently, if 'every member' as used in CT does not range over d, then, for all that [∀C] says, $\forall x\ \alpha$ might be true in C under A even though α is not true in C under A[*x*/d] and d is in the domain of C. But that would violate the intended meaning of $\forall x\ \alpha$, on which it is true in C only if it applies to everything in the domain of C. Thus [∀C] in CT is a suitable semantic clause for L only if, for each context C, in CT the quantifier in the meta-language ranges over every member of C. In terms of domains: the domain of the theoretical context CT must include the domain of every context C; CT must be at least as wide as every context for L. Anything that is relevant in any context is thereby relevant to the semanticist in the theoretical context of an attempt to articulate the cross-contextual truth-conditional semantics of the language. What is relevant to speakers is thereby relevant to semanticists who study them.

To say that the meta-linguistic context CT is at least as wide as every context for L is not yet to say that the meta-linguistic quantifiers as used in CT range over absolutely everything. If something is not in the domain of any context for L, and so cannot be talked about in L at all, then the meta-linguistic quantifier 'every member' as used in CT need not range over it. Such a thing is utterly ineffable in L. If L is a language of limited scope, there may indeed be such things: for example, if L is intended only for the expression of Peano arithmetic, its quantifiers may range only over the natural numbers (whatever they are). But natural languages are not so inflexible. Of course, it is self-defeating to say in English that something cannot be quantified over in English, just as it is self-defeating to say in a context that something cannot be quantified over in that context (see section V). The generality-relativist may reply that not even a natural language can adequately serve as its own meta-language; that is the lesson that Tarski drew from the semantic paradoxes. Let us grant that a truth-predicate for the language L is not translatable into L itself. But that conceptual limitation of the language does not imply any corresponding ontological limitation. According to the generality-relativist, one can truly say, speaking a meta-language for L in an expanded context CT*, that not everything was quantified over in CT. Some object d in the domain of CT* is not in the domain of CT. Thus, from the perspective of CT*, [∀C] gives the correct semantics for ∀ in L only if d is not quantified over in any context for L at all, and thereby constitutes a semantic pariah for L. But it is highly implausible to think that there are such semantic pariahs for English; natural language quantification seems too promiscuous for that. But there is no reason for quantification to be any more restricted in L than it is in English. Thus it is equally implausible to regard d as a semantic pariah for English. The generality-relativist might like to say that in the object-language we can quantify over anything in particular, but not over too much at once. Unfortunately, we have just seen that conception to be inconsistent with generality-relativism. Generality-relativist restrictions on the meta-language make the adequacy of the semantic clause [∀C] imply that there are semantic pariahs for the object-language.

Contemporary model theory faces a similar problem. Speaking in the meta-language of first-order model theory, one says: every model has a set for its domain; since no set contains everything, no model has everything in its domain; but each thing belongs to some sets (such as its own singleton) and therefore to the domain of some model or other. Consequently, no model has every model in its domain. Thus a formalization of the meta-theory in a first-order language has no intended model, in the standard sense. Speaking in a meta-meta-language, the generality-relativist reads disappointing restrictions into generalizations such as 'Each thing belongs to the domain of some model' in a meta-language for first-order logic. The generality-absolutist interprets the meta-language of model theory as expressing absolutely universal generalizations, but sees no reason to insist that quantification in the first-order object-language be restricted to a set domain. Without appeal to set-theoretic considerations, the discussion of [∀C] indicated, concerning the combination of restrictions on how much can be quantified over all at once with the absence of restrictions on what can be quantified over at all, that it characterizes a language only if it does not characterize an adequate meta-language for that language; so the combination is not a universal law —unless an adequate meta-language for the original language is impossible.

Could one reduce the disparity in expressive resources between the object-language L and its meta-language by giving something more like a homophonic semantics for L? One might try to suppress the explicit context-relativity in [∀C] by using a meta-linguistic quantifier in the meta-language subject to the same tacit contextual constraints as the corresponding quantifier of L. The result would look just like the generality-absolutist's semantic clause, but would differ in its intended interpretation:

[$\forall$!] $\forall x\ \alpha$ is true under A if and only if everything d is such that α is true under A[x/d].

'Everything' in [$\forall$!] ranges over whatever contextually restricted domain $\forall$ itself ranges over. Although [$\forall$!] is not literally homophonic, because it renders visible (audible?) the difference between $\forall$ in the object-language and 'everything' in the meta-language, the two quantifiers are intended to work in the same way.

Clause [$\forall$!] may be correct when one is speaking in the theoretical context CT about sentences of L as used in CT itself. Indeed, it may be correct as used in any context to speak about sentences of L as used in that very context. However, [$\forall$!] does not extend properly to cases in which we speak in one context about sentences of L as used in other contexts. For the result of generalizing [$\forall$!] over contexts without explicitly restricting 'everything' to their domains is this:

[$\forall$!+] For every context C, $\forall x\ \alpha$ is true in C under A if and only if everything d is such that α is true in C under A[x/d].

When [$\forall$!+] is used in a theoretical context CT, it interprets $\forall x\ \alpha$ as used in an arbitrary context C as a generalization over the domain of CT. Thus [$\forall$!+] misinterprets $\forall x\ \alpha$ as used in any context with a domain different from the domain of CT. Even if [$\forall$!+] could somehow be read like [$\forall$C], it would thereby lose any supposed advantage over [$\forall$C]. In order to describe contextual variation, serious semantic theorists want to generalize simultaneously about contexts that differ from each other in their domains; at least some of those contexts must therefore also differ in their domains from the context CT. Thus neither [$\forall$!] nor [$\forall$!+] is a theoretically adequate substitute for [$\forall$C].

The falsity of homophonic semantic generalizations for context-dependent languages is a familiar phenomenon. Consider a homophonic clause for the first-person pronoun:

[I!] I am the referent of 'I'.

Clause [I!] may be correct as used in any context about 'I' as used with its current meaning in that very context. However, it does not generalize properly to cases in which we speak in one context about 'I' as used with its current meaning in other contexts. For the result of generalizing [I!] over contexts without making the contextual variation explicit is this:

[I!+] In every context C, I am the referent of 'I'.

But [I!+] is the claim of an egomaniac who regards himself as the referent of 'I' as uttered by anyone else. To formulate a correct generalization about the reference of 'I' as used with its current meaning by other people, one must forsake homophony to lay down a rule like this:

[IC] In every context C, the speaker in C is the referent of 'I' as uttered in C.

Clause [IC] stands to [I!] and [I!+] as [∀C] stands to [∀!] and [∀!+]. To formulate correct semantic generalization across contexts, a meta-language often requires significantly greater conceptual resources than those of an object-language with context-sensitive expressions.

The analogy with [IC] suggests something further: that we should not regard [∀C] as merely a theoretical generalization about L made from a standpoint quite external to that of the speakers of L themselves. Although it would be over-optimistic to expect competent speakers of English explicitly to produce the rule [IC] in immediate response to the question 'What does "I" refer to?', competence with 'I' does require some sort of implicit sensitivity to the implications of the rule. Someone who simply judges the truth-values of utterances involving the word 'I' according to the rule [I!+], on which it refers to him no matter who uses it, rather than according to [IC], would not be competent with the word; he would be treating it as his own name. If he uses 'I' on his own behalf like that, he would not be using it in the normal way. Perhaps competence with 'I' even requires tacit knowledge of [IC]. It is, of course, deeply unclear what such implicit sensitivity or tacit knowledge amounts to. Nevertheless, it is plausible that a speaker who understands 'I' in English is thereby in some sense *committed* to something like [IC]. Similarly, it is plausible that a speaker who understands $\forall$ in L is thereby in the same sense committed to something like [$\forall$C]. More generally, competence with contextually restricted quantifiers requires commitment to contextual variation in their range. If I automatically treat the sentence 'Everything fits into the suitcase' as uttered by any speaker of English as concerning the same items that it concerns when I utter it, I betray a lack of competence with quantification in English. Although competence with L does not enable one to express unrestricted generalizations in L, it nevertheless appears to commit one to unrestrictedly general thought, by committing one to a semantic rule with unrestrictedly general content.

Could the generality-relativist treat the semantic clause [$\forall$C] as itself a schema, with different instances for different ranges of contexts? On that view, however many contexts for L one quantifies over in the meta-language for L, one can always quantify over more. After all, if quantification in L is always over the members of a set, and there is a set of all contexts for L, then there is a set that is the union of the domains of all contexts for L, by the axioms of Zermelo-Fraenkel set theory, in which case some things are semantic pariahs for L, since no set contains absolutely everything. Thus if quantification in L is always over the members of a set, and L has no semantic pariahs, then there is no set of all contexts for L. Consequently, if quantification in the meta-language for L is also always over the members of a set, then one cannot quantify simultaneously over all contexts for L in the meta-language for L, and only a schematic reading of [∀C] is left (but what about the quantification over contexts in this very sentence?).[17] The claim that there is no set of all contexts for L looks implausible if contexts are like spatio-temporal locations, but might be made palatable by appeal to a more abstract notion of context, perhaps one on which the domain is itself a constituent of the context, independent of the other dimensions, for then there will be as many contexts as sets.

One problem for the schematic interpretation of [∀C] is that when one replaces the variable 'L' by an empirical specification of an interpreted language, the result is not a stipulation but an empirical hypothesis. It can be denied as well as asserted; in testing it one considers conditionals of which it is the antecedent. The sentence embeds in more complex sentences in the usual ways. As already noted, schemas do not embed like that. Yet if [∀C] is schematic across ranges of contexts, then so is the result of filling in a value for 'L'. Thus the schematic interpretation does not do justice to the role of [∀C] in empirical semantics. In any case, analogues of the other problems noted in section VI for generality-relativists' use of schemas also apply to a schematic reading of [∀C]. When they come to explain why certain apparent instances of [∀C] are not real instances, a purely syntactic criterion will misclassify some candidates, while an adequate semantic criterion will involve exactly the quantification over absolutely all contexts for L in the meta-meta-language that schemas were invoked to avoid in the meta-language. Thus the use of schemas merely passes the buck higher up the hierarchy of meta-languages. It postpones without solving the generality-relativist's problem.

The upshot of this section is that generality-relativism blocks adequate semantic reflection in a meta-language even on a language for which generality is always contextually restricted, unless the language has semantic pariahs, things over which its quantifiers are semantically prevented from ever ranging. If natural languages lack semantic pariahs, as they seem to, then generality-relativism blocks adequate semantic reflection on natural languages, even in a meta-language, for such reflection would require absolutely general thought. Thus semantics and not just metaphysics requires the notion of absolutely everything. But however dispiriting the implication of generality-relativism that one cannot adequately state the semantics of quantification, it is not quite a *reductio ad absurdum*. Perhaps both metaphysics and semantics need what they cannot have. That conclusion cannot easily be dismissed. Reality may be intrinsically unsystematic or mysterious, essentially resistant to full theoretical understanding. Nevertheless, we may contrapose, and tentatively say: if we can adequately state the semantics of our own language in a suitable meta-language, then generality-absolutism is true.

VIII

Some versions of generality-relativism appeal to the overt syntactic structure of generalizations in natural language rather than to covert contextual domains. Analogues of section VII's arguments apply to them too.

As noted in section II, a generalization such as 'Every donkey brays' decomposes into a complex quantifier ('every donkey') and a verb phrase ('brays'); the quantifier decomposes into a determiner ('every') and a phrase of category N, a head noun ('donkey') optionally qualified by adjectives or relative clauses. The sense in which the complex quantifier DET N is *restricted* by N can be made more precise by the observation that if two verb phrases VP1 and VP2 coincide in application over those things to which N applies, then the sentences (DET N) VP1 and (DET N) VP2 coincide in truth-value.[18]

The (DET N) VP structure may have extra-linguistic significance. Perhaps it evolved to express generalizations over the members of a kind or sort (given by the head noun): sometimes over all members, sometimes over just those members that satisfied extra conditions (given by qualifying adjectives or relative clauses). The kind may be specific (donkeys) or more general (animals), but some generality-relativists think that in a coherent generalization it cannot be so general that it fails to supply any non-trivial principle of individuation whatsoever for the things over which the generalization is to range. It is not entirely clear what a non-trivial principle of individuation is supposed to be, but for the sake of argument let us at least pretend that we know. It is held that, typically, a sortal noun supplies the principle of individuation, although sometimes context is needed too. On this view, we can make sense of 'thing' in 'everything' only by tacitly restricting it to a contextually appropriate sort, perhaps a very wide one, but nevertheless specific enough to provide some non-trivial principle of individuation, and therefore too specific to support an absolutely universal generalization, since it is held that an absolutely universal principle of individuation would be trivial. Kinds or sorts play a similar role to that which sets play in other versions of generality-relativism. On this sortal view, when we try to use words like 'everything' in the way that the generality-absolutist defends, we talk non-obvious nonsense by using natural languages in ways with which they were not designed to cope. When generality-absolutists try to explain how to escape the confines of sortal thinking, their attempts are self-defeating because the only meanings that the semantics of the language can assign to the explanations are sortal ones.

Rather than being left as a specific empirical claim about the human language faculty, the sortal conception of generality may be underpinned by an Aristotelian or quasi-Aristotelian metaphysics of natural and artificial kinds, for generality-relativists as such are not committed to the claim that all forms of metaphysical inquiry depend on an illusion. They can hold that what depends on an illusion is only the use of words like 'everything' as though they transcended all sortal restrictions.[19]

Once again, the challenge to the relativity-generalist is to provide a meta-theory of adequate generality. A natural truth-conditional clause for the determiner 'every' goes something like this:

[Every] 'Every ν x φ' is true under A if and only if every compliant of ν under A, d, is such that 'x φ' is true under A[x/d].

The clause is intended as a generalization over all expressions ν and φ of categories N and VP respectively. For simplicity, reference to context has been omitted in [Every], but the bound variable 'x' has been made explicit in the object-language. A is an assignment. The term 'compliant' is to be understood in terms of the appropriate semantic clause for expressions of category N. The compliants of ν are the things to which ν applies; if ν stands for a kind, then the compliants of ν are the members of that kind.

The clause [Every] uses 'every' in the meta-language in a subsentence of exactly the form (DET N) VP, with 'compliant of ν under A' as the expression of category N. Syntactically, [Every] is faithful to the spirit of sortalism. Semantically, however, the head noun 'compliant' is not associated with any non-trivial principle of individuation in the theoretical context in which [Every] is laid down. It cannot be, for [Every] generalizes across expressions of category N with any head noun available in the language, so ‘compliant’ must be applicable in the context of [Every] to the compliants of all those nouns. If a single non-trivial principle of individuation individuated them all, the supposed obstacle to absolute generality would dissolve, for sortalism is not intended to impose any restriction on what can be a compliant of some noun or other in the language.

Of course, for any particular expression ν of category N and assignment A in an object-language of very limited expressive power, the compliants of the expression ‘compliant of ν under A’ under that assignment may all share a non-trivial principle of individuation, but that principle is not imposed by the noun ‘compliant’. Clause [every] has instances such as these:

[Every-donkey] ‘Every donkey *x* φ’ is true under A if and only if every donkey d is such that ‘*x* φ’ is true under A[*x*/d].

[Every-stick] ‘Every stick *x* φ’ is true under A if and only if every stick d is such that ‘*x* φ’ is true under A[*x*/d].

But of course we do not have to learn all expressions of the form DET N separately. You can understand ‘every’ and ‘stick’ separately and then understand ‘every stick’ on first hearing. The meaning of determiners such as ‘every’, ‘some’ and ‘no’ is neutral between different sortals. To capture this compositionality in our understanding of complex quantifiers such as ‘every stick’, we need a clause as general as [Every], which is what makes trouble for the generality-relativist. Using the resources of [Every], we can construct complex cross-sortal meta-linguistic generalizations such as ‘Every expression ν is such that every compliant of ν φ’.

Once we have seen that the meta-language violates the supposed sortal constraint on quantification, we lose any motivation to expect the object-language to obey the constraint. For the argument for the sortal constraint was quite general; if it had been any good, it would have applied to the meta-language just as to the object-language. Once we deny the sortal constraint, we can acknowledge the obvious fact that natural languages contain many nouns that are not associated with any non-trivial principle of individuation, even as they are coherently used in a particular context (‘thing’, ‘object’, ‘item’, ‘entity’, ‘member’, ‘element’, ‘instance’, ‘example’, ‘topic’, ‘compliant’, ...). Similarly, the natural semantics for the prefix ‘non-‘, as applied to expressions of category N, is that for any such expression ν, ‘non-ν’ is another expression of category N that applies to all and only those things to which ν does not apply. Whatever is not an animal is a non-animal. On the sortalist view, non-animals have no non-trivial principle of individuation in common. Yet ‘non-animal’ is a legitimate expression of category N, and the natural understanding of ‘non-‘ makes the conjunction ‘Every animal φ and every non-animal φ’ logically equivalent to the absolutely universal generalization ‘Everything φ’. Of course, the generality-relativist may claim that coherent uses of such expressions are always non-trivially restricted by context, but that is just to retreat to the contextualist version of the theory that encountered the problems explained in sections V-VII; it is not justified by a sortal feature of the grammars of natural languages. The foregoing arguments extend in obvious ways to languages that combine contextual limits on quantifier domains with sortal limits on determiner complements. We can describe their semantics adequately only in a meta-language that transcends both sorts of limit.

IX

Generality-relativism endangers more than metaphysics. As sections V-VIII have shown, it endangers the possibility of a reflective understanding of our own thought and language, even from the standpoint of a meta-language. Of course, we never expected to ascend all the way up through an infinite hierarchy of meta-languages, but generality-relativism is an obstacle to an adequate meta-language even at the very first level. If the unexamined life is not worth living, the credentials of a life without absolutely general thought are shaky. Yet, according to generality-relativism, the attempt to think with absolute generality inevitably leads us into paradoxes. How can we proceed with integrity in thinking about our language on the assumption that such thinking is doomed to paradox? (Is the postmodern answer to proceed without integrity?)

Generality-relativism leads to meta-linguistic pessimism: but meta-linguistic pessimism undermines the original case for generality-relativism. For the case depended on the charge that generality-absolutism inevitably generates a version of Russell's paradox. The representative version in section IV involved quantifying over generality-absolutist interpretations of a language. But, according to meta-linguistically pessimistic generality-relativists, such meta-linguistic reflection meets insuperable obstacles even when combined with a correct view of generality. They are therefore in a weak position to blame the crucial version of Russell's paradox on generality-absolutism rather than on its use of meta-linguistic reflection. If the generality-relativist is allowed to block paradoxes with meta-linguistic pessimism, why not the generality-absolutist too?

Fortunately, the generality-absolutist can respond to the Russellian paradox without resort to meta-linguistic pessimism. The paradox-inducing argument (1)-(4) of section IV assumes that there are such things as interpretations, in particular the interpretations of predicate letters. That assumption was justified by the need to generalize over such interpretations in a Tarskian definition of logical consequence. Suppose, for example, that we are interested in whether $\exists x\ Px$ is a logical consequence of $\forall x\ Px$. We might stipulate that the predicate letter P applies$_I$ to something if and only if it brays, and that it applies$_J$ to something if and only if it is red. Correspondingly, we have conditionals like these:

(19_I) If $\forall x\ Px$ is true$_I$ then $\exists x\ Px$ is true$_I$.

(19_J) If $\forall x\ Px$ is true$_J$ then $\exists x\ Px$ is true$_J$.

In effect, Tarski requires us to generalize into the subscript position of 'I' and 'J' in (19_I) and (19_J). Of course, in defining logical consequence for a reasonably expressive language our concern is not just with a single predicate, as above, but since the crucial difficulty arises even in that simple case, we can ignore the rest of the language for present purposes. We naturally ask ourselves the question: if we are to generalize, *what* are we to generalize over? The natural answer is 'interpretations'. If there are interpretations, then they are automatically things in the generality-absolutist's sense.

In asking ourselves what we are to generalize over, we took for granted that a subscript like 'I' on which we are to generalize is somehow to be treated as referring to something of some kind. Our task is then to generalize over everything of that kind. Semantically, we assimilate the position of a subscript to the position of a name or other singular term. The underlying assumption is that generalizing always amounts to generalizing into name position, that all quantification in the end reduces to first-order quantification. But that Quinean assumption is not forced on us. After all, the difference between the subscripts 'I' and 'J' was to mark the difference between the interpreting predicates 'brays' and 'is red'. It is therefore more natural, and more in the spirit of homophonic semantics, to think of subscript position in (19_I) and (19_J) as predicate position rather than name position. We might say that the predicate letter *P* *I*s an object o if and only if o brays, and that *P* *J*s o if and only if o is red. In defining logical consequence, we generalize into predicate position in a second- (or higher-)order meta-language. We reject the question 'What are we to generalize over?' because inserting a predicate in the blank in 'We are to generalize over ...' produces an ill-formed string.

More formally, consider a first-order language L with identity. For simplicity, assume that L has no function symbols, and that its primitive logical symbols are just ∀, ~ and &. Consider a dyadic meta-linguistic predicate 'D' (written between its arguments). Suppose that:

DEN(D) Each singular term *t* of L Ds exactly one thing.

Given DEN(D), let $\mathrm{den}_D(t)$ be the thing that *t* Ds. Informally, we can read 'D' as 'denote'; *n*-place atomic predicates will also be treated as denoting the *n*-tuples to which they apply (see [*P*2] below). For any variable *v* of L and anything d, say that a thing o D[*v*/d]s a thing o* if and only if either o is not *v* and o Ds o* or o is *v* and o* is d; on suitable assignments to '*v*' and 'd', 'D[*v*/d]' is another dyadic meta-linguistic predicate. DEN(D) entails DEN(D[*v*/d]): each singular term *t* of L D[*v*/d]s exactly one thing. More specifically, if *t* is not *v*, then *t* Ds just what it D[*v*/d]s, but *v* D[*v*/d]s d rather than $\mathrm{den}_D(v)$. Now define a new monadic meta-linguistic predicate 'true-D' recursively on the formulas of L (likewise for any other dyadic meta-linguistic predicate ‘D*’ in place of ‘D’, given DEN(D*)):

[P2] For any n-place atomic predicate P and singular terms $t_1, ..., t_n$ of L, $Pt_1,...,t_n$ true-Ds if and only if P Ds <$\text{den}_D(t_1), ..., \text{den}_D(t_n)$>.

[=2] For any singular terms t_1 and t_2 of L, $t_1=t_2$ true-Ds if and only if $\text{den}_D(t_1)$ is identical with $\text{den}_D(t_2)$.

[~2] For any formula α of L, ~α true-Ds if and only if α does not true-D.

[&2] For any formulas α and β of L, α & β true-Ds if and only if α true-Ds and β true-Ds.

[∀2] For any formula α and variable v of L, ∀v α true-Ds if and only if for everything d, α true-D[v/d]s.

‘Everything’ in clause [∀2] is to be read in the generality-absolutist way. The five clauses define a truth predicate ‘true-D’ for L from the denotation predicate ‘D’. It has not been assumed that a denotation predicate stands for a denotation relation or anything else. No such things as interpretations have been postulated.[21]

What happens to the Russellian paradox of section IV within this framework? Given a predicate ‘F’, in place of the singular term ‘I(F)’ for an interpretation in the paradoxical argument we have only a denotation predicate ‘D(F)’. Thus premise (1) of the paradoxical argument becomes this:

(1[2])   For everything o, *P* D(F)s o if and only if o Fs.

The expressions ‘F’ and ‘D(F)’ occupy predicate position in (1[2]). If we try to define a verb ‘R’ in imitation of step (2) of the argument, we must first make a predicate ‘P(o)’ out of the individual variable ‘o’ in order to have a well-formed result when we put it in predicate position:

(2[2])   For everything o, o Rs if and only if *P* does not D(P(o)) o.

As before, we can put ‘R’ for ‘F’ in the first premise and apply the definition of ‘R’:

(3[2])   For everything o, *P* D(R)s o if and only if *P* does not D(P(o)) o.

If we could find a value of ‘o’ for which ‘D(P(o))’ was equivalent to ‘D(R)’ we should still have a paradox: but we no longer have any reason to think that there is such a value of ‘o’. Thus the paradox dissolves.

Generality-absolutists may fall back into paradox if they somehow commit themselves to a reduction of quantification into predicate position to quantification into name position, by giving the semantics of the second-order meta-language in a first-order meta-meta-language or otherwise. But if they stick resolutely to the higher-order viewpoint throughout the hierarchy of meta-languages, they can avoid paradox. In particular, they should reject Quine’s insistence on regimenting a theory into a first-order language as the test of its ontological commitments.

Few philosophers of logic still feel bound by Quine’s ban on any but first-order quantification. Nevertheless, a tendency lingers to feel uncomfortable with any but first-order quantification in the meta-language in which one explains what the truth-conditions of sentences of the object-language really are. The generality-absolutist should regard that reversion to the familiar first-order viewpoint in the meta-language as a disastrous failure of nerve. Marshall Bazaine helped lose the Franco-Prussian war by retreating to the apparent safety of the fortress of Metz; he was compared to a sailor clinging to the anchor in a shipwreck. The appearance of security in first-order explanations is equally illusory.

It is often assumed that the only alternative to a first-order 'objectual' reading of the quantifiers is a substitutional one. On a substitutional reading, the second-order generalization $\forall F\ \alpha$ is true if and only if every uniform substitution of a predicate of the language for the variable 'F' in α yields a truth. Substitutional quantification is a fairly harmless device: but it is not what is wanted here. The paradox arose when we tried to apply Tarski's definition of logical consequence to a first-order object language without domain restrictions. As Tarski emphasized ([CTC]), a definition of logical consequence as truth-preservation over all substitution instances creates anomalous effects by making the validity of an argument irrelevantly sensitive to accidental expressive limitations elsewhere in the language. It is not sufficient for the truth of $\forall F\ \alpha$, interpreted suitably for the definition of logical consequence, that every uniform substitution of a predicate of the language for 'F' yield a truth. An objectual reading interprets the quantifiers as ranging over objects and merely revives the Russellian paradox. The generality-absolutist requires higher-order quantification that is neither objectual nor substitutional.

The generality-relativist will press the question: how are we to read second-order generalizations such as $\forall F\ \alpha$? If they are just uninterpreted formulas, we cannot use them in a serious definition of logical consequence. Yet natural languages seem to force us into reading all generalizations as generalizations over things, whatever we are generalizing over, with which the generality-absolutist falls back into a version of Russell's paradox. We can always ask 'What are you quantifying over?', 'What are the values of the variables?' The form of the questions demands a noun phrase in response.

George Boolos gave an ingenious and much-discussed response by interpreting second-order logic in terms of quantification into *plural* noun position.[22] For example, we might read the formula $\forall F\ (\forall x\ Fx \supset \exists x\ Fx)$ thus:

(20) Any things are such that if everything is one of them then something is one of them.

The initial quantifier 'any things' in (20) is not to be paraphrased as 'any set of things', 'any plurality of things' or with any other singular expression of category N as the complement of 'any'. Generality-absolutists can indeed formulate a definition of logical consequence for a first-order object-language in a second-order meta-language, plurally interpreted, that picks out just the arguments that they want. Boolos may very well be right that plural quantification does not collapse into any kind of singular quantification. Nevertheless, there are reasons to doubt that plurals are ultimately what we want for generalizing into predicate position.

The most obvious point is that plurals are not predicative (Simons [HOQOC]): the expressions 'red things' and 'is red' belong to grammatically distinct categories, plural N and VP respectively. To read atomic Px as 'it is one of them' is to impose more structure than appears to be present in the object-language. This difference is too fine-grained to matter for most purposes. In particular, it does not make the definition of logical consequence misclassify any arguments in a first-order object-language. Nevertheless, it hints that the plural approach is unnecessarily indirect.

The second point is a small one. Since ‘no things’ does not express a special case of ‘any things’ or ‘some things’, Boolos is forced to add an extra clause to the plural reading to handle the special case in which the predicate applies to nothing, a case whose legitimacy Frege emphasized. For example, to (20) we might add this extra conjunct:

(21) If everything is self-distinct then something is self-distinct.

The extra clauses make no technical trouble, but the need for them gives another hint of Procrustean activity. Analogy: although a philosopher of mathematics who denies that there is a natural number zero may have no technical trouble in adding clauses to handle the special case of apparent reference to zero, the need for those extra clauses is evidence that something has already gone wrong, that the theory is out of tune with the underlying reality.

A third point, also small, is that quantification into *n*-place predicate position has no natural plural reading for *n* greater than one, although it can be regimented as plural quantification over *n*-tuples.[23] Again, the unnaturalness does no more than hint that something may be wrong.

The fourth point is that the plural reading gives unwanted results when modal operators are introduced into second-order logic. For example, the following sentence is false, read as a first-order universal generalization:

(22) If anything could have been wet then it is wet.

My computer could have been wet, but it is not wet. Thus the formula $\forall x\ (\Diamond Fx \supset Fx)$ has a false reading, so the second-order generalization $\forall F\ \forall x\ (\Diamond Fx \supset Fx)$ is false. But the second-order generalization, read plurally, says something like this (with the extra conjunct for the empty case):

(23) Any things are such that if anything could have been one of them then it is one of them, and if anything could have been self-distinct then it is self-distinct.

But (23) seems to be true, since the natural reading of plural variables is a 'rigid' one: for any thing and any things, it is not contingent whether the former is one of the latter. No doubt plural definite descriptions such as 'the wet things' can easily be non-rigid, but that no more shows that the corresponding plural variables are non-rigid than the non-rigidity of singular definite descriptions such as 'the wettest thing in my room' shows that the corresponding singular variables are rigid. Indeed, the rigidity of plurals makes the plural reading attractively extensional. If every creature with a heart is a creature with a kidney and *vice versa*, then the creatures with a heart *are* the creatures with a kidney, even though there could have been a creature with a heart and no kidney. For a modal language, it is hard to make sense of the plural reading of the second-order quantifiers without validating this principle, the analogue of the necessity of identity for plurals:

(24) $\forall F\ \forall G\ (\forall x\ (Fx \equiv Gx) \supset \Box\ \forall x\ (Fx \equiv Gx))$

But if we want to be able to instantiate universally quantified second-order variables with predicates such as 'is a creature with a heart' and 'is a creature with a kidney', then we must reject (24) and seek a more intensional, predicative reading of second-order quantification than plurals seem to provide.

A fifth point is that the plural reading has no natural generalization to *n*th-order quantification for *n* greater than two. But, by a generalization of the Russellian paradox, we need to use (*n*+1)th-order quantification in the meta-language to define logical consequence for an *n*th-order object-language that contains *n*th-level non-logical predicates. For example, the formula Π(F) in which the second-level predicate Π is applied to the first-level predicate 'F' might have the reading 'the things that F collectively lifted a piano'. A Tarskian definition of logical consequence for a second-order language with the second-level predicate Π in the spirit of the account above would involve third-order quantification into the position of Π, for which the plural reading would not suffice.

The five points above in no way undermine the intelligibility of a plural reading of a second-order formal language. Such a reading is good enough for the purposes of an extensionally adequate account of truth and logical consequence for a first-order object-language with unrestricted quantification. That provides a sort of consistency proof of the generality-absolutist account of first-order languages, relative to the logic of plural quantifiers in natural languages. Since the paradoxes, the main threats to the consistency of generality-absolutism, already arises for the generality-absolutist account of first-order languages, we therefore have some evidence for the overall consistency of generality-absolutism. Moreover, since the most serious threats to the truth of generality-absolutism are threats to its consistency, defending its consistency goes quite some way towards defending its truth.[24]

Nevertheless, the five points above indicate that generality-absolutists should not rest content with the plural reading of second-order variables. They suggest that a more predicative reading may be more appropriate for present purposes. Perhaps we can go some way in that direction by thinking in Fregean terms.

A crude Fregean account of the difference between quantification into name position and quantification into predicate position is that it is the difference between quantification over objects and quantification over concepts. But that formulation immediately suggests an objection to generality-absolutism. If quantification into name position ranges over objects but not over concepts, then it is not absolutely universal after all, for concepts are things in generality-absolutists' unrestricted sense of 'thing'. If generality-absolutists respond by denying that concepts are things, they cheat by betraying the spirit of their original explanation. Rather, they should insist that the very use of the noun 'concept' makes the crude Fregean account inaccurate. What is really needed is a higher-level predicate that stands to first-level predicates as the first-level predicate 'is an object' stands to names. [25] For the same reason, the attempt to contrast objects and concepts as saturated and unsaturated respectively is deeply misleading, for 'unsaturated' is the negation of 'saturated' and the two adjectives belong to the same grammatical category; but whereas 'is saturated' is a first-level predicate, we need a higher-level predicate in place of 'is unsaturated' to do the required work. The distinction must remain one of grammar and not of ontology, because one cannot use first-level and second-level expressions in the same grammatical context to articulate an ontological distinction without violating constraints of well-formedness. Frege himself had some conception of the difficulty, as his uneasy discussion of the concept *horse* shows. He knew that his explanations required at least a pinch of salt, but vastly underestimated the depth of the problem. He fell into inconsistency by postulating Basic Law V in an attempt to avoid the mathematically inconvenient consequences of the distinction between first-level and second-level expressions.[26] It is quite within the spirit of Frege's philosophy to insist that one can state matters perspicuously only in a formal language such as his *begriffsschrift*. In that notation (without Basic Law V), quantification into predicate position is simply incommensurable with quantification into name position; the former presents no coherent threat to the absolute generality of the latter. Things in the generality-absolutist's sense are just Fregean objects; there is nothing else. Frege interpreted his quantifiers into name position as ranging over all objects whatsoever, not just over those in some limited domain; he was a generality-absolutist in spirit.

On this view of second-order quantification, we must reject as misconceived the questions 'What does quantification into predicate position quantify over?' and 'What are the values of variables in predicate position?'; in particular, we must not answer 'Concepts'. If second-level analogues of those questions make sense, the answer to them will be the second-level analogue of 'Objects' as an answer to the questions 'What does quantification into name position quantify over?' and 'What are the values of variables in name position?'

Perhaps no reading in a natural language of quantification into predicate position is wholly satisfactory. If so, that does not show that something is wrong with quantification into predicate position, for it may reflect an expressive inadequacy in natural languages. We may have to learn second-order languages by the direct method, not by translating them into a language with which we are already familiar. After all, that may well be how we come to understand other symbols in contemporary logic, such as $\supset$ and $\Diamond$: we can approximate them by 'if' and 'possibly', but for familiar reasons they may fall short of perfect synonymy, and we certainly do not employ $\supset$ and $\Diamond$ as synonyms for the complex discourses in which we explain how they differ subtly in meaning from 'if' and 'possibly'. At some point, we learn to understand the symbols directly; why not use the same method for $\forall F$? We must learn to use higher-order languages as our home language. Having done so, we can do the semantics and metalogic of a higher-order formal language in a higher-order formal meta-language of even greater expressive power.

Issues in philosophy often turn on what language we use as our home language, the language in which we are happy to work, at least for the time being, without seeing it through the lens of a meta-language, the language that we treat as basic for explanatory purposes. For example, questions about tense or modality look one way if we allow unreduced tense or modal operators in our home language; they look quite different if we insist on reducing them to other terms in a first-order language with quantification over times or worlds. What we are willing to take as our home language is partly a matter of what we feel comfortable with; unfortunately, it can be hard to argue someone into feeling comfortable. Generality-absolutists cannot expect a quick victory, for it takes time to become at home with semantic theories for absolutely universal first-order quantification in a higher-order meta-language. If we try to understand such theories by translation into other terms, the generality-absolutists may well appear to be cheating, by restricting the first-order quantifiers which they advertise as absolutely universal to exclude some things over which the higher-order quantifiers range. For generality-absolutists, that appearance is a mere artefact of bad translation.

Can such a stand-off ever be resolved? We saw in earlier sections how badly generality-relativists fare when they engage in semantic reflection: not so badly that we can be immediately sure that they are wrong without even checking how their opponents are doing, but still very badly. If generality-absolutists can engage in consistent and systematic semantic theory-building within a higher-order meta-language, and give principled replies from that perspective to generality-relativists' charges of cheating, that will be good evidence that the path of enlightenment leads in their direction. They have already begun the task of theory-building. If the best cases that they can offer of absolutely unrestricted generalizations turn out by their own criteria not to be absolutely unrestricted, then the generality-absolutist programme is bankrupt. But if the work of theory-building continues with no such disaster, then the appeal of generality-relativism will gradually become just obscurantist.[27]

NOTES

1 A more heroic line for the generality-absolutist to take is that the word 'everything' semantically always means absolutely everything, but pragmatically is sometimes used to convey contextually restricted generality (Bach [CI]: 138-9 and [QQC]; for criticism see Stanley and Szabó [QDR]: 239-45). A successful defence of that heroic line would facilitate the argument for generality-absolutism in this paper but is not required by it.

2 For a recent account of conditions that justify interpreting 'everything' in the generality-absolutist way see Rayo [WEME]; Glanzberg [QR] takes the opposing view. McGee [E] points out that the universal quantifier is uniquely characterized by its introduction and elimination rules in first-order logic, in the sense that two quantifiers both subject to those rules are provably equivalent (see also Harris [LLA] and Williamson [EE]). As McGee notes, what is thereby uniquely characterized is absolutely universal quantification only if no contextual restriction is in force on what can be named. The possibility of unrestricted naming raises similar issues to the possibility of unrestricted quantification. Something may be so elusive that humans cannot single it out for naming; nevertheless, unrestricted quantifiers range over it, and no restriction in the grammar of English *forbids* us to name it.

3 See Cartwright [SE] and McGee [TPTTC] for further discussion.

4 See Stanley and Szabó [QDR]: 249 and Stanley [NR] for bound contextual restrictions; Bach [QQC] defends a contrary interpretation. A less acute difficulty for the treatment of contextual restrictions as pre-sentential discourse features is that they can vary between quantifiers in the same sentence (Westerståhl [DCS], Soames [IDD], Stanley and Williamson [QCD]); such variation is easy to understand if the restrictions are constituents of semantic structure. If one conceives the restrictions as bare sets, they may seem unsuitable to act as such constituents; but it is independently plausible that contextual restrictions are intensional in a sense suitable for semantic constituents (see also Stanley and Szabó [QDR]: 252).

5 Contextual restrictions in natural language need not be restrictions to the members of a set. On one reading of 'Everything except the empty set has members', the initial quantifier ranges over all and only sets, even though there is no set of all sets.

6 Barwise and Cooper [GQNL] is a seminal discussion of the (DET N) VP structure of natural language quantification.

7 If contextual restrictions are constituents of semantic structure, the question arises whether they apply to the determiner ('every') or to the head noun ('thing'). Stanley [NR] argues for the latter view. However, it leads to complications; for example, if the sentence 'Every student is talking loudly' is true as uttered with a restriction to the people at a party, then the restriction must be applied to 'student'; but it must be applied a second time to 'non-student' as a whole in 'Every non-student is talking loudly', otherwise people who are not at the party will incorrectly count as falling

under ‘non-student’, since they do not fall under ‘student’, and therefore as falsifying ‘Every non-student is talking loudly’. Worse, in a context in which only cheap goods are relevant (say, those costing less than £10), the noun ‘antique’ would be restricted to cheap antiques, so ‘fake antique’ would presumably apply only to fake cheap antiques, which is the wrong result, since it should also apply to cheap fake expensive antiques. If the contextual restriction is applied to ‘antique’ in ‘Some antiques are enormous’ but only to ‘fake antique’ and not to ‘antique’ in ‘Some fake antiques are enormous’, the account appears to be non-compositional. These complications disappear if the contextual restriction is attached to the determiner rather than the head noun. Stanley criticizes that view by citing examples such as ‘The tallest person is nice’, for if ‘tallest person’ is contextually unrestricted, it applies only to the tallest person in the universe, and may therefore apply to nothing contextually relevant at all; but since the phrase ‘tallest person’ is implicitly quantified (‘person taller than every other person’), such examples are compatible with the view at issue. However, generality-absolutism can afford to be neutral on the detailed mechanisms of contextual restriction in natural languages.

8 The discussion in the text is amplified in Williamson [BP], [EC] and [NFO], which provide a defence of the Barcan formula and its converse. The denial of the Barcan formula becomes especially hard to maintain on an unrestricted reading of the quantifiers if the semantics for the modal object-language is given in a meta-language without modal operators, as with possible worlds semantics, for then one must be able to quantify in the meta-language over the members of the domain of any world; analogous points apply to the semantics of fiction and tense. For an example of the

use of unrestricted quantification to set up an issue about the metaphysics of modality see Lewis [PW]: 136-8.

9 It is natural for the generality-absolutist to treat $\forall$ (obligatorily interpreted as unrestricted) as a logical constant (contrast Cartwright [SE]: 9-10), in which case, on standard assumptions, for each natural number *n* the formalization of the claim that there are at least *n* things will be a logical truth (a result with which logicists will be happy). If those formulas are added as axioms to a standard formalization of first-order logic, the result can be shown to be sound and complete for the unrestricted reading, given a global choice principle (Friedman [CTE], Williamson [EC], Rayo and Williamson [CTUFL], on the last of which McGee [UUQ] is a commentary). By permitting both restricted and unrestricted interpretations of $\forall$, Cartwright secures the soundness and completeness of first-order logic without the extra axioms ([SE]: 9-10).

10 Much of Cartwright [SE] is an effective critique of Michael Dummett's attempts to use the set-theoretic paradoxes to discredit unrestricted quantification, and a more sympathetic account of Russell's treatment of the same issue. It is unnecessary to go over the same ground here. See also Boolos [RPSC].

11 Forster [STUS] is an account of set theories with a universal set.

12 McGee [TPTCC] and Williamson [EC] discuss Tarski's account of logical consequence in [CLC] with respect to absolute generality. Tarski does not mention

domains of quantification when he explains what a model is; he speaks only of assignments of values to variables. He may well have been assuming that the first-order quantifiers range over all individuals whatsoever.

13 It is not claimed that the paradox in the text is the only variant on the set-theoretic and semantic paradoxes to constitute a threat to generality-absolutism. However, the assumptions on which it depends are particularly economical; if generality-absolutism can withstand this threat without *ad hoc* concessions, then it is in good shape to withstand all such threats. For paradoxes of quantification that depend on the notion of a proposition and of a proposition's being about something, see Grim [IU]: 113-22.

14 Generality-absolutists will also reject some generality-relativist meta-linguistic claims. For example, they will reject (15) on the grounds that the English sentence 'Not everything is quantified over in C0' is false as uttered in any context when 'C0' refers to a context in which 'everything' expresses unrestricted generality.

15 Compare the set-theoretic notion of absoluteness. Roughly, if an absolute formula is true when its quantifiers are restricted to a transitive class domain D (that is, any member of a member of D is a member of D), then it is true when its quantifiers are restricted to any extension D* of D, under a fixed assignment of values in D to variables. For example, the formula $\forall y \sim y \in x$ is absolute because any member in D* of the set assigned to the variable x would already have been a member of D (since D is transitive).

16 For a fixed domain and interpretation of the constant terms and predicate letters, interpretations are equivalent to assignments of values to variables.

17 Kaplan's definition of validity for arguments in a language with demonstratives employs quantification over every context ([D]: 547), and therefore carries an unwanted contextual restriction on the proposed view.

18 In Barwise and Cooper's terminology, the operator expressed by DET N *lives on* the extension of N ([GQNL]: 178). In effect, their definition requires only the special case in which the extension of VP2 is the intersection of the extension of VP1 with the extension of N, but that special case is easily seen to be equivalent in any set model to the more general case in the text. They propose the linguistic universal that '*Every natural language contains basic expressions,* (called determiners) *whose semantic function is to assign to common count noun denotations* (i.e., sets) *A a quantifier that lives on A*' ([GQNL]: 179). This property of determiners has come to be known as conservativeness. The generality-absolutist should insist on expressing it without set-theoretic apparatus.

19 The view in the text is reminiscent of Geach's in [RG]. However, the present concern is not with the details of his view.

20 To be tidy, we could additionally require that nothing but a singular term or atomic predicate of L Ds anything, but doing so would make no difference to the semantic evaluation of expressions of L.

21 For semantic theories in a second-order meta-language see Boolos [NP], Rayo and Uzquiano [TTSC], Rayo and Williamson [CTUFL] and McGee [UUQ]. Quine [PL]: 66-68 was an influential critique of second-order logic, to which Boolos [LLL] and Shapiro [FF] provide sustained replies. For further debate see Shapiro [LL].

22 See Boolos [BBVV] and [NP]. For somr recent discussion see Higginbotham [HOLNL], to which Bostock [MHOL] replies, Rayo [WO] and Yi [LLP].

23 For discussion see the appendix by John P. Burgess, A.P. Hazen and Lewis to Lewis [PC] and Rayo and Yablo [NDN]: 75-77.

24 For a different attempt to render second-order quantification into a form of English see Rayo and Yablo [NDN].

25 According to Dummett's memory of Frege's Nachlass ([FPL]: 212-13), Frege later took the line proposed for him in the text. Wright [FGS] objects that the move does not resolve the problem because a meta-linguistic expression such as 'the referent of "is a horse"' still does occupy name position. However, the point of using a second-order meta-language is to avoid such expressions, because they regenerate Russell's paradox.

26 Dummett [FPM]: 217 blames the paradox on Frege's use of second-order quantification; Boolos [WC] replies.

27 Earlier versions of this material were presented at a British Logic Colloquium meeting in Manchester, the Fourth European Congress of Analytic Philosophy in Lund, Metaphysical Mayhem VII in Syracuse, a conference of the Italian Society for Analytic Philosophy in Bergamo, a meeting of the Eastern Division of the American Philosophical Association in Philadelphia and a conference on the Philosophy of Logic and Language in Birmingham, at the Institut d'Histoire et Philosophie des Sciences et des Techniques in Paris, and at the following universities: Lisbon, Liverpool, Michigan, Minnesota, Missouri, Nottingham, Oxford, Queen's (Kingston), Rutgers, St Andrews and Toronto. I thank the audiences at those events for their many helpful comments, and in particular Michael Glanzberg, Peter Pagin and Stewart Shapiro for their replies at Syracuse, Birmingham and Philadelphia respectively, Peter Lopston for written comments, and Kit Fine, Agustin Rayo and Byeon-Uk Yi for extended discussion of the issues.

BIBLIOGRAPHY

Bach, K. [CI] ‘Conversational impliciture’, *Mind and Language* **9** (1994): 124-162.

Bach, K. [QQC] ‘Quantification, qualification and context: a reply to Stanley and Szabó’, *Mind and Language* **15** (2000): 262-283.

Barwise, J., and Cooper, R. [GQNL] ‘Generalized quantifiers and natural language’, *Linguistics and Philosophy* **4** (1981): 159-219.

Boolos, G. [RPSC] ‘A reply to Parsons’ “Sets and Classes”’, in Boolos [LLL].

Boolos, G. [BBVV] ‘To be is to be a value of a variable (or to be some values of some variables’, *The Journal of Philosophy* **81** (1984): 430-449.

Boolos, G. [NP] ‘Nominalist platonism’, *The Philosophical Review* **94** (1985): 327-344.

Boolos, G. [WC] ‘Whence the contradiction?’, *Aristotelian Society* supplementary volume **67** (1993): 213-233.

Boolos, G. [LLL] *Logic, Logic and Logic*. Harvard, Cambridge MA: 1998.

Bostock, D. [MHOL] ‘On motivating higher-order logic’, in Smiley [PL].

Bradley, F.H. [AR] *Appearance and Reality*, 2nd ed. London: George Allen, 1897.

Cartwright, R. [SE] ‘Speaking of everything’, *Noûs* **28** (1994): 1-20.

Dummett, M.A.E. [FPL] *Frege: Philosophy of Language*, 2nd edition. London: Duckworth, 1981.

Dummett, M.A.E. [FPM] *Frege: Philosophy of Mathematics*. Cambridge MA, Harvard University Press, 1991.

Forster, T.E. [STUS] *Set Theory with a Universal Set: Exploring an Untyped Universe,* 2nd ed. Oxford: Clarendon Press, 1995.

Friedman, H. [CTE] ‘A complete theory of everything: validity in the universal domain’, www.math.ohio-state.edu/~friedman/ (1999).

Geach, P.T. [RG] *Reference and Generality: An Examination of Some Medieval and Modern Theories*, 3rd ed. Ithaca, NY: Cornell University Press, 1980.

Glanzberg, M. [QR] 'Quantification and realism', *Philosophy and Phenomenological Research*, forthcoming.

Grim, P. [IU] *The Incomplete Universe: Totality, Knowledge and Truth*. Cambridge, MA: MIT Press, 1991.

Harris, J.H. [LLA] 'What's so logical about the "logical" axioms?', *Studia Logica* **41** (1982): 159-171.

Higginbotham, J. [HOLNL] 'On higher-order logic and natural language', in Smiley [PL].

Kaplan, D. [D] 'Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals', in J. Almog, J. Perry and H. Wettstein, eds., *Themes from Kaplan*. Oxford: Oxford University Press, 1989.

Kripke, S.A. [WRPL] *Wittgenstein on Rules and Private Language*. Oxford: Blackwell, 1982.

Lewis, D.K. [PW] *On the Plurality of Worlds*. Oxford: Blackwell, 1986.

Lewis, D.K. [PC] *Parts of Classes*. Oxford: Blackwell, 1991.

McGee, V. [TPTTC] 'Two problems with Tarski's theory of consequence', *Proceedings of the Aristotelian Society* **92** (1992): 273-292.

McGee, V. [E] '"Everything"', in G. Sher and R. Tieszen, eds., *Between Logic and Intuition*. Cambridge: Cambridge University Press, 2000.

McGee, V. [UUQ] 'Universal universal quantification', in J.C. Beall, ed., *Liars and Heaps*. Oxford: Oxford University Press, 2003 (forthcoming).

Quine, W.V. [FLPV] *From a Logical Point of View*. Cambridge, MA.: Harvard University

Press, 1953.

Quine, W.V. [PL] *Philosophy of Logic*. Englewood Cliffs, NJ: Prentice Hall, 1970.

Rayo, A. [WO] 'Words and objects', *Noûs* **36** (2002): 436-464.

Rayo, A. [WEME] 'When does "everything" mean *everything*?', *Analysis* **63** (2003): 100-106.

Rayo, A., and Uzquiano, G. [TTSC] 'Toward a theory of second-order consequence', *Notre Dame Journal of Formal Logic* **40** (1999): 315-325.

Rayo, A., and Williamson, T. [CTUFL] 'A completeness theorem for unrestricted first-order languages', in J.C. Beall, ed., *Liars and Heaps*. Oxford: Oxford University Press, 2003, (forthcoming).

Rayo, A., and Yablo, S. [NDN] 'Nominalism through de-nominalization', *Noûs* **35** (2001): 74-92.

Shapiro, S. [FF] *Foundations without Foundationalism: A Case for Second-Order Logic*. Oxford: Oxford University Press, 2001.

Shapiro, S. [LL] *The Limits of Logic: Higher-Order Logic and the Löwenheim-Skolem Theorem*. Aldershot: Dartmouth, 1996.

Simons, P. [HOQOC] 'Higher-order quantification and ontological commitment', *Dialectica* **51** (1997): 255-271.

Smiley, T., ed. [PL] *Philosophical Logic*. Oxford: Oxford University Press for the British Academy.

Soames, S. [IDD] 'Incomplete definite descriptions', *Notre Dame Journal of Formal Logic* **27** (1986): 349-375.

Stanley, J. [NR] 'Nominal restriction', in G. Peters and G. Preyer, eds., *Logical Form and Language*. Oxford: Oxford University Press, 2002.

Stanley, J., and Szabó, Z.G. [QDR] 'On quantifier domain restriction', *Mind and Language* **15** (2000): 219-261.

Stanley, J., and Szabó, Z.G. [RBN] 'Reply to Bach and Neale', *Mind and Language* **15** (2000): 295-298.

Stanley, J., and Williamson, T. [QCD] 'Quantifiers and context dependence', *Analysis* **55** (1995): 291-295.

Tarski, A. [CLC] 'On the concept of logical consequence', in *Logic Semantics, Metamathematics*, translated by J.H. Woodger, 2nd edition, J. Corcoran, ed. Indianapolis: Hackett, 1983.

Westerståhl, D. [DCS] 'Determiners and context sets', in J. van Benthem and A. ter Meulen eds., *Generalized Quantifiers in Natural Language*. Dordrecht: Foris, 1985.

Williamson, T. [EE] 'Equivocation and existence', *Proceedings of the Aristotelian Society* **87** (1987/8): 109-127.

Williamson, T. [BP] 'Bare possibilia', *Erkenntnis* **48** (1998): 257-273.

Williamson, T. [EC] 'Existence and contingency', *Aristotelian Society*, sup. vol. **73** (1999): 181-203. Reprinted with printer's errors corrected in *Proceedings of the Aristotelian Society* **100** (2000): 321-343 (117-139 in unbound version).

Williamson, T. [NFO] 'The necessary framework of objects', *Topoi* **19** (2000): 201-208.

Wittgenstein, L. [TLP] *Tractatus Logico-Philosophicus*, translated by D.F. Pears and B.F. McGuinness. London: Routledge & Kegan Paul, 1961.

Wright, C.J.G. [FGS] 'Why Frege did not deserve his *granum salis*: a note on the paradox of "The Concept Horse" and the ascription of bedeutungen to predicates', *Grazer Philosophische Studien* **55** (1998): 239-263.

Yi, B-U. [LLP] 'The language and logic of plurals', *Journal of Philosophical Logic*,

forthcoming.