To appear in A. Rayo and G. Uzquiano (eds.), *Unrestricted Quantification: New Essays*
(Oxford: Oxford University Press, forthcoming 2007)

Absolute Identity and Absolute Generality

Timothy Williamson

The aim of this chapter is to tighten our grip on some issues about quantification by analogy with corresponding issues about identity on which our grip is tighter. We start with the issues about identity.

I

In conversations between native speakers, words such as 'same' and 'identical' do not usually cause much difficulty. We take it for granted that others use them with the same sense as we do. If it is unclear whether numerical or qualitative identity is intended, a brief gloss such as 'one thing not two' for the former or 'exactly alike' for the latter removes the unclarity. In this paper, numerical identity is intended. A particularly conscientious and logically aware speaker might explain what 'identical' means in her mouth by saying: ‘Everything is identical with itself. If something is identical with something, then whatever applies to the former also applies to the latter.’ It seems perverse to continue doubting whether ‘identical’ in her mouth means *identical* (in our sense). Yet other interpretations are conceivable. For instance, she might have been speaking an odd idiolect in which ‘identical’ means *in love*, under the misapprehension that everything is in love with itself and with nothing else (narcissism as a universal theory).

Let us stick to interpretations on which she spoke truly. Let us also assume for the time being that we can interpret her use of the other words homophonically. We will make no assumption at this stage as to whether ‘everything’ and ‘something’ are restricted to a domain of contextually relevant objects. We can argue that ‘identical’ in her mouth is coextensive with ‘identical’ in ours. For suppose that an object x is identical in her sense with an object y. By our interpretative hypotheses, if something is identical in her sense with something, then whatever applies to the former also applies to the latter. Thus whatever applies to x also applies to y. By the logic of identity in our sense (in particular, reflexivity), everything is identical in our sense with itself, so x is identical in our sense with x. Thus being such that x is identical in our sense with it applies to x. Consequently, being such that x is identical in our sense with it applies to y. Therefore, x is identical in our sense with y. Generalizing: whatever things are identical in her sense are identical in ours. Conversely, suppose that x is identical in our sense with y. By the logic of identity in our sense (in particular, Leibniz’s Law), if something is identical in our sense with something, then whatever applies to the former also applies to the latter. Thus whatever applies to x also applies to y. By our interpretative hypotheses, everything is identical in her sense with itself, so x is identical in her sense with x. Thus being such that x is identical in her sense with it applies to x. Consequently, being such that x is identical in her sense with it applies to y. Therefore, x is identical in her sense with y. Generalizing: whatever things are identical in our sense are identical in hers. Conclusion: identity in her sense is coextensive with identity in our sense.[1]

Of course, coextensiveness does not imply synonymy or even necessary coextensiveness. Thus we have not yet ruled out finer-grained differences in meaning between her use of ‘identical’ and ours. If we can interpret her explanation as consisting of logical truths, then, given that the principles that the argument invoked about identity in our sense (reflexivity and Leibniz’s Law) are also logical truths, we can show that the universally quantified biconditional linking identity in her sense with identity in ours is a logical truth, so that coextensiveness is logically guaranteed. If the relevant kind of logical truth is closed under the rule of necessitation from modal logic, then the necessitated universally quantified biconditional too is a logical truth, so that necessary coextensiveness is also logically guaranteed.[2] But not even a logical guarantee of necessary coextensiveness suffices for synonymy: we have such a guarantee for ‘cat who licks all and only those cats who do not lick themselves’ and ‘mouse who is not a mouse’, but they are not synonymous. Nevertheless, even simple coextensiveness excludes by far the worst forms of misunderstanding.

But now we must reconsider our homophonic interpretation of all our speaker’s other words. In her second claim ‘If something is identical with something, then whatever applies to the former also applies to the latter’, how much did ‘whatever’ cover? We can regiment her utterance as the schema (2) of first-order logic (for the record, we also formalize her first claim as (1)):

(1) $\forall x\, xIx$

(2) $\forall x \forall y\, (xIy \rightarrow (A(x) \rightarrow A(y))$

Here 'I' symbolizes identity in her sense; $A(y)$ differs from $A(x)$ at most in having the free variable y in some or all places where $A(x)$ has the free variable x. Our speaker will instantiate (2) only by formulas $A(x)$ and $A(y)$ of her language. But our argument for coextensiveness in effect involved the inference from xIy and $x=x$ to $x=y$, where '=' symbolizes identity in our sense. To use (2) for that purpose, we must take $A(x)$ and $A(y)$ to be $x=x$ and $x=y$ respectively. By what right did we treat $x=x$ and $x=y$ as formulas of her language, not merely of ours? Perhaps her language has no equivalent formulas, and both (1) and all instances of (2) in her language are true even though 'I' does not have the extension of identity in our sense.

We can be more precise. Let M be an ordinary model with a nonempty domain D for a first-order language L. Define a new model M* for L as follows. The domain D* of M* contains all and only ordered pairs $<d, j>$, where d is a member of D and j is a member of some fixed index set J of finite cardinality |J| greater than one; pick a member # of J. If R is an n-adic atomic predicate of L, the extension of R in M* contains the n-tuple $<<d_1, j_1>, \ldots, <d_n, j_n>>$ if and only if the extension of R in M contains the n-tuple $<d_1, \ldots, d_n>$. If the constant c denotes d in M, c denotes $<d, \#>$ in M*. More generally, if the n-place function symbol f of L denotes a function φ in M, f denotes the function φ^* in M*, where $\varphi^*(<d_1, j_1>, \ldots, <d_n, j_n>)$ is $<\varphi(d_1, \ldots, d_n), \#>$. Given this definition of M*, it is routine to prove that exactly the same formulas of L are true in M* as in M.[3] Now suppose that the dyadic atomic predicate I of L is interpreted by identity in M; its extension there consists of all pairs $<d, d>$, where d is in D. Thus (1) and all instances of (2) are true in M. Consequently, (1) and all instances of (2) are true in M*. Nevertheless, I is not interpreted by identity in M*, for the extension of I in M* contains the pair $<<d, j>, <d, k>>$ for any member d of D and members j and k of J. Thus in M* everything has the relation for which I stands to |J| things, and therefore to |J|−1 things distinct from itself, although one cannot express that fact in L. Nor does any formula, simple or complex, of L express identity in M*, whether or not identity is expressed in M by some formula.[4] We cannot rule out M* by adding a sentence of L that is false in M* to our theory of identity ((1) and (2)), for any such sentence will also be false in M, the 'intended' model.

We may object that M* is not a model for a first-order language with identity, precisely because it does not interpret any atomic predicate of the language by identity. What distinguishes first-order logic with identity from first-order logic without identity is that the former treats an atomic identity predicate as a logical constant. In standard first-order logic with identity, logical consequence is defined as truth-preservation in all models, and all models are stipulated to interpret that predicate by identity. Unintended interpretations of some basic mathematical terms can be excluded in first-order logic with identity but not in first-order logic without identity.

An example is the concept of linear (total) ordering. In first-order logic with identity, we standardly axiomatize the theory of (reflexive) linear orders (such as ≤ on the real numbers) by (3), (4) and (5):

(3) $\forall x \forall y \forall z\, ((xRy \,\&\, yRz) \rightarrow xRz))$ (transitivity)

(4) $\forall x \forall y\, (xRy \vee yRx)$ (connectedness)

(5) $\forall x \forall y\, ((xRy \,\&\, yRx) \rightarrow x{=}y)$ (anti-symmetry)

The models of this little theory are exactly those in which *R* is interpreted by a reflexive linear order (over the relevant domain). The use of the identity predicate in the anti-symmetry axiom (5) is essential. For if *R* is interpreted by a reflexive linear order, then the open formula *Rxy* & *Ryx* must express identity (over the relevant domain). But we have seen that in first-order logic without identity any theory with a model M has a model M* as above in which no formula expresses identity, therefore in which *R* is not interpreted by a reflexive linear order. Consequently, no theory in first-order logic without identity has as models exactly those in which *R* is interpreted by a reflexive linear order. Since the models of first-order logic with and without identity differ only over the interpretation of the identity predicate, even in first-order logic with identity no theory axiomatized purely by sentences without the identity predicate has as models exactly those in which *R* is interpreted by a reflexive linear order.

The position is substantially the same for irreflexive linear orders. To axiomatize the theory of irreflexive linear orders (such as < on the real numbers), we standardly retain axiom (3) but replace (4) and (5) by (6) and (7):

(6) $\forall x \forall y\, (xRy \lor yRx \lor x=y)$ (linearity)

(7) $\forall x \forall y\, (xRy \rightarrow {\sim}yRx)$ (asymmetry)

The use of the identity predicate in the linearity axiom (6) is essential. For if *R* is interpreted by an irreflexive linear order, then the open formula ~*Rxy* & ~*Ryx* must express identity (over the relevant domain). In first-order logic without identity, any theory with a model has a model in which no formula expresses identity, therefore in which *R* is not interpreted by an irreflexive linear order. Consequently, no theory in first-order logic without identity has as models exactly those in which *R* is interpreted by an irreflexive linear order. Even in first-order logic with identity, no theory axiomatized purely by sentences without the identity predicate has as models exactly those in which *R* is interpreted by an irreflexive linear order.

First-order logic with identity is superior in expressive power to first-order logic without identity in mathematically central ways.[5]

Nevertheless, the appeal to first-order logic with identity may not resolve the doubts of those who take the problem of interpretation seriously. Indeed, it may strike them as cheating. For how do we know that the speaker whom we are trying to interpret is using a first-order language with identity at all? For example, how do we know that she is trying to talk about linear ordering? To pose the problem in less epistemic terms: what makes it the case that the speaker is using a first-order language with identity?

Quine has a short way with bloated models such as M*. He excludes them by his methodology of interpretation, which requires us to interpret the language in such a way that the strongest indiscernibility relation expressible in it is identity. Roughly speaking, he applies *a priori* the inverse of the operation that took M to M*. As he says, we thereby 'impose a certain identification of indiscernibles', adding 'but only in a mild way' (1960: 230). The 'mildness' consists in this: indiscernibility in the relevant sense is the negation of weak discernibility, not of strong discernibility. Objects d and d^* in the domain of the interpretation are strongly discernible if and only if, for some open formula $A(x)$ of L with one free variable (x) and assignments a and a^* to variables of values in the domain, a^* is like a except that $a(x)$ is d while $a^*(x)$ is d^*, and the truth-value of $A(x)$ under a differs from its truth-value under a^*. The definition of weak discernibility is the same except that variables other than x are allowed to occur free in $A(x)$. For example, consider a language with just one atomic predicate, a dyadic one I, without constants or function symbols, and an interpretation with an infinite domain, over which I is interpreted by identity. Let d and d^* be distinct members of the domain. Then d and d^* are not strongly discernible, but they are weakly discernible by the formula xIy and assignments a and a^*, where $a(x)$ is d, $a(y)$ is d^* and a^* is like a except that $a(x)$ is d^*. In this case, Quine's methodology does not erase distinctions between members of the domain: doing so here would involve collapsing the domain to a single object and so switching the formula $\forall x \forall y\, (xIx \rightarrow xIy)$ from false to true. The identification of indiscernibles as he apparently intends it has the hermeneutically appealing feature that it does not alter the truth-value of any formula.[6]

In other examples, Quine's methodology has more radical effects on the model. For instance, consider another language with just two atomic predicates, the monadic F and G, without constants or function symbols, and an interpretation on which 1,000 members of the domain are in the intersection of the extensions of F and G, just one is in the extension of F but not of G, 1,000,000 are in the extension of G but not of F, and just one in the extension of neither. Objects are weakly discernible if and only if either one is in the extension of F while the other is not or one is in the extension of G while the other is not. Thus Quine's 'mild' identification of indiscernibles collapses the 1,000 objects in the intersection of the extensions of F and G into a single object, and the 1,000,000 objects in the extension of G but not F into another single object.[7]

Moreover, Quine's methodology does not preserve the truth-values of all formulas once we add generalized quantifiers to the language. For instance, let us add a binary quantifier M for 'most' to the language in the last example, where $\mathrm{M}x\ (A(x); B(x))$ is true under an assignment a if and only if most (more than half) of the members d of the domain such that $A(x)$ is true under $a[d/x]$ are such that $B(x)$ is true under $a[d/x]$, where the assignment $a[d/x]$ is like a except that $a[d/x](x)$ is d. The addition of M to the language makes no difference to weak discernibility. On the original interpretation, the sentence $\mathrm{M}x\ (Fx; Gx)$ is true, because 1,000 of the 1,001 objects in the extension of F are in the extension of G, while $\mathrm{M}x\ (Gx; Fx)$ is false, because only 1,000 of the 1,001,000 objects in the extension of G are in the extension of F. By contrast, after the identification of indiscernibles, both sentences are false, because exactly one of the two objects in the extension of F is in the extension of G and exactly one of the two objects in the extension of G is in the extension of F. Moreover, no attempt to reinterpret 'M' as a logical quantifier other than 'most' in line with the identification of indiscernibles would preserve the truth-values of all formulas. For the collapsed model is symmetrical between F and G: each applies to exactly one thing to which the other does not. Thus on any interpretation of 'M' as a logical quantifier $\mathrm{M}x\ (Fx; Gx)$ and $\mathrm{M}x\ (Gx; Fx)$ will receive the same truth-value in the collapsed model, whereas they have different truth-values in the new model (see Westerståhl 1989 for logical quantifiers).

Thus the ability of Quine's identification of indiscernibles to preserve the truth-values of all formulas depends on an unwarranted restriction of the language to the usual quantifiers $\forall$ and $\exists$.[8] In the presence of other quantifiers, his identification of indiscernibles does not preserve truth-values, and so is hermeneutically unappealing. Of course, the example was a toy one; the expressive resources of the language were radically impoverished by comparison with any natural human language. Nevertheless, it shows that Quine's methodology does not provide an adequate solution to the problem of interpreting an identity predicate. For if independent considerations have not eliminated interpretations of a natural language on which its domain contains distinct indiscernibles, the use of Quine's methodology to do so risks imposing on the language an interpretation far less charitable to its speakers than some of the 'bloated' interpretations are.

An alternative, anti-Quinean, proposal is to move from first-order to second-order logic. We could then replace the first-order schema (2) with a single second-order axiom:

(2+) $\quad \forall x \forall y\ (xIy \rightarrow \forall P\ (Px \rightarrow Py))$

In the usual models for higher-order logic, the second-order quantifier $\forall P$ is required to range in effect over all subsets of the first-order domain. Given any two objects *d* and *d**, some subset of the domain (for example, $\{d\}$) contains *d* but not *d**. Similarly, if one interprets $\forall P$ as a plural quantifier, there are some objects of which *d* is one and of which *d** is not one. Thus (2+), conjoined with (1), forces the predicate *I* to have the extension of identity over the domain. For practical purposes, one could even use the open formula $\forall P\,(Px \rightarrow Py)$ simply to *define* identity, although it is unlikely that second-order quantification is conceptually more basic than identity in any deep sense.[9] But the appeal to second-order quantification may not satisfy those who are seriously worried about the problem of interpreting the identity predicate. For how do we know, or what makes it the case, that the second-order quantifier $\forall P$ should be interpreted in the standard way? Consider an interpretation of the first-order fragment of the language (with *I* as an atomic predicate) on which not all pairs of distinct members of the domain are weakly discernible, and the extension of *I* contains exactly those pairs of members of the domain that are not weakly discernible. We can now construct a non-standard interpretation of the full second-order language by stipulating that the range of the second-order quantifiers is to be restricted to those subsets of the domain with the property that if *d* is a member and *d** is not weakly discernible from *d* with respect to the first-order fragment then *d** is also a member (a similar stipulation is available for the plural interpretation).[10] It can then be shown that, on this interpretation, objects are weakly discernible with respect to the full second-order language if and only if they are weakly discernible with respect to the first-order fragment. Consequently, not all distinct pairs of members of the first-order domain are weakly discernible with respect to the full second-order language. Nevertheless, (1) and (2+) come out true on such an interpretation even though *I* is not interpreted by identity over the domain. Thus invoking second-order logic only pushes the problem back to that of interpreting higher-order languages.[11]

If we conceive the hermeneutic problem as purely epistemic − how do we know whether another means *identical*? − then we may suppose that it does not arise in the first-person case: how can I be mistaken in thinking that by 'identical' I mean *identical*? But if the problem is constitutive − in virtue of what does another mean *identical* by 'identical'? − then it presumably arises just as much in the first-person case: in virtue of what do I mean *identical* by 'identical'? Despairing of an answer, someone might doubt the very conception of identity that underlies the question. One might even become a relativist about identity in the manner of Peter Geach, with a conception of a predicate's playing the role of *I* relative to a given language, by verifying (1) and all instances of (2) in that language, but reject any conception of its playing the role absolutely, by verifying (1) and all instances of (2) in all possible extensions of the language.[12]

Such a reaction would be grossly premature, resting on no properly worked out, plausible account of interpretation. Questions of the form 'In virtue of what do we mean *X* by "X"?' are notoriously hard to answer satisfactorily, no matter what is substituted for 'X' (Kripke 1982). It is therefore methodologically misguided to treat a particular expression (for instance, 'identical') as problematic merely on the grounds that the question is hard to answer satisfactorily for it.[13] Of course, the details of the alternative interpretations and surrounding arguments depend on the nature of the expression at issue, but that should not cause us to overlook the generality of the underlying problem. It is extremely doubtful that the sceptical reaction yields anything coherent when generalized. In the particular case of identity, few have found Geach's arguments for his local scepticism convincing or his relativism plausible. In any case, let us suppose that we do use 'identical' in an absolute way, and ask in virtue of what we do so. Given what has just been said, we should not expect more than a sketchy answer.

Somehow or other, 'identical' means what it does because we use it in the way that we do. A central part of that use concerns our inferential practice with the term, as summarized by (1) and (2) or (2+). What is crucial in our use of the first-order schema (2) (or a corresponding first-order inference rule) is that we do not treat it as exhausted by its instances in our current language. Rather, we have a general disposition to accept instances of (2) in extensions of our current language. That is not to say that in all circumstances in which we are or could be presented with an instance of (2) in an extension of our current language, we accept it. Obviously, we may reject it as a result of computational error, or die of shock at the sight of it, or roll our eyes as a protest against pedantry; some instances may be too long or complex to be presented to us at all. But the existence of a large gap of that kind between the disposition and the conditionals is the normal case for dispositions, including ordinary physical dispositions such as fragility and toxicity: external factors of all sorts (such as antidotes) can intervene between them and their manifestations. The link between the disposition to D if C and conditionals of the form 'If C, it Ds' is of a much looser sort. The failure of some of the associated conditionals does not show the absence of the disposition.[14] We have the general disposition because we respond, when we do, to the general form of the schema (2) (or of a corresponding inference rule) rather than treating each of its instances as an independent problem. Our non-intentionally described behaviour alone does not make it intelligible why we count as responding to the actual form of (2) and not to some gerrymandered variant on it with *ad hoc* restrictions for cases beyond our ken: it is also relevant that the actual form is more natural than the gerrymandered one in a way that fits it for being meant (it is a 'reference magnet').[15]

Our understanding of (2) as transcending the bounds of our current language is already suggested by our use of the phrase 'Leibniz's Law' for a single principle. We do not usually think of the phrase as ambiguously denoting lots of different principles, one for each language.

In the case of the second-order axiom (2+), what is crucial is that we do not treat the rule of universal instantiation for the second-order quantifier as exhausted by its instances in our current language. Rather, we have a general disposition to accept instances of universal instantiation for the second-order quantifier in extensions of our current language. Again, the presence of the disposition is consistent with the failure of some of the associated conditionals. We have the general disposition because we respond to the general form of universal instantiation rather than treating each of its instances as a separate problem.

The sort of open-ended commitment just described is typical of our commitment to rules of inference. For example, my commitment to reasoning by disjunctive syllogism is not exhausted by my commitment to its instances in my current language; when I learn a new word, I am not faced with an open question concerning whether to apply disjunctive syllogism to sentences in which it occurs. Indeed, open-ended commitment may well be the default sort of commitment: one's commitment is open-ended unless one does something special to restrict it.

Open-ended commitment is just what is needed to reinstate the argument given at the beginning for a homophonic interpretation of another's use of the word 'identical', given her commitment to (in effect) (1) and (2). We reason as follows. Let agents S and S* speak distinct first-order languages L and L* respectively. For simplicity, assume that S and S* coincide in their logical vocabulary, with the possible exception of an identity predicate. Their interpretation of the other logical vocabulary is assumed to be standard; although the present style of argument can be extended to the other logical vocabulary, that is not our present concern. Let I be a predicate of L but not of L* and I^* a predicate of L* but not of L. Suppose that S has an open-ended commitment to (1) and (2), while S* has an open-ended commitment to (1*) and (2*), the results of substituting I^* for I in (1) and (2):

(1*) $\forall x\, xI^*x$

(2*) $\forall x \forall y\, (xI^*y \rightarrow (A(x) \rightarrow A(y))$

Now merge L and L* into a single first-order language L+L* whose primitive vocabulary is the union of the primitive vocabularies of L and L*. Thus we can treat (1) and (1*) as sentences of L+L* and (2) and (2*) as schemas of L+L*. The interpretation of the logical vocabulary of L+L* is assumed to be standard, like that of L and L*, again with the possible exception of an identity predicate. The quantifiers of L+L* are interpreted as ranging over the intersection of the domains of the quantifiers of L and of L*, for our present question is in effect whether I and I^* can diverge for objects over which both are defined. By the open-endedness of their commitments, S is committed to (1) as a sentence of L+L* and to all instances of (2) in L+L*, while S* is committed to (1*) as a sentence of L+L* and to all instances of (2*) in L+L*. Here are instances of (2) and (2*) respectively in L+L*:

(8) $\forall x \forall y\, (xIy \rightarrow (xI^*x \rightarrow xI^*y)$

(8*) $\forall x \forall y\, (xI^*y \rightarrow (xIx \rightarrow xIy)$

Reasoning in L+L*, we deduce (9) from (1*) and (8) and (9*) from (1) and (8*):

(9) $\forall x \forall y\, (xIy \rightarrow xI^*y)$

(9*) $\forall x \forall y\, (xI^*y \rightarrow xIy)$

Thus, given the pooled commitments of S and S*, I and I^* are coextensive over the common domain.

The result should not be interpreted as concerning only the extensions of I and I^* in a new context created by the fusion of L and L*. For the open-ended commitments in play were incurred by S and S* in using I and I^* in the original contexts for L and L* respectively; that is the nature of open-endedness. Thus the result concerns the extensions of *I* and *I** in the original contexts for L and L* too.

In the way just seen, (1) and the open-ended schema uniquely characterize identity (recall that the other logical vocabulary in (9) and (9*) is being given its standard interpretation). A similar argument can be given for the second-order analogue of (2). The arguments are in fact a special case of a more general pattern of reasoning that shows all the usual logical constants to be uniquely characterized by the classical principles of logic for them.[16]

Of course, these remarks fall far short of a fully satisfying account of what makes ‘identical’ mean *identical*. The connection between which logical principles a speaker *accepts* for a given expression and which logical principles are *correct* (true or truth-preserving) for that expression is quite loose; no reasonable principle of charity in interpretation guarantees freedom from logical error. Misguided philosophers who reject standard logical principles for ‘identical’ probably still mean *identical* by the word, because they continue to use it as a word of the common language.[17] It is a fallacy to reason from the premise that someone has a wildly deviant theory to the conclusion that they speak a deviant language. Perhaps even more misguided philosophers could argue themselves into the view that the logic of the phrase ‘in love’ comprises analogues of the standard logical principles for ‘identical’, while still meaning *in love* rather than *identical* by the phrase, because they continue to use it as a phrase of the common language. Nevertheless, such examples do not suggest that the account above of how ‘identical’ means *identical* does not at least point in the right direction, for they may still be parasitic on a loose underlying connection between inferential practice and meaning.

Naturally, the account will not help if it is incoherent, as Geach would claim it to be. According to him, generalizing over all predicates in all possible extensions of the language generates semantic paradoxes (1972: 240; 1991: 297). Now the foregoing account does not assume that speakers who reason with (1) and (2) or (2+) must themselves have a conception of all predicates in all possible extensions of the language. For the sentences involved in such reasoning need not be metalinguistic. If speakers already have metalinguistic vocabulary in the language, they can use it in what they substitute for $A(x)$ and $A(y)$ in (2), but that does not imply that the schema itself is distinctively metalinguistic. Agents may lack the conceptual apparatus necessary to give a reflective account of their own practices. However, the foregoing theoretical account does deploy something reminiscent of the conception of all predicates in all possible extensions of the language on its own behalf, in explaining the nature of speakers' open-ended commitments. Thus Geach's charge is at least relevant.

Unfortunately, Geach does not bother to argue in detail that the theorist of absolute identity really requires conceptual resources powerful enough to generate semantic paradoxes. In fact, when we consider identity over a given set domain, we need only generalize over all subsets of that domain, or over all subsets of the Cartesian product of the domain with itself.[18] For the unique characterization argument above, we need merely consider expansions of the language by a single dyadic atomic predicate I^*, whose extension is a subset of the Cartesian product of the original domain with itself. Similarly, schema (2) forces I to have the extension of identity over the domain as soon as we consider its instances in expansions of the language by a single monadic atomic predicate, whose extension is a subset of the original domain. In standard (Zermelo-

Fraenkel) set theory, sets are closed under the power set operation and the formation of Cartesian products. Thus quantification over subsets of the original set domain or of its Cartesian product with itself is quantification over another set domain. Consequently, the kind of generalization required for the theorist of absolute identity over a given set domain is of a quite harmless sort. It poses no serious threat of semantic or set-theoretic paradox. Some mysteries about the power set operation remain unsolved, notably Cantor's continuum problem (how many subsets has the set of natural numbers?), but they are not paradoxes. In any case, they are largely independent of the application to identity, for they concern the size of the whole power set, whereas in order to characterize identity it suffices to have just the singleton sets of members of the original domain; there are no more singletons of members than members. Geach's argument for the incoherence of absolute identity theory does not withstand attention.

Suppose that the foregoing account of our grasp of absolute identity is correct, as far as it goes. What does it suggest about our grasp of absolute generality, generality over absolutely everything, without any explicit or implicit restrictions whatsoever?

II

Sympathetic readers will have felt little difficulty in understanding the words 'absolute generality, generality over absolutely everything, without any explicit or implicit restrictions whatsoever': but in principle those words are open to alternative interpretations. 'Absolute' might be read as itself relative to a background contextually supplied standard, and the quantifiers ‘any’ and ‘whatsoever’ over restrictions as themselves contextually restricted. One can find oneself saying ‘By “everything” I mean *everything*’ with the same desperate intensity with which one may say ‘By “identical” I mean *identical*’. How deep does the similarity of the interpretative challenges go?

Let us start with the question of unique characterization. Here are standard rules for a (first-order) universal quantifier:[19]

$\forall$-Introduction	Given a deduction of A from some premises, one may deduce $\forall v\, A(v/t)$ from the same premises, where $A(v/t)$ is the result of replacing all occurrences of the individual constant t in the formula A by the individual variable v, provided that no such occurrence of v is bound in $A(v/t)$ and that t occurs in none of the premises.
$\forall$-Elimination	From $\forall v\, A$ one may deduce $A(t/v)$, where $A(t/v)$ is the result of replacing all free occurrences of the individual variable v in the formula A by the individual constant t.

Consider a universal quantifier $\forall$ in a language L governed by those rules, and another universal quantifier $\forall^*$ in a language L* governed by exactly parallel rules, $\forall^*$-Introduction and $\forall^*$-Elimination. Suppose that the commitment of speakers of L and L* to their principles is open-ended in the way discussed above for the case of the identity rules. Merge L and L* into a single language L+L*, whose primitive vocabulary is the union of the primitive vocabularies of L and L*. Let A be a formula of L+L* in which the individual constant t does not occur and no variable except v occurs free. We reason in L+L*. From $\forall v A$ we can deduce $A(t/v)$ by $\forall$-Elimination. Therefore, since t does not occur in the premise and A is the result of replacing all occurrences of t in $A(t/v)$ by v, and no such occurrence of v thereby becomes bound in A, from $\forall v A$ we can deduce $\forall^* v A$ by $\forall$*-Introduction. Conversely, from $\forall^* v A$ we can deduce $A(t/v)$ by $\forall$*-Elimination, and therefore $\forall v A$ by $\forall$-Introduction. Thus, given the pooled commitments of speakers of L and L*, the two quantifiers are logically equivalent.[20]

The result should not be interpreted as concerning only the reference of $\forall$ and $\forall$* in a new context created by the fusion of L and L*. For the open-ended commitments in play were incurred by S and S* in using $\forall$ and $\forall$* in the original contexts for L and L* respectively; that is the nature of open-endedness. Thus the result concerns the reference of $\forall$ and $\forall$* in the original contexts for L and L* too.

Observe that the argument for unique characterization did not proceed by semantic analysis of the quantifier. It did not invoke the idea of unrestricted generality. In particular, the argument was not that only an unrestricted interpretation of $\forall$ validates $\forall$-Introduction and $\forall$-Elimination. Rather, it was a syntactic argument for interderivability. Thus it is not circular to use the unique characterization result to support the claim that we have an idea of unrestricted generality.

Nevertheless, it is tempting to suspect the argument for unique characterization of sophistry. For $\forall$-Introduction and $\forall$-Elimination are standard rules for a universal quantifier in standard first-order logic, for which the standard model theory interprets the quantifier as restricted to the domain of a model. It may therefore look as though the argument must prove too much, since $\forall$-Introduction and $\forall$-Elimination are valid if $\forall$

is interpreted over a domain D, while $\forall$ *-Introduction and $\forall$ *-Elimination are valid even if $\forall$ * is interpreted over a distinct domain D*.

In that form, the objection is unthreatening, for it neglects the stipulation that the commitment of speakers of L and of L* to the respective quantifier rules is open-ended in the sense explained in part I. If $\forall$ is restricted to a domain D, then speakers of L do not have an open-ended commitment to $\forall$ -Elimination, even if the latter has no counter-instance in L, since it has the potential for a counter-instance with a new term *t* that denotes something outside D in a language such as L+L*. In the setting of L+L*, $\forall$ -Elimination would require an extra premise involving *t* to the effect, concerning what *t* denotes, that it belongs to D. Then *t* would occur in one of the premises from which $A(t/v)$ was deduced, so the condition for the application of $\forall$ *-Introduction would not be met. Of course, if $\forall$ * were restricted too, to a domain D*, then one might modify $\forall$ *-Introduction accordingly, by allowing *t* to occur in one extra premise of the envisaged deduction of *A* to the effect, concerning what *t* denotes, that it belongs to D*. With no guarantee that D includes D*, however, the deduction of $\forall *v\, A$ from $\forall\, v\, A$ could not be carried through. The converse deduction faces an exactly analogous problem. It may sometimes be hard to know whether a given speaker's commitment is open-ended, but the considerations of part I indicated that open-ended commitment to a rule is a genuine, recognizable phenomenon, and no reason has emerged to view the quantifier rules as exceptional in that respect. Indeed, as before, open-ended commitment may be the default sort of commitment to $\forall$ -Introduction and $\forall$ -Elimination; it would be implausible to suggest that all speakers are always doing something special to override the default.

The pertinent objection is not to the argument for unique characterization from the open-ended understanding of the quantifier rules. Rather, it is to the open-ended understanding of the quantifier rules itself. More precisely: it is not straightforward that the open-ended versions $\forall$-Introduction and $\forall$-Elimination are really valid on the unrestricted reading of the quantifier.

Let us take $\forall$-Elimination first. Free logicians will object to it that, even if $\forall$ is supposed to be unrestricted, the rule is too strong because the individual constant *t* may be an empty name. For example, $\forall$-Elimination enables us to derive from the logically true premise $\forall y \exists x\, x=y$ the conclusion $\exists x\, x=t$, which is false on the unrestricted reading of the quantifier if *t* denotes nothing whatsoever.

One response to that objection to $\forall$-Elimination is that the only role of *t* in the unique characterization argument is as an arbitrary name, which functions like a free variable. The success of any empirical or conceptual process of reference-fixing for *t* is irrelevant to the argument. Thus, if one shares the free logicians' qualms, one can replace 'individual constant' by 'arbitrary name' or 'free variable' throughout $\forall$-Introduction and $\forall$-Elimination, and envisage *t* as denoting merely relative to an assignment.

A less concessive response can also be made. In assessing validity, our concern is with truth-preservation only when the relevant formulas are fully interpreted. For a sentence to be fully interpreted, it is not enough that it is a meaningful formula of the language; it must also express a proposition as used in the relevant context. For example, although 'This is that' is a meaningful sentence of English, it fails to express a proposition in a context in which no reference has been assigned to the demonstratives 'this' and 'that'. It would be foolish to object to the usual introduction rule for disjunction (deduce a disjunction from any of its disjuncts) that it takes one from the true premise '2+2 = 4' to the conclusion '2+2 = 4 or this is that', which is not true when no reference has been assigned to 'this' and 'that' because it expresses no proposition (plausibly, a disjunction expresses a proposition only if each of its disjuncts does). Truth-preservation is required only once any singular terms in the argument have been assigned a reference. Consequently, the free-logical objection to $\forall$-Elimination fails. Let us adopt this conception of validity and therefore leave $\forall$-Elimination unmodified.

'Inclusive' logicians object to $\forall$-Elimination because it does not allow for the empty domain: given axiom (1) for identity, one can prove $\exists x\, x=x$ (as an abbreviation of $\sim\forall x \sim x=x$), which is false in that domain.[21] But once the language contains unbound singular terms, as ours does, then it cannot be fully interpreted in the sense just sketched over the empty domain. For such languages, our notion of validity excludes the empty domain. More controversially, one can argue that $\exists x\, x=x$ is a logical truth on the unrestricted interpretation of the quantifiers, by appeal to Tarski's model-theoretic account of logical truth: it is true on all models (interpretations) that preserve the intended interpretations of the logical constants in it, for it is true and it contains no nonlogical constituents.[22] Since Tarski (1936) understood a model as an assignment of reference to the nonlogical atoms of the language (more exactly, as an assignment of values to variables, which replace those atoms), his treatment of interpretation is consistent with the notion above of a fully interpreted formula. The relevance of Tarski's conception of logical truth and logical consequence to the logic of unrestricted quantification will be discussed more fully below.

Given an appropriate notion of validity, open-ended $\forall$-Elimination is valid on the unrestricted reading of the quantifier. What of $\forall$-Introduction? The obvious worry is this. Suppose that $\forall$ is unrestricted while $\forall^*$ is restricted to a domain D*, and that the term t is constrained by a 'meaning postulate' to be in D*. Thus from $\forall^*x\, x \in D^*$ alone we can infer $t \in D^*$ by the restricted elimination rule for $\forall^*$; since t does not occur in the premise, $\forall$-Introduction therefore permits us to infer $\forall x\, x \in D^*$, which is false because D* does not contain absolutely everything, from the true $\forall^*x\, x \in D^*$. What has gone wrong here is that $t \in D^*$ was not *freely* deduced from $\forall^*x\, x \in D^*$, in a sense of 'free' that has nothing to do with free logic. The deduction was unfree in the sense that it invoked special rules that required t to satisfy constraints beyond simply being a singular term that occurs in none of the premises. That is still a syntactic feature of the deduction. Let us therefore read 'deduction' in $\forall$-Introduction as 'free deduction', and 'deduced' in both $\forall$-Introduction and $\forall$-Elimination as 'freely deduced'. With that understanding, both $\forall$-Introduction and $\forall$-Elimination are valid on the unrestricted reading of the quantifier.

Given that the two rules are materially valid on the unrestricted reading, someone might still worry that they are not strictly logically valid, because the unrestricted reading of the quantifier is not part of its logic. The discussion so far has been framed in terms of the background assumption that the unrestricted universal quantifier as such should be classified as a logical constant, subject to rules of inference that exploit its unrestrictedness. For those who admit the coherence of unrestricted quantification, the salient alternative is to have as a logical constant a universal quantifier such that, for any things whatsoever, for purposes of defining logical truth and logical consequence the quantifier can be interpreted as ranging over those things and nothing else (interpretations here play the role of models). It does not matter whether there are too many of those things to form a set or set-like domain. The quantifier can be legitimately interpreted as ranging over all things whatsoever, or over all sets whatsoever, but it can also be legitimately interpreted as ranging over just the books on my shelves. A conclusion is a logical consequence of some premises only if, however they are legitimately interpreted, the conclusion is true if the premises are.[23] On this view, open-ended $\forall$ -Elimination is invalid, since in a language such as L+L* with different sorts of quantifier a term *t* may denote something in the domain of one of those other quantifiers that is not in the domain of $\forall$ on some unintended interpretation. Let the *constant account* be that on which the quantifier is mandatorily interpreted as unrestricted and the *variable account* be that on which, for any things, it can legitimately be interpreted as ranging over just those things. Both accounts are framed within a broadly Tarskian approach to the concept of logical consequence.

The difference between the two accounts has dramatic implications for first-order logic. Let $\exists n$ be the usual first-order formalization of the claim that there are at least n things, where $n > 1$. On the variable account, $\exists n$ is not a logical truth, because the quantifier can legitimately be interpreted as ranging over fewer than n things. On the constant account, $\exists n$ is a logical truth, because the quantifier must be interpreted as ranging over absolutely everything and there are in fact at least n things: for example, at least n symbols occur in $\exists n$ itself. The sentences $\exists n$ for all natural numbers n turn out to exhaust the extra logical consequences generated by the constant account, in the sense that the result of adding them as extra axioms to first-order logic can be proved complete as well as sound on the constant account.[24] By contrast, the variable account is logically conservative: it delivers exactly the same logical consequence relation for first-order logic as does the standard model theory with set domains.[25]

The difference between the two accounts is robust. Even if we relativize interpretations to various parameters for the context of utterance or circumstance of evaluation, the sentences $\exists n$ still all come out as logical truths on the constant account, because the unrestricted reading of the quantifier forbids us to interpret it as ranging only over a domain associated with the context of utterance or circumstance. Someone might reply that if the domain contains everything that exists in the relevant possible world then the restriction is merely apparent, because in that world there is nothing else to quantify over. But that objection in effect treats the semantic clause for $\forall$ as though it were a misleading approximate translation of a more fundamental semantic clause in which an unrestricted universal quantifier of a more fundamental modal meta-language occurs within the scope of a modal operator. But the Tarskian framework for the theory of logical consequence is not a modal one. It defines logical consequence without using modal operators, interpreted metaphysically or epistemically. The non-modal meta-language should therefore be taken at face value: a semantic clause according to which $\forall$ ranges only over the domain of some world is inconsistent with the unrestricted reading, just as it appears to be. Indeed, it is part of Tarski's great achievement to have cleanly separated the concept of logical consequence from metaphysical and epistemic clutter. Not that there is anything wrong with metaphysical and epistemic modalities in their place: but it is methodologically wrong-headed to mix them up with the simple but powerful non-modal concept of logical consequence that Tarski painstakingly isolated, which compels and rewards investigation in its own right. This chapter works within the Tarskian paradigm.

The dispute between the variable and constant accounts raises deep questions about the metaphysical and epistemological status of logic that we cannot hope to answer here. But the analogy with identity does help us to see what is wrong with one argument against the constant account. It is sometimes urged that the variable account is preferable because it has greater generality, since every legitimate interpretation on the constant account is also legitimate on the variable account (for it allows us to interpret the quantifier as ranging over absolutely everything), but not *vice versa*. However, there is an analogous argument against first-order logic with identity, according to which first-order logic without identity is preferable because it has greater generality, since every legitimate interpretation in first-order logic with identity is also legitimate in first-order logic without identity (for it allows us to interpret a dyadic predicate by identity), but not *vice versa*. The latter argument clearly fails, because the point of making identity a logical constant is to capture its distinctive logic by excluding unintended interpretations. No significant generality is thereby lost, because all the other interpretations can be shifted to other dyadic predicates. The former argument against the constant account fails similarly, because the point of making the unrestricted universal quantifier a logical constant is to capture its distinctive logic by excluding unintended interpretations. No significant generality is thereby lost, because all the other interpretations can be captured by complex restricted quantifiers consisting of the simple unrestricted quantifier and a restricting predicate.

Of course, we have some sense of which expressions deserve to be treated as logical constants: very roughly, those whose meaning is 'purely structural'.[26] By that standard, the unrestricted universal quantifier is at least as good a candidate as identity is. Moreover, like identity, the unrestricted quantifier has the kind of stark simplicity in meaning that we seek in a logical constant that is to be treated as basic (some purely structural meanings are very complicated). Although unrestricted quantification is less central to mathematical reasoning than identity is, it does enable us to capture the generality that principles of a set theory with ur-elements (non-sets) such as ZFU need if mathematics is to have its full range of applications: for absolutely *any* objects x and y, there is a set of which x and y are members, for example.

In using absolute identity to support absolute generality, we must be careful to check that the latter does not squash the former. For Geach's arguments against the coherence of absolute identity look superficially more formidable in the context of absolute generality. The previous section considered absolute identity over a set-sized domain; despite Geach's threats, no danger of paradox arises in characterizing it by quantifying over subsets of the domain, or plurally over members of the domain, or the like. But if our first-order quantifiers are absolutely unrestricted, then the interpretation of the corresponding second-order quantifiers is a much trickier business. There is no set domain over whose subsets they could range. Indeed, any attempt to interpret them in terms of values of the second-order variables will generate a version of Russell's paradox, given an adequately strong comprehension principle concerning the existence of such values, since an absolutely unrestricted first-order quantifier must range over them too. But not even considerations of this kind can rescue the charge of incoherence against absolute identity. First, the paradox results from the attempt to interpret the second-order quantifiers of the object-language in a first-order meta-language, by first-order quantification over sets. If one interprets the second-order quantifiers more faithfully, in a second-order meta-language, by second-order quantification read plurally or in some other non-first-order way, then no paradox results. Second, even if one interprets the second-order quantifiers as ranging over 'small' sets, with an appropriately qualified comprehension principle, that suffices for characterizing identity, although not for all other purposes. Third, if absolute identity is coherent for each set-sized domain, then it is simply coherent: for any objects o and o^*, absolute identity is coherent over the set-sized domain $\{o, o^*\}$ by hypothesis, which is all that we need coherently to ask whether o and o^* are absolutely identical.

If there is a threat of paradox, it comes from the idea of absolute generality, not from that of absolute identity. A 'paradox' here is a proof of an explicit contradiction from premises to which generality absolutists are committed by rules of inference to which they are also committed, something like the Russell or Burali-Forti paradox. The greatest strength of generality relativism is the suspicion that generality absolutism is ultimately inconsistent because it leads to such a paradox.

Generality relativists also tend to use a second sort of argument: that generality absolutism is inarticulate, in the sense that whatever utterances generality absolutists assent to or dissent from in trying to articulate their position, on some generality relativist interpretations all the assents were to truths and all the dissents from falsehoods (according to the generality relativist). In that sense, absolutism about identity is also inarticulate: whatever utterances identity absolutists assent to or dissent from in trying to articulate their position, on some identity relativist interpretations all the assents were to truths and all the dissents from falsehoods (according to the identity relativist). It does not follow that absolute identity is inexpressible, for all those interpretations were incorrect: as explained in section I, they misidentified our dispositions to make inferences using the identity predicate. Similarly, if generality absolutism is inarticulate, it does not follow that absolute generality is inexpressible, for the generality absolutist can argue that all the generality relativist interpretations were incorrect, because they misidentified our dispositions to make inferences using the universal quantifier. The generality absolutist may endorse a reflection principle to the effect that any quantified sentence true on the intended unrestricted interpretation of the quantifier is also true on some unintended restricted interpretation of the quantifier. Nevertheless, such a result does make the generality absolutism and its negation somewhat elusive for purposes of theoretical dispute.

It may be less widely appreciated than it should be that the generality relativist cannot combine the two objections to generality absolutism, by charging that it is *both* inconsistent *and* inarticulate. For suppose that generality absolutism is both inconsistent and inarticulate. Since it is inconsistent, there is a proof of an explicit contradiction from premises to which generality absolutists are committed by rules of inference to which they are committed. By hypothesis, generality absolutism is also inarticulate, so on some generality relativist interpretation the premises of the proof are true (according to the generality relativist) and the rules of inference are truth-preserving (according to the generality relativist). Thus the generality relativist is committed to the truth of the conclusion of the proof on the generality relativist interpretation. But the conclusion is a contradiction, and so is not true even on that interpretation (according to the generality relativism). Thus generality relativism is inconsistent. To sum up: if generality absolutism is inarticulate, then it is inconsistent only if generality relativism is also inconsistent. Therefore, the generality relativist is ill-advised to accuse generality absolutism of being both inconsistent and inarticulate. In effect, the assumption that generality absolutism is inarticulate yields a consistency proof for generality absolutism relative to generality relativism.

Perhaps we can generalize the result. Given any argument against generality absolutism, why can't it be reinterpreted as an argument against something that the generality relativist accepts, if generality absolutism really is inarticulate?

Generality relativists seem to be faced with a choice. They can drop the charge of inarticulacy, try to explain why it seemed compelling, and then, treating generality absolutism as an articulated theory, try to prove a contradiction in it. If they succeed in that, they win (dialetheism is a fate worse than death). But if they cannot produce such a proof, then they had better drop the charge of inconsistency: put up or shut up. Alternatively, they can drop the charge of inconsistency right away, and press the charge of inarticulacy. But that charge is hardly damaging to absolutism about generality, for it applies equally to absolutism about identity.[28]

Notes

1 The argument goes back to Quine 1961; see the reprinted version in Quine 1966 at 178.

2 The logic of indexicals is arguably not closed under the rule of necessitation (Kaplan 1989). Such problems do not seem to arise for (1) and (2).

3 Sketch of proof: Let *a* be any assignment of values in D* to variables. Let *a*^ be the assignment of values in D to variables such that for any variable *v*, if *a*(*v*) is <*d*, *j*> then *a*^(*v*) is *d*. It is routine that for any term *t*, if *t* denotes <*d*, *j*> in M* relative to *a* then *t* denotes *d* in M under *a*^ (by induction on the complexity of *t*). Thus any atomic formula $Rt_1\ldots,t_n$ is true in M* under an assignment *a* if and only if it is true in M under *a*^, by definition of the extension of *R* in M*. We show that any formula *A* of L is true in M* under an assignment *a* if and only if it is true in M under *a*^ by induction on the complexity of *A*. The induction step for the truth-functors is trivial. For the universal quantifier, the induction hypothesis is this: for every assignment *b* of values in D*, *A* is true in M* under *b* if and only if it is true in M under *b*^. Suppose that $\forall x A$ is not true in M* under an assignment *a*. Then, under some assignment *b* of values in D* that differs from *a* at most over *x*, *A* is not true in M*. By induction hypothesis, *A* is not true in M under *b*^. By construction, *b*^ differs from *a*^ at most over *x*. Thus $\forall x A$ is not true in M under *a*^. Conversely, suppose that $\forall x A$ is not true in M under *a*^. Then, under some assignment *a*^! of values in D that differs from *a*^ at most over *x*, *A* is not true in M. Let *b*

be the assignment of values in D* like *a* except that *b*(*x*) is <*a*^!(*x*), #>. Thus *b*^ is *a*^!, so *A* is not true in M under *b*^. By induction hypothesis, *A* is not true in M* under *b*. Since *b* differs from *a* at most over *x*, ∀ *x A* is not true in M* under *a*. Thus any formula *A* is true in M* under an assignment *a* if and only if it is true in M under *a*^. Finally, we show that a formula is true in M* iff it is true in M. If *A* is not true in M*, then, under some assignment *a* of values in D*, *A* is not true in M*, so *A* is not true in M under *a*^, so *A* is not true in M. Conversely, if *A* is not true in M, then, under some assignment $ of values in D, *A* is not true in M; but $ is *a*^ for some assignment *a* of values in D* (since we can set *a*(*v*) to be <$(*v*), #>), so *A* is not true in M* under *a*, so *A* is not true under *a*. The whole construction is adapted from the standard proof of the upward Löwenheim-Skolem theorem for first-order logic without identity. All it really requires is a homomorphism in a suitable sense from M* onto M; thus it is inessential that an equal and finite number of members of D* are mapped to each member of D.

4 Proof: Suppose that *A*(*x,y*) expresses identity in M*: for every assignment *a* of values in D*, *A*(*x,y*) is true in M* under *a* if and only if *a*(*x*) is *a*(*y*). For some member *d* of D, let *a* be an assignment of values in D* such that both *a*(*x*) and *a*(*y*) are <*d*, #>. By hypothesis, *A*(*x,y*) is true in M* under *a*. Let *b* be an assignment of values like *a* except that *b*(*y*) is <*d*, ##> for some member ## of J distinct from #. By hypothesis, *A*(*x,y*) is not true in M* under *b*, since *b*(*x*) is not *b*(*y*). By the argument of the previous footnote and in its notation: *A*(*x,y*) is true in M* under *a* if and only if it is true in M under *a*^; *A*(*x,y*) is true in M* under *b* if and only if it is true in M under *b*^. But *a*^ is *b*^. Thus *A*(*x,y*) is true in M* under *a* if and only if it is true in M* under *b*. Contradiction.

5 For the view that identity is not a logical constant, which puts first-order logic with identity in an anomalous position, see Peacocke 1976.

6 The text does not follow Quine in inessential details. He speaks of the satisfaction of a formula with a given number of free variables by that number of objects in a given order, rather than of the truth of a formula under an assignment of objects to all variables. Quine 1960: 230 incorrectly claims that the relevant kind of discernibility is relative discernibility: satisfaction of an open formula with two free variables by the objects in only one order. Quine 1976 implicitly corrects the mistake. In the example in the text (taken from that article), no two objects are even relatively discernible. The text attributes a domain to the interpretation, perhaps contrary to Quine's intentions, in order to make it clear that the argument here does not rely on contentious premises about unrestricted quantification.

7 See Wiggins 2001: 185 for discussion of a related example, attributed to Wallace 1964.

8 The addition of generalized quantifiers impacts on the 'bloated' model M*. No problem arises for 'many', since the construction preserves the ratios between the cardinalities of the extensions of predicates (for the index set J was stipulated to be finite), but numerical quantifiers must be reinterpreted in order to give all formulas the same truth-values as in M: 'at least *m*' is interpreted as 'at least $m|J|$'.

9 See also Shapiro 1991: 63. Since the values over which the second-order quantifier ranges are closed under complementation relative to the individual domain, strengthening the conditional in the *definiens* to a biconditional would make no difference.

10 The construction can be generalized to polyadic predicate variables and to orders greater than two, if desired.

11 Such non-standard models are models in the sense of the non-standard semantics with respect to which Henkin 1950 proves the completeness of higher-order logic. As Shapiro (1991: 76) remarks, 'Henkin semantics and first-order semantics are pretty much the same'. It is therefore no surprise that, once Henkin models are allowed, higher-order logic is no advance on first-order logic in solving the interpretation problem. Similarly, the substitution of Henkin semantics for the standard semantics throws away the advantages of second-order logic over first-order logic as a setting for mathematical theories. For example, the result that all models of second-order arithmetic are mutually isomorphic holds only for the standard semantics; non-standard models of first-order arithmetic can be simulated by appropriate Henkin models of second-order arithmetic.

12 Geach gives his views on identity in his 1967, 1972, 1980, 1991 and elsewhere. For critical discussion of them see Dummett 1981: 547-583 and 1993, Noonan 1997 and Hawthorne 2003: 111-123. See also Wiggins 2001: 21-54.

13 Geach can allow that in a suitable context 'identical' may means *same F*, understood as identity relative to the informative sortal expression *F*; the question concerns the use of 'identical' in contexts that supply no such sortal.

14 On the relation between dispositions and conditionals see Martin 1994, Lewis 1997, Bird 1998 and Mumford 1998. For the specific application to rule-following see Martin and Heil 1998.

15 On the role of naturalness in the constitution of meaning see Lewis 1983. Of course, we may have to check putative instances of (2) to ensure that they really have the form that they appear to have and do not involve intensional or quotational contexts, shifts of reference in indexicals and so on. But such problems are not at the heart of the dispute between absolutists and relativists about identity.

16 For further discussion and references see Harris 1982, Williamson 1987/8 and McGee 2000.

17 See Williamson 2003b for arguments of this type.

18 The Cartesian product of sets X and Y is the set of all ordered pairs whose first member belongs to X and second member belongs to Y.

19 For simplicity, functional terms are ignored. To qualify a variable or constant as 'individual' is just to say that it occupies singular term position.

20 See McGee 2000 and Rayo 2003 for related discussion. For opposed views see Dummett 1981 and Glanzberg 2004.

21 Logic for the empty domain is somewhat trickier than the remark in the text indicates; see Williamson 1999b.

22 See Williamson 1999a and Rayo and Williamson 2003 for this approach to the logic of unrestricted quantification. For reasons explained in the latter, the first-order quantification over interpretations in the text is a loose rendering of the higher-order quantification that is needed for an accurate metalogic of unrestricted quantification.

23 Cartwright 1994 takes this view.

24 See Friedman 1999, Williamson 1999a and Rayo and Williamson 2003. The result depends on a global choice assumption; Friedman discusses non-theorems of this logic that are logically true on some anti-choice assumptions. As an alternative to axioms of the form $\exists\, n$, the proponent of the constant account can use the following structural *Rule of Atomic Freedom*. Let $\Gamma \cup \{\gamma\}$ be a set of sentences, and $\Delta \cup \{\delta\}$ ($\delta \notin \Delta$) a set of atomic sentences each consisting of a non-logical predicate that does not occur in Γ and individual non-logical constants that do not occur in Γ such that $\Gamma, \Delta \vdash \delta$; then $\Gamma \vdash \gamma$ (Γ is

inconsistent; it entails everything). To see why this rule preserves validity on the constant account, suppose that all members of Γ are true on some interpretation. Then all members of $\Gamma \cup \Delta \cup \{\sim\delta\}$ are also true on some interpretation, for we can stipulate that each constant in $\Delta \cup \{\sim\delta\}$ denotes itself and that the extension of each *n*-place predicate in $\Delta \cup \{\sim\delta\}$ is the set of *n*-tuples of singular terms with which it is concatenated in Δ; thus every member of Δ is true, and δ is false because $\delta \notin \Delta$ (a simplified Henkin construction, which does not itself assume the existence of infinitely many things). Since the vocabulary of Γ is disjoint from that of $\Delta \cup \{\sim\delta\}$, it can be interpreted as originally; thus every member of Γ is true on the new interpretation too (this part of the argument would not work on the variable account, since there may be more constants in $\Delta \cup \{\sim\delta\}$ than things quantified over on the original interpretation). By contraposition, Atomic Freedom preserves validity. To see how to use Atomic Freedom to derive all sentences of the form $\exists\, n$, it suffices to look at the case $n = 3$. Let R be a triadic predicate, a, b and c distinct constants, $\Gamma = \{\sim\exists\, 3\}$, $\Delta = \{Raac, Raba, Rabb\}$ and $\delta = Rabc$. By ordinary first-order reasoning, $\sim\exists\, 3, Raac, Raba, Rabb \vdash Rabc$ (unless there are at least three things, the three constants cannot all have distinct denotations); therefore, by Atomic Freedom, $\sim\exists\, 3 \vdash \perp$, so $\vdash \exists\, 3$. Note that the rule of Atomic Freedom is formulated without reference to any particular logical constant (contrast axioms of the form $\exists\, n$). Another way to think of the constant account is therefore as freeing up the interpretation of atomic formulas.

25 See Cartwright 1994; the argument goes back to Kreisel 1967.

26 Tarski 1986 proposes the more precise criterion of invariance under all permutations of individuals in this spirit.

27 For the plural interpretation see several of the essays in Boolos 1998 (and 48-49 and 54 for brief remarks on the logic of identity). Williamson 2003a argues in favour of an interpretation that takes more seriously the idea of quantification into predicate position.

28 Thanks to Kit Fine, Øystein Linnebo, Agustín Rayo and Gabriel Uzquiano for helpful written comments on a draft of this chapter. A version of the material was presented at the conference at the Central European University in Budapest; thanks also to the audience there, especially the commentator Katalin Farkas, for useful discussion.

References

Bird, A. 1998. ‘Dispositions and antidotes’, *Philosophical Quarterly* 48: 227-234.

Boolos, G. 1998. *Logic, Logic, and Logic*. Cambridge, Mass.: Harvard University Press.

Cartwright, R. 1994. ‘Speaking of everything’, *Noûs* 28: 1-20.

Dummett, M.A.E. 1981. *Frege: Philosophy of Language*, 2nd ed. London: Duckworth.

Dummett, M.A.E. 1993. ‘Does quantification involve identity?’, in Dummett, *The Seas of Language*. Oxford: Clarendon Press.

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

Geach, P.T. 1967. ‘Identity’, *Review of Metaphysics* 21: 3-12.

Geach, P.T. 1972. *Logic Matters*. Oxford: Blackwell.

Geach, P.T. 1980. *Reference and Generality*, 3rd ed. Ithaca, NY: Cornell University Press.

Geach, P.T. 1991. ‘Replies’, in H.A. Lewis (ed.), *Peter Geach: Philosophical Encounters*. Dordrecht: Kluwer.

Glanzberg, M. 2004. ‘Quantification and realism’, *Philosophy and Phenomenological Research* 69: 541-572.

Harris, J.H. 1982. ‘What’s so logical about the “logical” axioms?’, *Studia Logica* 41: 159-171.

Hawthorne, J. 2003. ‘Identity’, in M. Loux and D. Zimmerman (eds.), *The Oxford Handbook of Metaphysics*. Oxford: Oxford University Press.

Henkin, L. 1950. ‘Completeness in the theory of types’, *Journal of Symbolic Logic* 15:

81-91.

Kaplan, D. 1989. ‘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.

Kreisel, G. 1967. ‘Informal rigour and completeness proofs’, in I. Lakatos (ed.), *Problems in the Philosophy of Mathematics*. Amsterdam: North-Holland.

Kripke, S. 1982. *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.

Lewis, D. 1983. ‘New work for a theory of universals’, *The Australasian Journal of Philosophy* 61: 343-377.

Lewis, D. 1997. ‘Finkish dispositions’, *Philosophical Quarterly* 47: 143-158.

Martin, C.B. 1994. ‘Dispositions and conditionals’, *Philosophical Quarterly* 44: 1-8.

Martin, C.B., and Heil, J. 1998. ‘Rules and powers’, *Philosophical Perspectives* 12: 283 -312.

McGee, V. 2000. ‘”Everything”’, in G. Sher and R. Tieszen (eds.), *Between Logic and Intuition*. Cambridge: Cambridge University Press.

Mumford, S. 1998. *Dispositions*. Oxford: Oxford University Press.

Noonan, H. 1997. ‘Relative identity’, in R. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*. Oxford: Blackwell.

Peacocke, C. 1976. ‘What is a logical constant?’, *Journal of Philosophy* 73: 221-240.

Quine, W.V. 1960. *Word and Object*. Cambridge, Mass.: MIT.

Quine, W.V. 1961. ‘Reply to Professor Marcus’ “Modalities and intensional languages”’, *Synthese* 13: 323-330.

Quine, W.V. 1966. *The Ways of Paradox and Other Essays*. New York: Random House.

Quine, W.V. 1976. 'Grades of discriminability', *Journal of Philosophy* 73: 113-116.

Rayo, A. 2003. 'When does "everything" mean *everything*?', *Analysis* 63: 100-106.

Rayo, A., and Williamson, T. 2003. 'A completeness theorem for unrestricted first-order Languages', in J.C. Beall (ed.), *Liars and Heaps*. Oxford: Clarendon Press.

Shapiro, S. 1991. *Foundations without Foundationalism: A Case for Second-order Logic*. Oxford: Clarendon Press.

Tarski, A. 1936. 'O pojciu wynikania logicznego', *Przegląd Filozoficzny* 39: 58-68. English translation by J. H. Woodger, 'On the concept of logical consequence', in Tarski, *Logic, Semantics, Metamathematics*, Indianapolis: Hackett, 2nd ed. 1983.

Tarski, A. 1986. 'What are logical notions?', *History and Philosophy of Logic* 7: 143 -154.

Wallace, J. 1964. *Philosophical Grammar*. Stanford University Ph.D. Published by University Microfilms Ltd, Ann Arbor, Michigan, 1969.

Westerståhl, D. 1989. 'Quantifiers in formal and natural languages', in D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, vol. 4. Dordrecht: Reidel.

Wiggins, D. 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.

Williamson, T. 1987/88. 'Equivocation and existence', *Proceedings of the Aristotelian Society* 87: 109-127.

Williamson, T. 1999a. 'Existence and contingency', *Proceedings of the Aristotelian Society*, sup. vol. 73: 181-203; reprinted with printer's errors corrected, 100: 321-343.

Williamson, T. 1999b. 'A note on truth, satisfaction and the empty domain', *Analysis* 59:

3-8.

Williamson, T. 2003a. 'Everything', *Philosophical Perspectives* 17: 415-465.

Williamson, T. 2003b. 'Understanding and inference', *Aristotelian Society* sup. vol. 77: 249-293.