Logique & Analyse **236** (2016), 367-387

CAN THE CUMULATIVE HIERARCHY BE CATEGORICALLY CHARACTERIZED?

Luca Incurvati

Abstract

Mathematical realists have long invoked the categoricity of axiomatizations of arithmetic and analysis to explain how we manage to fix the intended meaning of their respective vocabulary. Can this strategy be extended to set theory? Although traditional wisdom recommends a negative answer to this question, Vann McGee (1997) has offered a proof that purports to show otherwise. I argue that one of the two key assumptions on which the proof rests deprives McGee's result of the significance he and the realist want to attribute to it. I consider two strategies to deal with the problem — one of which is outlined by McGee himself (2000) — and argue that both of them fail. I end with some remarks on the prospects for mathematical realism in the light of my discussion.

Keywords: Categoricity, Set theory, McGee, Mathematical realism.

Mathematical realists take every mathematical statement to have a determinate truth-value, and consider that truth-value to be determined by the meaning of their words and the way the world is. And the world, they contend, is such as to make true all or most of the accepted mathematical axioms. To aid their favoured understanding of disciplines such as arithmetic and real analysis, realists have long invoked the categoricity of axiomatizations of these mathematical systems. Categoricity results, they claim, can be used to explain how we manage to fix the meaning of the vocabulary of the mathematical theories in question with enough precision that every sentence in their language has a determinate truth-value.

Of all mathematical disciplines, set theory is of special interest for the realist. For not only is set theory now a branch of mathematics in its own right; the fact that mathematics can be embedded in current standard set theory could be used by the realist to argue that a realist understanding of that discipline vindicates realism across the mathematical board.[1] However,

[1] Of course, the claim that mathematics can be embedded in standard set theory is not uncontroversial. Category theory, for instance, is sometimes taken to provide an example of a mathematical theory which cannot be reconstructed within set theory, although the jury is still out on the issue. For the purposes of this paper, I shall set these complications aside and grant that mathematics can be embedded in set theory.

doi: 10.2143/LEA.236.0.3186062.

traditional wisdom has it that the structure of pure sets cannot be categorically characterized, so that the realist appears to be left without an explanation of how we manage to affix a determinate truth-value to every mathematical sentence. Against this, Vann McGee (1997) has offered a proof that purports to show that the realist can have exactly what she wants: a true axiomatization of set theory which is categorical with respect to the pure sets. As a result, a new consensus seems to have arisen: the structure of the pure sets *can* be categorically characterized, and the realist's strategy in the case of arithmetic and analysis can be extended to the whole of mathematics.[2]

Unfortunately, the new consensus is wrong, and McGee's proof does not show what it purports to, or so I shall argue. I begin by explaining how the realist intends to use categoricity results to account for the fact that we succeed in fixing the intended interpretation of certain pieces of mathematical vocabulary. Then, I turn to McGee's proof, highlighting the two key assumptions on which it rests. Although discussion of the proof has focused on the first assumption, the real culprit is the second one. This assumption, I shall contend, deprives the result of the significance the realist wants to attribute to it.

1. Realism and Model Theory

The realist holds that all mathematical statements have a determinate truth-value determined by the meaning of her words together with the mathematical facts, and that these facts make true all or most of the accepted mathematical axioms. Hence — granting that the meaning of a mathematical expression is wholly dependent upon the thoughts and practices of speakers — the realist needs to explain how our use of mathematical vocabulary manages to fix the meaning of mathematical terms in such a way that determinacy of truth-value can obtain. She needs to explain, in other words, what it is that 'we think, do, or say that fixes the intended meaning of mathematical terms' (McGee 1997, 36) to such a degree. The proposed explanation is that mathematical terms have their meaning fixed by their role in mathematical *theories*: if they do, the realist claims, every mathematical statement must have a determinate truth-value.

To give substance to this explanation, the realist turns to model theory. She wants to use model theory to show that the relevant mathematical theories have a property which ensures that every sentence in their language has a determinate truth-value. The property in question is categoricity, where a theory is *categorical* iff all of its models are isomorphic. For it is easy to

[2] The new consensus includes the likes of Horsten (2007), Rayo and Uzquiano (1999, 315), Uzquiano (2002, 181) and Williamson (2003, 289–290, fn. 43).

see that a categorical theory is also *semantically complete*: all of its models make the same sentences in the language of the theory true. And semantic completeness, in turn, ensures determinacy of truth-value. For say that a sentence in the language of a theory is *determinately true* just in case it is true in all the models of the theory (and *determinately false* iff its negation is determinately true): assuming a bivalent semantics, a semantically complete theory is also one in which every sentence is either determinately true or determinately false. By equating determinate truth with truth in all models (and defining determinate falsity in terms of determinate truth), the realist can argue from categoricity to determinacy of truth-value.

Thus, the categoricity of a certain theory serves to explain how we manage to fix the meanings of our mathematical vocabulary in such a way that every statement of a certain area of mathematical discourse has a determinate truth-value: we learn the theory, and the theory singles out a particular class of models having the same structure — an isomorphism class. As a result, every sentence in the language of the theory must have a determinate truth-value.

It should be stressed, however, that the categoricity of a theory does not ensure that its axioms are *true*. For one thing, a categorical theory might have no model at all: determinacy of truth-value is compatible, for instance, with all universal sentences of the theory being vacuously true and all its existential sentences being false. For another, there are categorical theories which are usually taken to have models and to be false: consider, e.g., the theory consisting of the sentence stating that there are exactly seven things. Thus, categoricity cannot, by itself, vindicate the realist's conviction that every sentence of the theory has a determinate truth-value *in such a way that the axioms of the theory turn out to be true*. That said, categoricity might provide support for the realist's claim that every statement of a certain area of mathematical discourse has a determinate truth-value compatible with the truth of most of the accepted mathematical axioms. For, in the presence of categoricity, her reasons for regarding the axioms of a certain theory as true will also be reasons for thinking that every statement of that theory is either determinately true or determinately false.[3]

In what follows, we shall focus on the categoricity of a theory as a tool in the realist's hands to explain how her thoughts and practices manage to fix the intended interpretation of that theory.

[3] Walmsley (2002) argues that the realist's reasons for regarding the axioms of the theory as true, which would establish the existence of a model of the theory, would also suffice to establish the existence of its intended model. Thus, he concludes, categoricity is of no use in establishing that every statement in the language of a theory has a determinate truth-value compatible with the truth of most of the accepted mathematical axioms. For a rejoinder, see Paseau 2005.

2. McGee's Proof

Realists regard arithmetic as dealing, up to isomorphism, with a particular mathematical system — the natural number system — which determines the truth-value of every arithmetical statement. At first sight, however, the prospects for using model theory to explain how we manage to fix the intended meaning of arithmetical vocabulary look bleak. For, as is well known, first-order Peano Arithmetic (PA) — the standard axiomatization of number theory — is non-categorical. The realist, however, attributes this failure of categoricity to the fact that we have restricted attention to a *first-order* axiomatization of arithmetic, which, she claims, is ill-suited to account for our practices. Instead, she suggests considering PA_2, the second-order theory of arithmetic in which induction is formulated as a second-order axiom rather than as a schema. And, as expected, all models of this theory *are* isomorphic.

Having offered an explanation of how arithmetical sentences get their truth-values, the realist would like to do the same for every mathematical sentence. To this end, she turns to set theory, and here matters are more complicated. What the realist wants is a theory which categorically characterizes the universe of sets, which is strong enough that pure mathematics can be interpreted within it, and whose axioms she has reasons to regard as true. To this end, she typically focuses on set theories which are somehow based on the *iterative conception of set*.

According to this conception, sets are obtained by iterating the *set of* operation starting with the individuals, usually taken to be objects that are not sets and sometimes called *Urelemente*.[4] The result is a picture of the set-theoretic universe as a *cumulative hierarchy* divided into levels. This picture, the realist thinks, provides us with reasons to believe the axioms of ZFCU — i.e. first-order Zermelo-Fraenkel set theory (including the Axiom of Choice) with *Urelemente*[5] — a theory strong enough to interpret pure mathematics. But what about categoricity? Being a first-order theory with an infinite domain, ZFCU has nonisomorphic models. Reminiscent of the case of arithmetic, the realist decides to focus on $ZFCU_2$, the second-order theory obtained by replacing ZFCU's Replacement Schema with the corresponding second-order axiom. The appeal to second-order logic does

[4] In the philosophical literature, the iterative conception is often understood as involving the idea that sets are *metaphysically dependent* upon their members (see, e.g., Parsons 1977, Potter 2004 and Linnebo 2008). Incurvati (2012) defends an account of the iterative conception which dispenses with metaphysical dependence and takes sets simply to be the objects obtained by iterating the *set of* operation.

[5] Of course, it is not uncontroversial whether all of the ZFCU axioms are justified on the iterative conception. I am setting such issues aside in this paper. For discussion, see, e.g., Boolos 1971; 1989 and Paseau 2007.

go some way towards establishing what the realist needs. For let the *characteristic* of a model of set theory (representing its height) be the supremum of the ordinals represented in it. Then we have the following partial (*quasi*-) categoricity result:

THEOREM 1 (ZERMELO 1930) Any two models of ZFCU_2 with the same number of individuals and the same characteristic are isomorphic.[6]

In particular, any two models of ZFC_2 (the pure version of ZFCU_2) with the same characteristic are isomorphic. Can Zermelo's result be improved upon and the apparent contrast between arithmetic and set theory eliminated?

It turns out that if we stick with ZFCU_2 and assume that it is consistent with the existence of inaccessible cardinals, the result is best possible: by varying either (a) the universe of discourse of the models or (b) their individuals, we can construct models of ZFCU_2 which contain non-isomorphic copies of the pure sets. In case (a), we take a model of ZFCU_2 with inaccessibles and construct a new model by taking the range of the quantifiers to include only those sets whose cardinality in the original model is less than that of an inaccessible. (Note that this construction does not rely on the use of individuals, and so can be applied to models of ZFC_2.) In case (b), we construct the new model by letting the sets whose cardinality is equal or more than that of an inaccessible be individuals in the new model. Either way, this means, from the realist viewpoint, that the sentence stating that there exists an inaccessible cardinal will be true in some models of set theory and false in others, and hence will not have a determinate truth-value. The problem cannot be overcome simply by adjoining this sentence, since the same problem will arise for the resulting theory concerning the sentence stating, say, that there are measurable cardinals, at least if this theory is consistent with it. If the realist wants to provide a categorical characterization of the structure of the pure sets, she needs a different strategy.

McGee's idea is to rule out construction (a) by requiring ZFCU_2's quantifiers to range over everything and construction (b) by laying down an axiom stating (on the intended interpretation) that there is a set of all individuals (non-sets):

Urelemente. $\exists x(\mathrm{Set}(x) \wedge \forall y(\neg \mathrm{Set}(y) \rightarrow y \in x))$.

He then goes on to show that the resulting theory is categorical as far as the pure sets are concerned. In other words, we have the following:[7]

[6] Ernst Zermelo's original result concerns a theory without the Axiom of Infinity and having models of characteristic ω.

[7] In the same paper, McGee also attempts to show that the result continues to hold if we replace the second-order Axiom of Replacement with the corresponding open-ended schema. (More on open-endedness below.) For doubts about the significance of this strategy, see

THEOREM 2 (MCGEE 1997) Any two models of $ZFCU_2$ + **Urelemente** in which the quantifiers range over everything have isomorphic pure sets.[8]

In broad outline, the proof proceeds as follows. One first shows that $ZFCU_2$ + **Urelemente** implies the *Completeness Principle*, which states, on the intended interpretation, that the pure sets cannot be embedded into the "pure sets" of any other model of $ZFCU_2$. Then, one proves that this principle is equivalent, over the axioms of $ZFCU_2$, to the *Maximality Principle*, which asserts, again on the intended interpretation, that any model of $ZFCU_2$ is such that *its* pure sets are isomorphic to the pure sets. Finally, one shows that any two models of $ZFCU_2$ + the Maximality Principle (and hence any two models of $ZFCU_2$ + **Urelemente**) with the same universe of discourse have isomorphic pure sets. In particular, McGee observes, any two such models whose variables range over everything have isomorphic pure sets. QED. Notice that whilst the assumption that the non-sets form a set is used at the outset to derive the Completeness Principle, the assumption that the quantifiers range over everything is not needed until this very last step.

Now McGee's result, as a piece of mathematics, is beyond dispute. What is disputable is whether, notwithstanding its additional assumptions, the result retains the significance the realist wants to attribute to categoricity results. After all, we already have examples of set theories extending ZFC_2 ($ZFCU_2$ without *Urelemente*) which *are* categorical, such as the theory obtained by adding to it the negation of the sentence stating that there exists an inaccessible cardinal. The categoricity of these theories depends on assumptions which the realist, by her own admission, is not entitled to make and which jeopardize the significance of the result. Is the situation different in the case of McGee's proof?

Pedersen and Rossberg 2010. Here we shall set the issue aside and grant that the realist understands second-order quantification.

[8] A note on terminology. In standard model theory, an *interpretation* for a language $\mathcal{L}$ is an ordered pair $\mathfrak{A} = \langle \mathcal{D}, \mathcal{I} \rangle$, where $\mathcal{D}$ is the domain of interpretation and $\mathcal{I}$ is an interpretation function mapping the non-logical symbols of $\mathcal{L}$ to appropriate relations, functions and constants in $\mathcal{D}$. A *model* of a theory is then an interpretation which makes every sentence of the theory true. However, since McGee's result concerns a theory whose quantifiers range over everything, the models his result is about are ones whose domains contain all sets. Since realists typically accept that there is no set of all sets, these models cannot be sets. To circumvent the problem, realists have explored the possibility of developing model theory with the help of higher-order logic by, roughly, taking models to be not sets but rather to be given by the objects which a monadic second-order variable is true of (see, e.g., Rayo and Williamson 2003). Without committing ourselves to the viability of this proposal, we shall follow McGee (1997, 49–50) and assume it in the remainder of the paper, unless context indicates otherwise. Thus, although we shall often talk of models as if they were set-theoretic objects, it should be kept in mind that, officially, they will be variable assignments.

3. ‘Every Set’ and ‘Every Thing’

As anticipated, discussion of the proof has tended to focus on the assumption that the non-sets form a set. Philosophers have tried to put pressure on this assumption by putting forward arguments to the effect that the non-sets are bound to be too many. For instance, Ted Sider (2009, 256) asks us to suppose that we admit properties in our ontology and that, for each set *x*, there is a distinct property of having the members that *x* has. Assuming that properties are not sets, a simple cardinality argument then seems to show that there cannot be a set of all non-sets. There are, however, a number of things that the realist could say in response to this sort of worry. She could question the assumption that properties are the kind of things that fall within the range of the first-order quantifier. Or she could challenge the assumption that there is a distinct property for each set *x*. For instance, she might claim that the argument simply shows that properties should be constructed as *natural* properties (as in Lewis 1983), and argue that it is not the case that for each set *x* there is a natural property of having the members that *x* has.

We will not pursue these issues here, however. For the major problem with McGee’s result lies in the assumption that the quantifiers range over everything, an assumption which is often mentioned just in passing. Indeed, it is not mentioned at all in the Stanford Encyclopedia of Philosophy entry on the philosophy of mathematics (Horsten 2007), where it is suggested that ‘McGee has shown that if we consider set theory with Urelements, then the theory is *fully* categorical with respect to pure sets if we assume that there are only set-many Urelements’. Yet, the assumption that the quantifiers range over everything is crucial to McGee’s proof. For, as we have seen, it is only when the quantifiers of ZFCU_2 + **Urelemente** are taken as ranging over everything that we can conclude that any two of its models have isomorphic pure sets.

To see what the problem with the assumption is, we first need to clarify what it amounts to. When people argue over whether it is possible to quantify over everything, what they are sometimes arguing about is whether there are *universal* models of set theory — models in which the entire universe participates.

However, assuming that such models exist is not enough for McGee’s purposes. For it is true that some philosophers have denied the possibility of quantifying over an all-inclusive domain which includes all sets, typically on the grounds that the concept of set is indefinitely extensible (see, e.g., Lear 1977 and Parsons 1974). And it is also true that the resulting view of set theory fits quite well with the one sketched by Zermelo at the end of the very same paper where he proves his quasi-categoricity result (Theorem 1 above). His idea seems to be that set theory does not deal with an intended domain of all sets, but is rather about an open-ended and well-ordered sequence of models, each of which is strictly larger than the preceding one. Inspired by this, one could take the set-theoretic quantifiers as never ranging over all sets, but only over the sets in one of the models in the 'unbounded double-series of essentially different set-theoretic models, in each of which the whole classical theory is expressed' (1930, § 5). This might be because one thinks that there is no domain of absolutely all sets or, more modestly, because one takes the set-theoretic quantifiers to be always restricted, as a matter of fact, to range over the sets in one of the models in the sequence proposed by Zermelo. Either way, the resulting model theory would refrain from countenancing models whose domain contains all sets and, *a fortiori*, universal models.

But one can deny that a theory such as $ZFCU_2$ characterizes the structure of pure sets whilst granting that it has universal models and hence models whose domain contains all sets. In denying that a theory such as $ZFCU_2$ characterizes the structure of pure sets, one might simply be denying that this theory determines whether 'every set' includes, say, inaccessible cardinals or not. Against this, McGee is assuming that the theory's quantifiers *do* range over everything, so that the theory *only* has universal models.

Making this assumption does deliver the result that all models of the theory have isomorphic pure sets. However, it also means that the result will show that the sentence stating that there is an inaccessible cardinal has a determinate truth-value only if it is presupposed that our thoughts and practices determine whether 'everything' includes inaccessible cardinals or not. But the usual reason (see, e.g., Field 1994, 394 and Isaacson 2009, §6) for denying that second-order set theory is capable of fixing the meaning of 'every set' to the extent that the statement that there exists an inaccessible cardinal has a determinate truth-value is precisely that this theory cannot fix the meaning of 'every *thing*' to such an extent. Thus, rather than providing an answer to the question of what it is that we think, do, or say that enables us to single out the intended model, McGee's result seems to just assume there to be one.

The problem can be put as follows. The point of denying that $ZFCU_2$ + **Urelemente** characterizes the structure of the pure sets was to highlight that it is compatible with the theory that (i) the smaller model has left out things that we intended to quantify over but also that (ii) the larger one has included things that we did *not* intend to quantify over. Taking the quantifiers of $ZFCU_2$ + **Urelemente** to range over everything is tantamount to assuming that the theory only has universal models and hence that (i) is the case. But this does not explain *what in our practices* ensures that we did intend to quantify over the things left out by the smaller model. And this is precisely what McGee's result was meant to be doing.

One might reply that we do have reasons for restricting attention to universal models. For

> [m]odels in which we clip off the construction of the universe at the first inaccessible […] are unintended models […] because the intended models of set theory are ones in which the universe of pure sets is as large as possible. It isn't possible to go on from an intended model of set theory to any other model with a still larger universe of pure sets. (McGee 1997, 53–54)

But this reply would just compound the original mistake. For the issue is not what the intended model[9] of set theory is — it can be granted, for present purposes, that it is the one in which the universe of sets is as large as possible. The issue, rather, is *whether our mastery of second-order set theory enables us to single it out*. And *this* issue is not settled in the positive simply by stipulating that the theory's quantifiers range over everything — by stipulating, that is, that the theory only has universal models. A categoricity result is of some significance for the realist as long as it enables her to say in virtue of what the intended model is intended and the unintended models are unintended: in so doing, the result puts her in a position to say what it is that we think, do, or say that enables us to single out the intended model. McGee's result, on the other hand, does not put her in such a position, since it simply assumes that models that are not universal are unintended.

To make the point more vivid, consider the following analogy. Almost everyone agrees that the intended model of arithmetic is characterized by a *minimality* condition — it is the smallest model of PA, included in any other model as an initial segment. Thus, PA fails to be categorical because it does not determine whether the series of natural numbers includes elements after all the standard natural numbers. This raises the question: in virtue of what can we say that a model including such elements has included things which we did not intend to quantify over when quantifying over 'all numbers' rather than the model not including such elements has left out things which we *did intend* to quantify over? It would be highly unsatisfactory to answer the question by considering a theory of arithmetic whose quantifiers have minimal range in the sense that if two models have different universes of discourse, then the one with the larger universe is one in which the variables do not have minimal range. For this amounts to stipulating that the models that are not minimal are unintended.

This is why the use of second-order logic in the case of PA_2 is important: it enables us to formulate a second-order version of the Induction Axiom which puts us in a position to say in virtue of what the minimal model, among all the models of PA, is the intended one — it is the model for which induction applies with respect to all subsets of the domain, and not only those definable by conditions expressible in the language of arithmetic. And

[9] McGee talks of 'intended model*s*' because the best we can hope for is to single out models up to isomorphism. I have chosen to follow what is perhaps standard usage and to talk of 'intended model', meaning *intended model up to isomorphism*.

this is what makes the categoricity result for PA_2 of some significance: it shows that if the realist understands second-order quantification and accepts basic truths about the numbers, she can single out the intended model of arithmetic.

Of course, the significance of the answer provided by a categoricity result may (and usually will) be challenged. In the case of arithmetic, for instance, it will typically be asked in virtue of what the second-order quantifiers are to be interpreted as ranging over all subsets of the first-order domain, and not over a collection of not necessarily all of its subsets (in which case the categoricity proof does not go through).[10] The point, however, is that the categoricity result for PA_2 at least shows that if besides first-order quantification — i.e. quantification over *all objects in the domain* — we also understand second-order quantification — taken to be quantification over *all subsets of the domain* — then there is something we do, say, or think that enables us to single out the intended model of arithmetic. By contrast, even granting that the realist understands first- and second-order quantification is not enough for McGee's proof to go through. Because of this, McGee's result does not even offer the resources to provide an answer that can itself be challenged. Or at least, it does not offer such resources unless we have *independent* reasons to assume that our practices determine whether 'everything' includes everything in the universe. Unsurprisingly, this is what McGee sets out to show in a later paper.

4. The Meaning of 'Everything': Open-Endedness

McGee (2000) argues that there is something in our practices which enables us to fix the meaning of 'everything' to the extent that it determinately includes the entire universe. If successful, the argument would provide the required independent reason for restricting attention to universal models from the outset, thereby showing that McGee's result has indeed the significance he wants to attribute to it.

The argument goes as follows (see McGee 2000, 68–69). Suppose that our first-order quantifiers do not range over everything that there is in reality, so that there exists some object *a* lying outside their range. And suppose,

[10] Similar challenges will arise, *mutatis mutandis*, if one gives a different interpretation of the second-order quantifiers, for instance as ranging over pluralities. This is important because the interpretation of the second-order quantifiers as ranging over subsets of the first-order domain will not be available if, as in the case of McGee's proof, among the first-order domains there are some which contain everything. For otherwise we could use Russell-style reasoning to define a set — the subset of all things in the universe which do not belong to themselves — which cannot, on pain of contradiction, belong to the universe, thereby invalidating the assumption that the first-order domain did include everything there is. See McGee 1997, 46–47 for details.

further, that we can name anything we like by an individual constant. Then, we can extend the language with a new individual constant c and let c denote a. But if we take a predicate F which picks out all and only the things in the range of our first-order quantifiers,[11] $\forall x Fx$ will be true and Fc will be false. And this violates the *open-ended* rule of $\forall$-Elimination, according to which one can infer $\phi(\tau)$ from $\{\forall x \phi(x)\}$ not only in the language of the theory, *but in any extension of the language*. Thus, the argument concludes, since our acceptance of the inference rules for our logical vocabulary is indeed open-ended, our quantifiers range over everything.

To see what is wrong with the argument, we need to recall what is at issue here: the realist is presented with two models — a universal model $\mathfrak{U}$ and less comprehensive model $\mathfrak{V}$ — and has to say what in our practices rules out $\mathfrak{V}$ as unintended. McGee's proposed answer is that somebody whose quantifiers only range over everything in $\mathfrak{V}$ will not use the rule of $\forall$-Elimination in an open-ended fashion. This conclusion, however, follows only if the person quantifying over everything in $\mathfrak{V}$ can name an object a which is not in the domain of $\mathfrak{V}$ with a constant c. For if the person quantifying over everything in $\mathfrak{V}$ can only name objects in the domain of $\mathfrak{V}$, her use of the inference rules will indeed conform to the open-ended rule of $\forall$-Elimination.

What this shows is that what the inference rules do, at best, is fix the meaning of the quantifiers *to the extent that they range over everything that we can name*. But they do not ensure that the quantifiers range over absolutely everything *unless we assume that we can name everything in the universe*. But the problem is that the assumption that we can name everything in the universe begs the question. The question, to repeat, is how we succeed in singling out the all-inclusive domain as our domain of quantification and in ruling out the less comprehensive ones. McGee's argument simply assumes that the all-inclusive domain is already our domain of discourse, since we can pick out any object belonging to it with an individual constant.[12] In other words, the most McGee has established is that if we can name something, we can quantify over it: the desired conclusion follows only if we assume that we can name everything in the universe. And to assume this is to assume much of what had to be shown.

McGee recognizes that the inference rules, *by themselves*, are unable to ensure that the quantifiers range over everything, and that they are only

[11] One such predicate is $\exists y\, y = x$. For this reason, McGee (2006, 187) presents a version of the argument in which the counterexample to the rule of $\forall$-Elimination is the inference from $\{\forall x \exists y\, y = x\}$ to $\exists y\, y = c$. I have kept the presentation of the argument closer to the original one here.

[12] Note that the complaint is not that it is question-begging to assume that there is an all-inclusive domain. The complaint is that it is question-begging to assume that that domain is already our domain of discourse.

capable of doing so in conjunction with the assumption that we can name everything in the universe. In the same passage, he also seems to attempt to offer an argument for the latter assumption:

> The rules of inference do not determine the range of quantification. What they ensure is that the domain of quantification in a given context includes everything that can be named within that context. This includes even contexts in which there are no restrictions on what can be named. In such contexts, the quantifiers range over everything. (McGee 2000, 69)

Thus, McGee argues, we can name everything in the universe because we can do so in contexts where there are no restrictions on what can be named. Clearly, 'restriction' here must mean *restriction of any sort*. For it is perfectly compatible with the claim that there are, say, no *syntactic* restrictions on what can be named that there are restrictions of some other kind on what can be named. Thus, the claim has to be that we can name everything because we can do so in contexts where there are no restrictions *of any sort* on what can be named. This is clearly the case, but the problem, again, is that to assume that there are indeed such contexts is to assume much of what is at stake.

To sum up, McGee's argument shows that the quantifiers range over everything — and not simply over everything in the given domain of discourse — only on the assumption that we can name everything in the universe. But this assumption begs the question, and so does McGee's attempt to argue for it.

5. The Meaning of 'Everything': Logicality

It might be insisted that even if McGee's argument fails, it is nonetheless legitimate for the realist to restrict attention to universal interpretations. The reason, it might be argued, is that the logical character of the quantifiers ensures that they range over the entire universe — a view famously defended by Timothy Williamson (2000). Indeed, Stewart Shapiro (2003, 483) explicitly says that 'McGee's philosophical interpretation of his result [i.e. his claim that his result establishes the determinacy of mathematical language] depends on the presence of a quantifier which, as a matter of logic, ranges over everything (in all interpretations)'.

Let us spell out the idea in more detail. It is a familiar fact that the standard recursive definition of truth in an interpretation assigns the connectives their intended meaning (and the same is true when interpretations are taken to be variable assignments). So, for instance, $\phi \wedge \psi$ is declared to be true in an interpretation $\mathfrak{A}$ iff ϕ is true in $\mathfrak{A}$ and ψ is true in $\mathfrak{A}$, and $\neg\phi$ is declared to be true in $\mathfrak{A}$ iff ϕ is not true in $\mathfrak{A}$. As a result, $\neg(\phi \wedge \neg\phi)$ is logically true. This way of proceeding is usually justified by saying that logical truth is truth in all interpretations in which the logical vocabulary — which includes the connectives — receives its intended meaning.

What about the quantifiers? The standard recursive definition of truth tells us that $\forall x\phi$ is true in $\mathfrak{A}$ under a certain assignment α of values to variables iff ϕ is true for *every* assignment β which differs from α at most in the value it assigns to x. Thus, it is usually claimed, the quantifiers retain their intended meaning across interpretations, as they should. According to the view under consideration, however, for the quantifiers to retain their intended meaning, it is not enough that the truth-conditions of sentences containing them be formulated in terms of every assignment; it must also be the case that their range encompasses the entire universe. Thus, when characterizing logical truth, we must restrict attention to interpretations in which the quantifiers range over everything. As a result, '$\exists !^k x\, x = x$' ('There exist exactly k things'), where k is the cardinality of the universe, is a logical truth. If such a view is correct, a defender of the significance of McGee's result might therefore suggest, the realist is justified in restricting attention to universal interpretations.

The idea that certain unintended interpretations can be ruled out by widening the scope of what is settled as a matter of logic is not new. Alfred Tarski (1958) already suggested ruling out the countable models of ZFC by treating '$\in$' as part of the logical vocabulary, thereby restricting attention to interpretations in which '$\in$' has its intended meaning. Indeed, as Paul Benacerraf (1985, 107) noticed, for Tarski's proposal to work it is not enough to treat '$\in$' as part of the logical vocabulary; it must also be the case that the quantifier domain contains enough sets so that the Downward Löwenheim-Skolem Theorem does not go through, which is, effectively, to take the quantifier as ranging over the entire universe in all interpretations. And, Benacerraf claimed, this would make '$\in$' sufficiently different from other pieces of logical vocabulary to make it doubtful that it should count as one at all.

Benacerraf's reply is unsatisfactory because the issue is precisely whether the realist can legitimately restrict attention to interpretations in which the quantifiers range over everything. To simply point out that standard pieces of logical vocabulary do not make demands on the quantifier's range seems just to refuse to engage with the proposal the realist is advancing.

So let us grant, for the sake of argument, that which individuals the truth-value of a quantified sentence depends upon is a logical matter, as Williamson and others contend. It follows that, *for the purpose of delimiting the class of logical truths*, one should restrict attention to universal interpretations, at least if one wishes to hold on to the standard account of logical truth. It does not follow, though, that it is legitimate to restrict attention to such interpretations *when trying to provide a categoricity result which is to serve the realist's purposes*. For what we are using model theory for in the latter case is not to delimit the class of logical truths but rather to offer an explanation of how we manage to fix the meaning of mathematical vocabulary to such an extent that every mathematical statement has a determinate truth-value. In so doing, we standardly restrict attention to interpretations in which the logical vocabulary receives its intended meaning.

But this is because it is usually granted, for the sake of argument, that the meaning of *what is typically taken to be logical vocabulary* is determinate.[13] However, if we are willing to count more as logical, then we can no longer assume that a piece of vocabulary keeps its intended meaning across interpretations just in virtue of being logical: a logicist about arithmetic who is also a realist is not exempted from the task of showing how our practices enable us to single out the intended model of arithmetic. Thus, even if it is conceded that '$\in$' is a logical constant and that it stands for membership only if the domain contains enough sets, the realist still has to explain how our thoughts and practices ensure that '$\in$' does stand for membership in this sense and not for some other relation. Similarly, even if it is granted that the quantifiers receive their intended meaning only if their range consists of everything, the realist still needs to say what it is that we think, do, or say which makes sure that they do receive their intended meaning. Hence, taking the range of the quantifiers to be determined by logic does not solve the problem the realist set out to solve, but simply relocates it.

6. Concluding Remarks: The Prospects for Mathematical Realism

Despite being a second-order theory, $ZFCU_2$ has models with non-isomorphic pure sets. At least until recently, the received view was that this is due to the nature of the mathematical system it attempts to describe: any theory based on the cumulative hierarchy will have models with non-isomorphic pure sets because, no matter how many axioms concerning its height this theory comprises, it will not be able to determine whether the hierarchy extends just as far as its axioms force it to or further up.

Against this received view, McGee suggested proving categoricity with respect to the pure sets by adding two novel assumptions to $ZFCU_2$, namely the axiom **Urelemente** and the assumption that the quantifiers range over

[13] Indeed, Carnap (1943) famously pointed out that there are non-standard interpretations of the standard connectives compatible with the standard rules for classical propositional logic. Thus, even the assumption that what is typically taken to be logical vocabulary has a determinate meaning is not as trivial as it might seem. McGee (2000) is, of course, well aware of the issue, and suggests dealing with it by appealing to a result of Harris (1982). This makes me suspect that he would agree with our claim that we cannot just *assume* a certain piece of vocabulary to have a determinate meaning just in virtue of its being logical.

everything. But this latter assumption, I have argued, jeopardizes the significance of McGee's result: the realist cannot help herself to the result to explain how our thoughts and practices succeed in fixing the meaning of mathematical vocabulary with sufficient precision to ensure that every mathematical sentence has a determinate truth-value. What is more, a theory which categorically characterizes the universe of pure sets in a way that does enable her to provide the required explanation does not appear to be forthcoming — in fact, it is not if the old received view as to why ZFC_2 fails to be fully categorical is correct.

Could the realist resort to ways other than categoricity to characterize structures up to isomorphism? The only available alternative, related to categoricity in a manner yet to be fully clarified (see Awodey and Reck 2002, 91–93), is provided by the category-theoretic notion of a universal mapping property, since any two structures satisfying a certain universal mapping property are isomorphic. Thus, Lawvere (1969) provided a characterization of the natural number structure as the free successor algebra, and Joyal and Moerdijk (1991; 1995) applied this approach to set theory, initiating *algebraic set theory*. Their approach is to consider a suitable category $\mathcal{C}$ together with a designated class of arrows which are said to be small and satisfy certain natural conditions. In such a category, they identify certain algebras which they call ZF-algebras. One of their main results is then a characterization of the cumulative hierarchy of pure sets as the free ZF-algebra on $\emptyset$. Significantly, this characterization is compatible and in fact supports the main claims of this paper, since it is relative to the background category $\mathcal{C}$ and its class of small maps. And indeed *the* cumulative hierarchy in the sense at stake in McGee's paper is obtained only when $\mathcal{C}$ consists of *all* the classes. Thus, in order to use the characterization of the cumulative hierarchy provided by Joyal and Moerdijk, the realist must assume that the meaning of 'all classes' is determinate, just as in order to use McGee's Theorem she had to assume that the meaning of 'everything' is.[14]

Our conclusions, thus, seem to leave the realist with two options. The first option would be for the realist to amend her view so as to allow for any amount of indeterminacy compatible with the quasi-categoricity of ZFC_2. Rather than holding that every mathematical sentence has a determinate truth-value, the realist might more modestly take to have a determinate truth-value only those sentences in the language of pure set theory which concern the hierarchy up to the first inaccessible. Call this position *modest (mathematical) realism*.

The idea is that whilst, say, the Continuum Hypothesis has a determinate truth-value — being either true in all models of ZFC_2 or false in all of them

[14] For more on algebraic set theories, see http://www.phil.cmu.edu/projects/ast/ and references contained therein.

— the sentence stating that there exists an inaccessible cardinal is, if consistent with ZFC_2, neither determinately true nor determinately false — being true in some models of the theory and false in others. And, the thought goes, although it is indeed indeterminate whether 'every set' includes inaccessible cardinals or not, this level of indeterminacy is one that it is possible to live with. In particular, the modest realist still considers most, if not all sentences of *ordinary* mathematics to have a determinate truth-value. (In fact, those sentences typically deal at most with the first few infinite levels of the hierarchy.) And, crucially, Zermelo's result would be sufficient to explain how our thoughts and practices succeed in fixing the meaning of mathematical vocabulary to the extent that every sentence the modest realist holds to have a determinate truth-value does have a determinate truth-value.

The modest realist, then, holds on to the idea that the meanings of mathematical terms are fixed by their role in her favourite axiomatic set theory, *viz.* ZFC_2, but acknowledges that these meanings fail to affix a determinate truth-value to sentences such as 'There exists an inaccessible cardinal' (assuming, again, that such sentences are consistent with ZFC_2). This, it is worth stressing, means that in her view *there is no fact of the matter* as to whether there are inaccessibles. For if what fixes the meaning of mathematical terms is their role in ZFC_2, and if some models of this theory contain inaccessibles and some don't, then there is nothing to make interpretations on which 'There exists an inaccessible cardinal' is false unintended.[15] True, the modest realist could avoid this verdict by slightly modifying her view and taking her favourite axiomatic set theory to include 'There exists an inaccessible cardinal' as an axiom. But, as emphasized throughout the paper, the same consequence would arise, *mutatis mutandis*, with regards to any stronger large cardinal axioms which can be consistently added to this new axiomatic set theory.

Thus, whilst endorsing modesty excuses the realist from explaining what it is that we do, think, or say which determines the truth-value of sentences such as 'There exists an inaccessible cardinal', it also results in a significant retreat from her original position. To appreciate how significant the retreat is, consider the search for new large cardinal axioms to settle questions — such as whether there exist non-constructible subsets of ω — which although semantically decided by ZFC_2, are proof-theoretically independent of it. This enterprise represents an important part of current set-theoretic practice. But what is the modest realist going to make of it? She cannot say that what set theorists are doing is trying to find axioms that hold in the universe of sets. For, in her view, there is no fact of the matter as to whether

[15] For an analogous point in the context of a discussion of the Axiom of Constructibility, see Putnam 1980, 469.

what these axioms say is the case (assuming their consistency with ZFC_2). So what *is* she going to say?

She might, of course, insist that the enterprise is ill-conceived. Besides the clash with set-theoretic practice, however, this suggestion is in tension with the modest realist's view that statements such as 'There exists a non-constructible subset of ω' do have a determinate truth-value. For if one holds a certain statement to have a determinate truth-value, one might devote some effort to trying to discover what this truth-value is. But by taking the debate over large cardinal axioms to be ill-conceived, the modest realist has prevented herself from resorting to these axioms to do so. True, it is open to her to say that one can use novel *proof-theoretic* principles, rather than set-theoretic ones, to find out the truth-value of those statements that are semantically, but not proof-theoretically decided by ZFC_2. But this would involve, again, a substantial departure from current set-theoretic practice (see Incurvati 2008, 89).

What the modest realist is likely to do, instead, is try to make sense of the search for new axioms by saying that these are simply tools to see what goes on at lower levels of the hierarchy. Take, for instance, the statement that there exists a measurable cardinal. If consistent with ZFC_2, this statement has, according to the modest realist, no determinate truth-value, being true in some models and false in others. But the fact that the statement implies 'There exists a nonconstructible subset of ω' tells us that the latter is true in all ZFC_2 models with measurable cardinals, and hence in all ZFC_2 models *tout court*. For 'There are non-constructible subsets of ω' is either true in all ZFC_2 models or false in all of them.

Thus, whilst only sentences concerning the hierarchy up to the first inaccessible have a determinate truth-value, the modest realist may claim, it is *useful* to resort to large cardinal axioms to find out what these truth-values are. Large cardinal axioms would then have a role which is remarkably similar to the one mathematics has in Field's (1980) nominalist philosophy of mathematics. In his view, as clarified in a debate with Shapiro (1983), 'mathematics is useful because it is often easier to see that a nominalistic claim follows from a nominalistic theory plus mathematics than to see that it follows from the nominalistic theory alone' (Field 1985, 241), where the notion of consequence in play here is semantic and second-order. According to the modest realist's suggestion, large cardinal axioms are useful because it is often easier to see that a sentence concerning the hierarchy up to the first inaccessible (semantically) follows from ZFC_2 plus large cardinal axioms than to see that it (semantically) follows from ZFC_2 alone.

Hence, the realist might indeed succeed in accounting for the search for large cardinal axioms. But she would do so by taking them to have a mere *instrumental* role, namely that of enabling us to discover more about what is already semantically decided by the ZFC_2 axioms. This brings to light the extent to which the realist, by embracing modesty, has backed away from her original position: short of dismissing the search for large cardinal axioms as ill-conceived, she is now forced to relegate these axioms — whose investigation is central to much research in set theory — to the role that mathematics has on a nominalist view such as Field's. The main challenge for the realist, if she chooses the first option, seems therefore that of explaining in what sense her position still counts as a form of realism about set theory at all.

The second option — diametrically opposed to the first one — would be for the realist to say that whilst set-theoretic terms *do* get their intended meaning, this meaning is *not* fixed in its entirety by their role in axiomatic set theory. A few remarks expressing sympathy for a position along these lines are offered by Benacerraf at the end of his 1985 paper. For instance, he writes:

> Even if we agree [...] that the determinants of mathematical meaning must lie somewhere in our *use* of mathematical language, why think use to be captured (or capturable) by axioms (we *do* argue about new axioms)?' (Benacerraf 1985, 111).

Clearly, though, going for the second option does not exempt the realist from the task of offering an explanation of how our thoughts and practices — now taken to resist axiomatizability — fix the intended meaning of mathematical vocabulary. What could this explanation look like?

A well-known suggestion is that we are capable of determinately quantifying over all sets through our grasp of the intuitive model for set theory — the cumulative hierarchy — together with the very intention to quantify over all sets. But for this suggestion to count as an explanation, an account would have to be provided of what our grasp of the cumulative hierarchy comes to — an account which makes it clear how this grasp makes deviant interpretations of axiomatic set theory unintended.

Hilary Putnam (1980) famously expressed pessimism about the possibility of undertaking such a task without postulating mysterious faculties which are not amenable to naturalistic treatment. The appeal to capacities like that of grasping intuitive models, he said, is 'both unhelpful as epistemology and unpersuasive as science' (1980, 471). And Dummett (1991), for his part, insists that intuitive models are grasped through language — rather than, say, direct intellectual apprehension — and, as a result, 'have no more content, and are no more definite, than the verbal or symbolic descriptions by means of which they may be communicated' (1991, 311). Thus, he concludes, it is doubtful whether they can be taken as conveying 'a conception of a domain of quantification sufficiently definite to warrant attributing to statements involving quantification determinate truth-values' (1991, 320). If the realist wants to pursue the second option by appealing to intuitive models, therefore, she will have to show Putnam's and Dummett's pessimism to be unwarranted. If, on the other hand, she settles for the second option but does not want to resort to intuitive models, she will have to offer a different account of the way our thoughts and practices, whilst possibly transcending what is explicitly capturable by axioms, determine the intended meaning of mathematical terms.

A comprehensive assessment of the merits of the two options the realist is faced with — and of the value of using full second-order logic (or its schematic variant) to explain how mathematical language can be determinate — will have to await another occasion. One thing seems clear though: although we can keep calling McGee's result a 'categoricity result' if we want, we should bear in mind that it does not have the significance the realist typically wants to attribute to results of this kind.

Acknowledgements

Many thanks to Tim Button, Hannes Leitgeb, Alex Oliver, Michael Potter, Florian Steinberger, Rob Trueman and two anonymous referees for comments and feedback on earlier versions of this material.

References

[1] Awodey, S. and Reck, E. H. (2002), 'Completeness and categoricity, part II: Twentieth century metalogic to twenty-first-century semantics', *History and Philosophy of Logic* **23**, 77–94.

[2] Benacerraf, P. (1985), 'Skolem and the skeptic', *Proceedings of the Aristotelian Society* **59**, 85–115.

[3] Benacerraf, P. and Putnam, H. (1983), *Philosophy of Mathematics. Selected Readings*, 2nd edn, Cambridge University Press, Cambridge.

[4] Boolos, G. (1971), 'The iterative conception of set', *Journal of Philosophy* **68**, 215–231. Reprinted in Boolos 1998, 13–29.

[5] Boolos, G. (1989), 'Iteration again', *Philosophical Topics* **17**, 5–21. Reprinted in Boolos 1998, 88–104.

[6] Boolos, G. (1998), *Logic, Logic, and Logic*, Harvard University Press, Cambridge, Massachusetts.

[7] Carnap, R. (1943), *Formalization of Logic*, Harvard University Press, Cambridge, Massachusetts.

[8] Dummett, M. (1991), *Frege. Philosophy of Mathematics*, Harvard University Press, Cambridge, Massachusetts.

[9] Ewald, W. B., ed. (1996), *From Kant to Hilbert: A Source Book in the Foundations of Mathematics*, Oxford University Press, Oxford.

[10] Field, H. (1980), *Science Without Numbers*, Oxford University Press, Oxford.

[11] Field, H. (1985), 'On conservativeness and incompleteness', *Journal of Philosophy* **82**, 239–260.

[12] Field, H. (1994), 'Are our logical and mathematical concepts highly indeterminate?', *in* P. A. French, T. E. Uehling, and H. K. Wettstein, eds, *Midwest Studies in Philosophy*, Vol. 19, pp. 391–429.

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

[14] Horsten, L. (2007), 'Philosophy of mathematics', *in* E. N. Zalta, ed., *Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*. Available at http://plato.stanford.edu/archives/fall2008/entries/philosophy-mathematics/.

[15] Incurvati, L. (2008), 'On adopting Kripke semantics in set theory', *Review of Symbolic Logic* **1**, 81–96.

[16] Incurvati, L. (2012), 'How to be a minimalist about sets', *Philosophical Studies* **159**, 69–87.

[17] Isaacson, D. (2009), 'The reality of mathematics and the case of set theory', *in* Z. Novák and A. Simonyi, eds, *Truth, Reference and Realism*, Central European University Press, Budapest, pp. 1–75.

[18] Joyal, A. and Moerdijk, I. (1991), 'A categorical theory of cumulative hierarchies of sets', *Comptes Rendus Mathématiques de l'Académie des Sciences* **13**, 55–58.

[19] Joyal, A. and Moerdijk, I. (1995), *Algebraic Set Theory*, Cambridge University Press, Cambridge.

[20] Lawvere, W. (1969), 'Adjointness in foundations', *Dialectica* **23**, 281–296.

[21] Lear, J. (1977), 'Sets and semantics', *Journal of Philosophy* **74**, 86–102.

[22] Lewis, D. (1983), 'New work for a theory of universals', *Australasian Journal of Philosophy* **61**, 343–377.

[23] Linnebo, Ø. (2008), 'Structuralism and the notion of dependence', *Philosophical Quarterly* **58**, 59–79.

[24] McGee, V. (1997), 'How we learn mathematical language', *Philosophical Review* **106**, 35–68.

[25] McGee, V. (2000), 'Everything', *in* G. Sher and R. Tieszen, eds, *Between Logic and Intuition. Essays in Honor of Charles Parsons*, Cambridge University Press, Cambridge, pp. 54–78.

[26] McGee, V. (2006), 'There's a rule for everything', *in* A. Rayo and G. Uzquiano, eds, *Absolute Generality*, Oxford University Press, Oxford, pp. 179–202.

[27] Parsons, C. (1974), 'Sets and classes', *Noûs* **8**, 1–12.

[28] Parsons, C. (1977), 'What is the iterative conception of set?', *in* R. E. Butts and J. Hintikka, eds, *Logic, Foundations of Mathematics, and Computability Theory. Proceedings of the Fifth International Congress of Logic, Methodology and the Philosophy of Science (London, Ontario, 1975)*, Reidel, Dordrecht. Reprinted in Benacerraf and Putnam 1983, 503–529.

[29] Paseau, A. (2005), 'On an application of categoricity', *Proceedings of the Aristotelian Society* **105**, 395–399.

[30] Paseau, A. (2007), 'Boolos on the justification of set theory', *Philosophia Mathematica* **15**, 30–53.

[31] Pedersen, N. and Rossberg, M. (2010), 'Open-endedness, schemas and ontological commitment', *Noûs* **44**, 329–339.

[32] Potter, M. (2004), *Set Theory and Its Philosophy*, Oxford University Press, Oxford.

[33] Putnam, H. (1980), 'Models and reality', *Journal of Symbolic Logic* **45**, 464–482.

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

[35] Rayo, A. and Williamson, T. (2003), 'A completeness theorem for unrestricted first-order languages', *in* J. C. Beall, ed., *Liar and Heaps*, Clarendon Press, Oxford, pp. 331–356.

[36] Shapiro, S. (1983), 'Conservativeness and incompleteness', *Journal of Philosophy* **80**, 521–531.

[37] Shapiro, S. (2003), 'All sets great and small: And I do mean all', *in* J. Hawthorne and D. Zimmerman, eds, *Language and Philosophical Linguistics*, Vol. 17 of *Philosophical Perspectives*, Blackwell, Oxford, pp. 467–490.

[38] Sider, T. (2009), 'Williamson's many necessary existents', *Analysis* **69**, 250–258.

[39] Tarski, A. (1958), 'Intervention sur le rapport de M. le Professeur Skolem', *in Le raisonnement en Mathématiques et en Sciences Expérimentales: Actes du Colloque de Logique Mathématique (Paris, 1955)*, C.N.R.S., Paris, pp. 17–18.

[40] Uzquiano, G. (2002), 'Categoricity theorems and conceptions of set', *Journal of Philosophical Logic* **31**, 181–196.

[41] Walmsley, J. (2002), 'Categoricity and indefinite extensibility', *Proceedings of the Aristotelian Society* **102**, 239–257.

[42] Williamson, T. (2000), 'Logic and existence', *Proceedings of the Aristotelian Society* **100**, 321–343.

[43] Williamson, T. (2003), 'Understanding and inference', *Proceedings of the Aristotelian Society* **sup. vol. 77**, 249–293.

[44] Zermelo, E. (1930), 'Über Grenzzahlen und Mengenbereiche: Neue Untersuchungen über die Grundlagen der Mengenlehre', *Fundamenta Mathematicae* **16**, 29–47. Reprinted and translated in Ewald 1996, 1209–1233.

Luca Incurvati
Universiteit van Amsterdam
Afdeling Filosofie
Institute for Logic,
Language and Computation
Oude Turfmarkt 141–147
1012 GC Amsterdam
L.Incurvati@uva.nl
https://sites.google.com/site/lucaincurvati/