

TALKING ABOUT IT ALL

ABSTRACT. A well-known argument of Hilary Putnam's seems to suggest that our quantifiers might not possess the generality we take them to have. This paper argues that we ought not to be worried about this, and that the contrary impression stems from a certain approach to the relationship between model theoretic semantics and the theory of meaning. Once we see that the argument to the sceptical conclusion relies on our being able to assume an external perspective on language, and that this is not possible, the worry dissolves. The paper first lays out Putnam's argument, before meeting it with McGee's open-endedness response, motivating it by means of an inferentialist account of meaning. This is defended against objections from Williams, before a discussion of the use of model theory in philosophy of language. The paper concludes with a brief consideration of second-order languages.

'The ideal, as we think of it, is unshakable. You can never get outside it; you must always turn back. There is no outside; outside you cannot breathe. – Where does this idea come from? It is like a pair of glasses on our nose through which we see whatever we look at. It never occurs to us to take them off'.

Wittgenstein – Philosophical Investigations §103

Does 'all' mean *all*? The question might seem to embody the kind of mischief-making that gives philosophy a bad name. Of course 'all' means all, one is tempted to respond in exasperation, *what else could it possibly mean?* The argument of this paper is that this bad-tempered response is basically correct. But there is something to the philosopher's question. For universal quantification is put to use in articulating some of the most important claims we make. 'All people should be treated equally', somebody might sum up their core ethical belief. 'God is the creator of all that is', another might articulate a central religious doctrine. The physicalist metaphysician, meanwhile, assures us that 'everything is physical'. If it turns out that the egalitarian's 'all' picked out everyone except for some poor individual on an island somewhere, or that the theist excludes some distant galaxy from the scope of her assertion, or that the physicalist leaves some entities out of consideration when articulating her position, then their utterances do not have the intended generality. Ought we to be worried that this kind of scenario in fact obtains?

A well-known argument of Hilary Putnam's might be thought to provide reasons for a positive answer to this question [13]. The argument of the present paper is that it does not, and that the contrary impression stems from a certain approach to the relationship between model theoretic semantics and the theory of meaning, which functions within the philosophical community rather like a pair of tinted glasses, conditioning how we view the philosophical terrain. Once we see that the argument to the sceptical conclusion relies on our being able to assume an external perspective on language, and that this is not possible, the worry dissolves. Section One lays out Putnam's argument, as developed in *Models and Reality*. Section Two presents McGee's open-endedness response, motivating it by means of an inferentialist account of meaning. Section Three describes, and responds to, an objection to McGee's approach from Williams. Section Four uses the favoured inferentialist account of meaning as a basis for discussing the role of model theory in the Putnamian argument. Section Five applies briefly the developed position to the case of second-order logic. Finally, a concluding section draws morals for the theory of meaning and re-examines Putnam's own conclusions from his model-theoretic argument.

Consider some speaker of a language \mathcal{L} , call her Sara. We'll assume for present purposes that \mathcal{L} can be represented faithfully using the language of first-order logic. We'll consider the second-order case below. Consider all the sentences Sara is prepared to assert; we might call this her *theory*. With an eye to human limitations, we will take Sara's theory to contain only finitely many sentences. Following Williams' development of Putnam, we can also assume that the reference of the non-quantificational parts of Sara's language are given and 'intended': so, for example, the predicate 'is red' applies to all and only red things, and 'Sara' refers to Sara herself [20]. Now we can enter upon the kind of strategy deployed by Putnam in *Models and Reality*.

Let's supply a model-theoretic semantics for Sara's theory. Make this the intended model, where not only are predicates and names interpreted as above, but also the domain of quantification contains absolutely everything Sara intends to talk about.¹ Since Sara's theory contains a reasonable amount of science, and the mathematics seemingly necessary to formulate it, the domain of the intended model is going to be uncountable.

Now Putnam's move can be made. By the Löwenheim-Skolem theorem, the details of which can be found in any introductory mathematical logic textbook, since Sara's theory has an infinite model, it has a countable model. That is to say there is a model that makes every sentence in Sara's theory true, but the domain of which contains no more entities than there are natural numbers. This is clearly not the intended model of Sara's theory, since that theory includes sentences that, as Sara would put it, say things like 'there are uncountably many reals'. Yet it is a model nonetheless, and nothing manifest in Sara's use of language would seem to rule out her sentences having the meanings encoded by the 'new' countable model rather than those encoded by the intended model. What could decide the matter? The countable model is after all a model of Sara's theory, and so satisfies the constraint of making every sentence in that theory true.

Perhaps, though, Sara has an answer. Nothing counts as a model of her theory unless the names and predicates are interpreted as in the intended model. At this point appeal could be made to some kind of causal constraint on reference. Unfortunately the Putnamite sceptic can still get purchase. A strengthening of the Löwenheim-Skolem result shows that the intended model of Sara's theory will have a countable submodel, the domain will be a subset of the domain of the intended model, and the interpretations of names and predicates restrictions of interpretations in the intended model. In other words, if Sara's sentences are interpreted in accordance with the countable submodel she will get nothing wrong; not only do those sentences come out true, but also her expressions have the referents we might ordinarily suppose them to have. However, if the countable submodel codifies the meanings Sara's sentences in fact have her quantifiers do not bind variables ranging as widely as we might have supposed from the perspective of the intended model. From that perspective, 'all' as uttered by Sara does not mean *all*, but merely *all things in this countable subset*.

There is nothing obvious in Sara's use of language that rules out the following scenario: the countable submodel correctly encodes the meaning of her quantified locutions. Suppose now that this is the case but the extra-linguistic world is accurately represented by the original uncountable model. Then it would seem that there is a clear sense in which which Sara's quantified utterances fail to have the intended meaning, indeed in which her 'all' doesn't mean all. She systematically fails to quantify over all but countably many of the

¹I ignore here complications arising from the possibility of non set-sized domains.

entities that exist. Can this scenario be excluded? Is there even a fact of the matter? The aim of the rest of the paper is to suggest that we should not be worried by these questions.

2

Putnam's argument undermines our confidence that our quantifiers have determinately the meanings we naturally suppose them to possess. A natural thought is to look for some constraint on admissible interpretations that rules out Putnam's interpretations as unintended, and an obvious place to look for such a constraint is to the rules governing the use of quantificational expressions by competent language users. As we will discuss in greater detail in due course, an appealing and current position in the philosophy of language views the meaning of expressions to be constituted by their use in the language. There is a strong case moreover that – in the case of logical vocabulary at least² – the use of expressions, in the salient sense,³ is to be identified with their deductive role [4] [2]. In the case of quantifiers this is made explicit by their introduction and elimination rules.

Or so I want to suggest. Some care is required here though, since unlike the language of a formal logic equipped with a proof system, English does not come packaged with introduction and elimination rules for its logical vocabulary.⁴ There is a certain social practice, the use of English, which speakers engage in without in the main giving any thought to the explicit formulation of the norms and rules governing that use. Any appeal to introduction and elimination rules to resolve a question about the meaning of English quantifiers will, then, be vulnerable to the objection that the rules to which appeal is made do not formalise accurately natural language use. The path of least resistance in response to such an objection involves dropping the claim that the formalisation is adequate for English whilst maintaining that there is some value in securing the determinacy of '∀', perhaps also suggesting that this might contribute to future progress with respect to English quantifiers. Thus McGee,

To pretend that standard formalizations adequately represent the logical structure of English oversimplifies on a vast scale, for the inadequacies of the formalization techniques are infamous. (One discrepancy that is particularly relevant to present concerns is that English permits and the formalized languages forbid nondenoting proper names.) The motive for nonetheless working with the formalized languages is a methodology of starting with the very simplest cases and advancing to the more complex. The formal languages are the linguistic equivalent of frictionless planes. [12, 56]

I do not share McGee's pessimism about the adequacy of '∀' as a formalisation of at least a certain use of an English quantifier. In particular, there are several responses available to the concern about nondenoting proper names. One could, for example, simply deny that there are any such names in our language, after the example of Russell [18];⁵ another is to propose that reference need not imply existence [1]. If McGee is correct

²The difficulty of isolating the logical vocabulary is hereby noted only to be set aside for present purposes [6]. I do not think the problem here is insurmountable; indeed I think that an inferentialist approach to semantics could prove a promising tool for resolution. On specifically *logical* inferentialism see for example [7].

³That is, their use when constituents of token sentences subject to assertoric force, and insofar as this is relevant to the reasoned assertion or denial of the proposition expressed by the sentence. On the view under consideration this is meaning-constitutive in that it determines the sense of expressions. There are, of course, broader conceptions of meaning, but these are not what is at issue in Putnam-type cases.

⁴And *mutatis mutandi* for other natural languages

⁵At least if 'proper name' is supposed to be a semantic, rather than a surface grammatical category. If the latter is intended in McGee's objection, then the Russellian rejoinder is that no sentence containing a nondenoting proper name is appropriately rendered in a formal language in terms of a wff containing a constant corresponding to the name (rather, the name get analysed away in favour of quantificational apparatus).

about the inadequacy of ‘ \forall ’ as a formalisation of English universal quantification, however, what follows should be understood as provisional pending adaptation to the natural language case. Putnam’s argument is troubling because it calls into doubt the determinacy of *English* quantification. We want to be reassured that when we say ‘everything is either a particular or a universal’ a determinate thought is expressed which involves quantification with unrestricted generality.⁶ Nothing purporting to be a non-concessive response that doesn’t secure this is successful. With that brought into the open, the rest of the paper will proceed in the light of my contentment with the natural deduction rules as capturing English quantification.⁷

Ignoring for present purposes the case, to my mind a strong one, for a bilateral formulation of proof theory [16], the rules for the universal quantifier are these:

$$\frac{\phi(t)}{\forall v \phi(v)} \forall I$$

$$\frac{\forall v \phi(v)}{\phi(t)} \forall E$$

We place the usual restrictions on substitutends for t in $\forall I$. As regards the elimination rule, Vann McGee, with an eye to our project of answering the generality sceptic, advocates an expanded understanding of the range of t there. Under the usual development of first-order languages, the intended range is over all terms in the language. But this, thinks McGee, concedes too much to the sceptic. For suppose that there is some entity, say Bob, for whom there is no name in our language. Then nothing about our inferential practice will ever manifest that our universal generalisations range over Bob. We can never eliminate a quantifier in favour of a name for Bob,⁸ for there is no such name. And if we use a free variable as substitutend for t in the conclusion of an application of $\forall E$, then neither will this in any way communicate that Bob is one of the entities we mean when we say ‘everything’. Let’s understand free variables as tacitly bound by a prenex universal quantifier, which seems a natural gloss on both the rules of inference and model-theoretic semantic treatments.⁹ Now the question of whether Bob is within the range of ‘everything’ reduces to the question whether Bob is in the range of ‘everything’, which indeed it does, but which doesn’t take us any further forward! Names, in other words will have to do all the heavy lifting in manifesting that Bob is one of the entities we mean when we say ‘everything’. Yet, as we have seen, it appears that they cannot, since we have *ex hypothesi* no name for Bob.

How then, we might worry, can we be sure that ‘all’ really means all, that Bob doesn’t get left out when we quantify? The answer, thinks McGee, involves recognising that our acceptance of quantifier rules is open-ended. When a speaker accepts $\forall E$ she does so on the basis that it remains truth-preserving if the language is extended by the addition of further names. Thus my assent to the validity of $\forall E$ yields confidence not simply that any instance of it will preserve truth *given the names available to me to substitute for t right*

⁶We want, in other words, to be able to engage with texts like [15] intelligibly.

⁷I think at least some reluctance on this point – although not McGee’s – issues from a failure to observe the semantics/ pragmatics distinction and to recognise that formalisation need only respect the former. We are concerned with assertoric content, not with the whole rich tapestry of linguistic communication

⁸In the spirit of taking the quantifier rules to formalise successfully the rules of use for English quantifiers, I take English names to correspond to individual constants.

⁹My use of ‘inference’ is the standard one within the theory of meaning literature. But see [17, 50-56] for discussion of the term in relation to *deduction*.

now but also that this will not cease to be the case if my nominal repertoire is extended. If this commitment is implicit in my acceptance of $\forall E$ then, in particular, if I am in a position to assert $\forall x\phi(x)$ then I incur a commitment to the assertibility of $\ulcorner\phi(\text{Bob})\urcorner$ when the name ‘Bob’ becomes available.¹⁰ Yet openness to this commitment is already present in my quantificational practice here and now, which is to say – we might reasonably conclude – that Bob is within the range of my quantifier. In particular, any model-theoretic interpretation of the quantifier on which $\forall E$ ceases to be valid if a name for Bob is introduced to the language is shown up as unintended. The Putnamian worry is defused.

Or is it? In unpublished work, Robbie Williams argues that McGee’s line of argument fails by being insensitive to the distinction between actual and counterfactual evaluation [20, 18]. To see this, consider a language user, Amy. Let us suppose that Amy’s quantifier is restricted so as not to include Bob within its range. Does McGee’s appeal to open-endedness undermine this supposition? Williams thinks not. Call w the world at which Amy is going about her everyday quantificational business, in blissful unawareness of Bob’s existence and so lacking any name for Bob. Here in w , Williams suggests, there is no counterexample to $\forall E$ even given the supposition that Amy’s quantifier is restricted. For sure, Amy *could* speak an expanded language which contains a name for Bob. But in order to assess the truth of Amy’s quantified utterances in this situation, which is counterfactual, we need to evaluate them with respect to a world v , where $v \neq w$ and Amy speaks a language at v which contains a name for Bob. Williams maintains that there is no good reason to think that the restriction on Amy’s quantifier is constant over w and v . It is consistent with Amy’s quantifier being restricted so as to exclude Bob at w that it includes Bob within its range at v . There is no counterexample to $\forall E$ at either world, and the invocation of open-endedness does not suffice to quieten Putnamian concerns.

3

Is this right? Does McGee fail to block Putnam’s scepticism through not marking the distinction between actual and counterfactual evaluation? Let’s review the case of Amy. Let ‘ c ’ be a name for Bob. Amy does not have the use of ‘ c ’ available to her. There is more than one reason why this could be the case. It could be that the linguistic community of which Amy is a part has ‘ c ’ as a lexical item, but that Amy hasn’t acquired it. Alternatively it could be that the name ‘ c ’ has not yet been introduced to the language, but will be – or at least, to put the point more concessively to the open future, it could be. Then again, perhaps Bob sits outside of Amy’s future light cone, and that of every member of her linguistic community, along with that of every member of every linguistic community which could, through contact with Amy’s community, bring it about that Amy’s language acquires a name for Bob.

Starting with the last case; if Bob lies forever beyond the shores of language I am happy just to concede that Amy’s quantification does not include Bob in its reach. Were it otherwise, there would be some part of the meaning of Amy’s quantifier that could never be manifested, and it is mysterious how Amy could have acquired a lexical item with meaning of this sort in a fashion compatible with learning it as part of a public language.¹¹ This might be thought to be to concede *everything* to the Putnamian sceptic; after all isn’t it precisely this kind of scenario that the deviant models beloved by the sceptic are supposed

¹⁰I am ignoring complications about word individuation here. If you think that ‘Bob’ as used to refer to Bob (as talked about in the previous paragraph) and ‘Bob’ as used to talk about that object over there, distinct from Bob, then think of my commitment as being to the assertibility of ϕ with ‘Bob’ as used to refer to Bob substituted.

¹¹It is, of course, entirely possible that there are nameable entities not within any language user’s future light-cone - those characteristically classed as abstracta, such as numbers, provide an obvious example (although it also seems wrong to say that these lie *outside* any light-cone). It is better, then, to say that temporal open-endedness is with respect to initial baptisms rather than baptised entities. We’ll ignore this complication in what follows.

to represent? There's Bob – over there in the outer reaches of the universe – we'll never get to name him; how do we exclude the possibility that our quantifiers aren't correctly interpreted according to one of the undesirable models? Many issues are in the vicinity here, including the thorny ones which fall within various realism-antirealism debates. For present purposes, I issue a promissory note to be discharged in a subsequent section of this paper¹² – once we have been persuaded that we ought to be worried about this, we are well down the path towards meaning scepticism. The remedy is to realise that, on any reading where attention to the model-theoretic argument prompts an intelligible question which invites an answer of us, that question is about the meaning of our word 'all' and is asked in our language.¹³

What if Amy's linguistic community have the expression '*c*' but Amy herself hasn't acquired it? Here we should take seriously the fact that Amy's quantifiers belong to a public language. In acquiring the language, Amy comes to participate in a social practice governed by conventions which are expressed by the rules.¹⁴ Furthermore since it is a social practice that she participates in, she is tacitly committed to any instance of ($\forall E$) expressible within the legitimate bounds of the practice. So now imagine that Amy is in conversation with Darryl. Amy asserts $\ulcorner \forall x \phi(x) \urcorner$. Darryl, who is sceptical on the matter, challenges her by asking her whether she is committed, therefore, to $\ulcorner \phi(c) \urcorner$. Not having acquired the relevant singular term Amy asks what *c* is. Darryl duly explains, perhaps by offering some explication of the referent in terms of an indentifying description. Amy can either accept Darryl's word, or else express doubt ('Oh, *come off it*, there can't be anybody like *that!*') at which point Darryl will be likely to appeal to the wider community of language users for attestation of the usage; these days this can be done by a simple search on a smartphone. Once convinced that '*c*' is indeed part of the language, Amy will do one of two things. She can either withdraw her initial claim, either absolutely or replacing it with a modified claim (perhaps a quantification restricted to the satisfiers of some sortal expression), or else she can accept the inference to $\ulcorner \phi(c) \urcorner$. Either way, what is manifest is that the applicability of ($\forall E$) over every name in the language is *already* implicit in Amy's knowing how to use the quantifier. She recognises, once she has accepted that *n* is a name in her language, for any *n*,¹⁵ then she is committed to accepting the inference from a universally quantified claim to the relevant instance involving *n*.

The reason for this is that understanding an expression consists in knowing how to participate, to some level of competency, in a *social practice* which by its very nature extends beyond the language user. The meaning of Amy's quantifier is given by that practice as a whole, not simply by Amy's inevitably partial grasp of the practice: were this not the case, each expression would have a different meaning as used by each language user, and the

¹²§3 below.

¹³This is a slight oversimplification, since I could ask you about the meaning of English 'all' in another natural language. No worry: translate the question into English and proceed as before.

¹⁴Which is not to say, of course, that she must be able to state the inferential rules for quantifiers, or even know that she is abiding by them or abide by them all the time. It is rather that in knowing how to quantify in English, she knows how to follow the rules.

¹⁵*Objection:* this account is circular, since it requires universal quantification to be stated. *Response:* the knowledge expressed in Amy's recognition is practical, not propositional. Stating, as part of a theory of meaning, what it is that Amy understands when she understands quantified sentences requires universal quantification. But in order to know how to use quantifiers, Amy doesn't need to possess such a theory, and therefore doesn't mean prior access to quantificational resources. I do not think there is any further deep question about why Amy's linguistic community engages in this practice rather than another. Wittgenstein's *On Certainty* is instructive here: 'You must bear in mind that the language-game is so to say something unpredictable. I mean: it is not based on grounds. It is not reasonable (or unreasonable). It is there - like our life..' [21, 559]

possibility of communication would be rendered mysterious. Instead, in acquiring a quantificational expression, a language user comes to possess an implicit grasp of the forms of inferential move it licences and a commitment to the licit status of all moves of these forms within the language, whether or not making them is presently within her competency. By partial analogy, a language user who understands a natural kind term knows how to use it on at least some occasions: she will deploy it appropriately in her discourse, and be prepared to move from sentences containing it to other sentences (perhaps containing some determinable under which its referent falls). Yet her grasp of the term's conditions of application will almost invariably be highly partial. There is, as Putnam realised, a social division of linguistic labour. The chemist has the knowledge and equipment to identify definitively whether or not a substance is water, and we duly defer to her in our usage of 'water'. In both cases, an adequate account of the practice of language users can only be had by paying due attention to the social nature of language.

Of course, we have not always possessed the capacity to identify definitively whether or not some substance is H_2O , and the use of the word 'water' predates this capacity. The practice of language is extended not just synchronically across individual language users, but diachronically as well. What then are we to say about the case mooted above where a name for some entity has yet to be introduced to the language, or at least *might* be? Here is where we need to insist that acceptance of rules for quantifiers is *temporally* open ended. To see this, consider a conversation between Amy and Bernie. Amy opines that all planets other than Earth are uninhabited. We may formalise her claim thus,

$$\forall x ((Px \wedge \sim Ex) \rightarrow Ux)$$

Having shared her opinion, Amy leaves the conversation. A month later, Amy and Bernie meet for their occasional coffee. In the interim, Donald Trump having diverted funding from Medicare to aggressive space exploration, a new planet has been discovered, inhabited by complex lifeforms. Reflecting on this development, Amy says to Bernie, 'you know, it turns out I was wrong. Not every planet other than Earth is uninhabited'. Does she speak truly when she says that her earlier statement was wrong? Surely she does, but this can only be the case if the meaning of her quantifier was, on the earlier occasion of use, such as to be truth-preserving not just with respect to names currently in the language but with respect to names yet to be introduced to the language. Language serves the function of communicating across time, an aspect which is most evident in the practice of writing.¹⁶ Moreover, since our linguistic practices are conducted as though the future were open, and therefore as though any entity within the future light-cone of some user of the language, is potentially nameable, temporal open-endedness secures a generous determinacy of quantification.

There are, then, various ways in which Amy might not have some name '*c*' available to her without her quantifier rule failing to be open-ended with respect to '*c*'. So far we have no counter-example to open-endedness. Williams' example might tempt us to think, however, that we are ignoring the most pressing case. Even if we can be satisfied that quantifiers are robust with respect to epistemic and temporal modalities, isn't the worry one about metaphysical modality, hence the talk of possible worlds? Suppose that we are satisfied that Amy's use of quantifiers is open-ended relative to both her present state of knowledge and her location in time. Couldn't it still be the case that it is *possible* that there is some un-named entity which could create Putnam-style problems for Amy's quantifiers?

¹⁶And which can raise practical problems, for instance with the design of warning signs on long-term nuclear waste storage facilities. See [].

What could this mean? The objection cannot be simply that there could be an entity such that whilst it is physically possible that it be named in the future by some member of Amy's linguistic community in fact this never occurs. This eventuality is already incorporated into our account of temporal open-endedness. If the worry is that it is possible that there be some entity that is not in actual fact nameable, then this was acknowledged above as raising questions about realism and anti-realism. All I claim is that it is not a peculiar problem for the philosopher claiming that the meaning of quantified expressions is adequately represented by open-ended rules, since it is a puzzle for any account of quantifier meaning how (if at all) we can quantify over entities that forever escape the prospect of individual naming.¹⁷

If these are not the salient senses of 'possible' then how else could it be possible that an object be unnameable? We are surely now in the realm of merely metaphysical possibility. To use the popular heuristic, there is some possible world in which there is some entity which the linguistic community is kept from naming for reasons other than time and distance. Perhaps, for example, an evil demon intervenes to prevent any initial baptism of the entity. Doesn't this show that open-ended quantifiers don't secure that 'all' means all? No it does not. Our concern is with our actual practice of naming, not counterfactual ones, nor is it a sensible constraint on philosophy of language that it have a response to the kind of meaning scepticism that would issue from considering possibilities such as that now under consideration as being actual. My own view is that in the world in which naming is prevented by unusual means, the situation is relevantly similar to that in which distance prevents naming and in order to theorise about quantification in those circumstances one would need to ask hard questions about realism. Given, however, that we are concerned with the quantifiers we actually use, we need not be detained by this worry. Merely possible problems do not undermine accounts of actual quantification, and we should be satisfied that open-ended rules provide a good avenue of response to Putnam's argument. The suggestion of the next section is that they provide more than this.

4

The inferentialist holds that the meaning of quantifiers is given by their use, and in the this use is made explicit by their introduction and elimination rules. To understand a quantifier, one needs to grasp which moves it licences within the language game (as Brandom puts it, the game of giving and receiving reasons [2]). Absent from this account is any mention of a language-external domain. It is most natural to see the equivalent of the domain existing for the inferentialist at one level of semantic ascent, in the form of the singular terms of the language. A natural philosophical outworking of this perspective would be the view, developed by Bob Hale amongst others since Frege, that objects just are possible subjects of naming [8]. Language has a certain priority over ontology. By comparison if a model-theoretic semantics is understood as a theory of meaning for a language, the contents of the domain as such will feature heavily in all but the most sparse account of that meaning, and it is here that the existence of countable submodels gets its bite as presenting a problem for quantifier *meaning*. For it is only if the domain of models represents something integral to the meaning of quantified formulae, and in particular if the cardinality of the domain does this, that the availability of countable submodels for theories raises a worry about meaning. In what follows I will draw out an entirely natural picture of how domains do indeed represent meaning in a suitable way and then, in spirit of the epigraph from Wittgenstein, diagnose the naturalness of the picture and suggest that once we have seen what is wrong with a certain kind of model theoretic orientation in the

¹⁷Look at the matter like this: if she wants to support quantification over the unnameable, the realist philosopher is going to have to offer an account of quantifier meaning which makes sense of this and satisfies other reasonable constraints in terms of explaining how quantificational vocabulary could be intelligibly acquired etc.

theory of meaning we will no longer be troubled by Putnam's argument.

What am I doing when I utter a declarative sentence with assertoric force? A natural thought is that I am describing the world as being some way. A further thought is that the success of the description, and so the truth of the utterance, consists in correspondence to the world. Language in some sense pictures the world. It is tempting to add to this broad account, unobjectionable as far as it goes, the further thought that the world is given prior to language, that there is so to speak a domain of objects out there waiting for us to name and group together. This claim is strictly stronger than the simple denial of linguistic idealism: it is one thing to think that there is a reality independent of our linguistic practices; it is another thing altogether to think that this reality is carved up determinately into discrete objects of potential reference and thought. It is this latter claim that is contained in the account of model theoretic semantics now on the table. In the beginning was the object, then came the word. On this picture the domain of a model represents the world, out there prior to language, and interpretation and assignment functions represent the subsequent picturing of the world by means of language. Truth consists in correspondence between picture and reality. Meaning meanwhile consists in the picture presented by a sentence: it is what would be required of the world (or of a model, as representing the world) in order for the picture to be a true one.

Notice that in order for this account to cohere, we need to be able to at least grasp what it would be for there to be a determinate reality consisting of discrete objects prior to the imposition of language. The use of model-theoretic semantics as a theory of meaning on the present account involves taking an *external perspective on our language*. The meaning of names is identified with particular elements in the given domain, of predicates as collections of such elements and of logical connectives as functions. In the case of the universal quantifiers standard model-theoretic satisfaction clauses are most naturally read as assigning to these functions from elements of the domain (picked out by assignments to the bound variable) to truth-values, given intensionally by the matrix. A more Fregean recourse would be to understand quantifier meaning in terms of functions from subsets of the domain to truth-values, given intensionally by the predicate(s) concatenated with the bound variable, and this could easily be accommodated within the framework of model-theoretic semantics. Either way, the domain is constitutive of quantifier meaning.

Enter Putnam. Now the worry is this: suppose we live in a world that is truly represented by a domain of cardinality 2^{\aleph_0} . Nothing about our use of language, and in particular nothing about those sentences we are prepared to assert (if you like, our theory), assures us that our quantifiers don't have the meanings associated with a countable sub-model of the 'real' model. (Remember that we are operating within a picture on which the world is 'out there' prior to language). So now we face the worry: does 'all' mean all? What if it doesn't? Perhaps my claim that everything is physical, or that everyone deserves respect are less all-encompassing than I may have thought. And this surely matters.

It would indeed be a problem if our attempts at universal statements fell short of generality, but notice what is required for the strengthening of the downward Lowenheim-Skolem theorem to motivate the thought that this is a genuine possibility. First, it is necessary that model theoretic semantics provides a correct theory of meaning for the language (distinguish providing a theory of meaning from being a useful tool: we'll return to this below). Second, model-theoretic semantics should be interpreted in such a fashion that the domain is given prior to language. Encompassing both these points, and crucial for the appeal to countable submodels, is the external perspective on language. We can, on this view, stand outside our own linguistic practices sufficiently to make sense of the thought that they come

apart radically from the world, as would be the case if they were correctly represented by a countable model whilst the world itself was correctly represented by a continuum-sized model.

Why are we led to believe this could be the case? Because we can go through the process for a formal language of our own devising. We set up, say, the language and theory of first-order analysis, identify one interpretation (namely $(\mathbb{R}, <)$) as being the intended one, and then proceed to demonstrate the existence of the problematic submodel. We then wonder whether we might not be in the unfortunate position of using expressions with restricted meanings (as in the submodel) to talk about a more expansive reality (as in the domain of the intended model). Yet in order to even state the problem we have needed to approach the language of first-order analysis from the outside, as competent users of another language – in this case the mathematicians’ English (or other natural language) used to express the model theoretic semantics. Within this language we talk about the domain and its contents, and using it we prove the existence of the relevant submodel. It is because we have a prior grasp on reality by means of the metalanguage for the model theory that we can raise puzzles about the object language’s expressive capabilities. We have not, in other words, stepped out of language altogether when working through Putnam’s argument, but only out of a language we ourselves brought about for the purposes of logical study and research.

The naturalness of the external model-theoretic picture of meaning arises from forgetting that the natural languages in which we are immersed are relevantly dissimilar from the formal languages of mathematical logic. Certain aspects of the declarative parts of our language, used assertorically, are modelled well by model theoretic semantics, and this can be useful for various purposes in philosophy and linguistics. What is not the case, however, is that we have an understanding of reality as divided up into objects of various sorts that we can articulate other through natural language. The very fact that we pose the question in terms of whether ‘all’ genuinely means *all* should alert us to this. Of course ‘all’ means all; English, after all, disquotes. If we were setting up a version of Putnam’s argument for English we will pick out the intended model using English quantification – the domain is (all) the reals. Moreover, supposing for the sake of argument that we can make sense of the thought that ‘all’ might not mean *all*, then if this were the case then the sceptical worry about meaning could not even be framed, since the question in English whether ‘all’ means all would not succeed in describing a sceptical scenario. We simply cannot get the external perspective on the languages in which we routinely think and speak about the world required to create difficulty here.¹⁸

Nor do we have any cognitive handle on the inventory of the world independent of language, which, in Dummett’s phrase ‘may be a distorting mirror but [is] the only mirror we have’ [5, 6]. The inferentialist alternative to model theoretic semantics scores well on both points. It does not suggest the availability of theoretical knowledge about the languages we use independent of the possession of those languages, but nonetheless presents us with a substantial theory of meaning for part of language through specifying rules, adherence to which is constitutive of competent use of the language. The focus is shifted from knowledge *that* to knowledge *how*, a theory of meaning serving to explicate the latter, as distinguished from stating the former. Nor does the inferentialist make any appeal to an extra-linguistic domain. Her account of quantification does not involve a domain, but instead uses metavariables ranging over the names through which we reach out linguistically

¹⁸What about creating difficulties for one natural language using another as the metalanguage? I think the ready translatability of natural languages into one another blocks this move. If I want to explain to an English speaker what ‘tous’ means, I will say ‘all’, and *vice versa* for a French speaker.

to the world. Once we have diagnosed concern about Putnam’s argument as involving a non-compulsory perspective on language and reality, an inferentialist account of quantifier meaning, which avoids any suggestion of this perspective, looks attractive.

Where does this leave model-theoretic semantics, however? They are an established part of mathematics, and it would be a foolhardy philosopher who sought to displace them from this position. The question is how model-theory relates to the study of language, in the sense in which philosophers of language are concerned with, not simply the elements of some class of strings satisfying some minimal constraints, but rather the means by which agents such as ourselves in fact communicate with one another. Given a regimentation of such a language, or part thereof – for example of part of English into the language of first-order logic – along with a set of rules of use for the language, a soundness result licences the use of model-theory to study the relationship of logical consequence on the language, constructing counter-models rather than demonstrating proof-theoretic invalidity directly. Where obtainable, a completeness result ensures the coextensiveness of proof-theoretic and model-theoretic consequence, and further licences the use of model-theory to study the consequence relation. This is invariably more straightforward than working proof-theoretically and is a genuine gain obtained in model-theory. Nor is there any harm in using model-theory to picture relations between language and world, so long as one doesn’t read into this usage the possibility of a language-external perspective or the availability of a domain apart from language. Moreover, this application of model-theory needs to proceed on a case by case basis, if it is not to entangle us in unwarranted metaphysics. Whilst approaching a non-modal language in this way might be unobjectionable, for example – names, after all, refer to objects – it would be a mistake to look at the extension of model-theory to a modal language and on that basis alone infer the existence of possible worlds. Here, as elsewhere, model theory is a powerful tool but a poor master.

5

Something should be said about second-order logic and its use to characterise mathematical structures categorically. Issues arise here that have a resemblance to those raised by Putnam [13, 481]. Take a second-order language equipped with rules of use. It is well-known that such languages can be used, given the standard model-theoretic semantics, to characterise the naturals and reals, under their standard orderings, up to isomorphism [19]. However as well as the standard semantics, on which n -adic second order variables range over $\wp(\mathcal{D}^n)$, with \mathcal{D} the domain, there are Henkin semantics for second-order languages. On Henkin semantics, there is a distinct second-order domain $\mathcal{D}^* \subseteq \mathcal{D}$. We generally restrict our attention to faithful Henkin models, which satisfy every instance of the second-order comprehension schema. The resulting semantics are complete, compact and possess the Löwenheim-Skolem properties. There are therefore nonstandard models of second-order arithmetic and analysis on Henkin semantics. Yet, since these semantics are sound for the usual second-order rules of proof, there seems to be nothing in the use of second-order vocabulary that rules out, for example, our second-order arithmetical practice¹⁹ picking out some nonstandard model, against our intentions. We seem to be trapped in a position analogous to that discussed above with respect to first-order languages. Yet it is precisely in order to escape such difficulties that appeal is often made to second-order logic [19]. What is to be said?

First notice that the analogy with our initial problem is exact in some respects. The present worry is that nothing in the practice of second-order quantification assures us, from a model-theoretic perspective, that the quantifiers bind variables ranging over *all* subsets

¹⁹Which, very plausibly, just is our arithmetical practice. Why do we accept every instance of the first-order induction schema? Because we antecedently accept the second-order axiom, surely.

of the relevant domain. Once again, the question is whether ‘all’ really means *all*. And once again the solution is to see that it couldn’t mean otherwise, if our concern is with the languages we routinely use to speak and think about the world. At this point there are complications peculiar to the second-order case. Most importantly, it is not straightforward to identify second-order quantification in natural languages. These do not seem to contain quantification into predicate position, and natural language quantification over *ways, properties* and so on looks grammatically first-order. Maybe some principled distinction can be made which underwrites the formalisation of some parts of language using second-order quantification; alternatively, maybe a case can be made that the meaning of some collection of sentences, concerning arithmetic for instance, can only be adequately explicated if those sentences are given a second-order regimentation. Either way, assuming there is some part of natural language that is appropriately regarded as second-order (and if there is not, the Henkinite sceptic has no purchase in any case) we can appeal to that part in making our anti-sceptical case.

How seriously, though ought the model-theoretic semantics to be taken as an account of the *meaning* of natural language second-order quantification? In particular, ought second-order quantification to be thought of as over sets of the entities named in the language? Or is it over something else – things in plurality or properties? Alternatively, is it a mistake to regard second-quantification as in any way governing a referential position? [14] [22]²⁰ It is not necessary to have a firm answer here in order to defeat the meaning sceptic. We can set up a dilemma. Either the meaning of second-order quantifiers is not appropriately represented by the model theoretic semantics, in which case the availability of nonstandard models poses no challenge to the determinacy of meaning, or else it is, in which case we avail ourselves of our *actual* second-order quantificational resources when specifying the range of the variables bound by ‘object language’ quantifiers. It could be objected here that the salient range is specified by the use of *first-order* quantification in the set-theoretic metalanguage, but this is surely an artefact of an assumed external perspective in our language. If we do succeed in second-order quantifying, and that interpretation should be interpreted in terms of a domain of second-order entities, then unrestricted *first-order* quantification over those entities in the same language is barred on pain of Cantor’s paradox, at least given plausible assumptions about the second-order domain.

When I am asked about the meaning of my second-order universal quantification, there is no better answer than ‘all *F*’. This suffices to rule out a Henkin-style interpretation of the language as unintended. It is no objection to this response that it doesn’t supply what is not available, namely an account of quantificational practice which is both available to the respondent and yet external to her linguistic perspective. She succeeds in meaning all that she ever intended to mean. And in particular, assuming that second-order Peano Arithmetic and analysis capture much of our everyday preformal mathematics, she does so when she talks about key mathematical structures.

In any case, even if the model-theoretic semantics does not mislead us considered as an account of the meaning of second-order quantifiers, the approach developed above with respect to the first-order quantifiers suggests that this account will be at best partial, and in particular that it will not be full-blooded.²¹ If we take seriously the nature of language as a human practice, then there is a powerful pull towards thinking that our theory of meaning ought to be use-orientated. For second-order quantifiers, as for their first-order relatives,

²⁰With Rayo and Yablo I take the key issue bearing on whether second-order quantification is committing to be whether predicates are committing. For discussion see [10].

²¹I use ‘full-blooded’ here in Dummett’s sense [3].

this involves their meaning being given by the introduction and elimination rules. For the universal quantifier these can be formulated straightforwardly:²²

$$\frac{\phi(P)}{\forall V \phi(V)} \forall SI$$

$$\frac{\forall V \phi(V)}{\phi(P)} \forall SE$$

Here P ranges over predicate letters, and the restrictions on substitutends are the obvious ones. Following McGee's manoeuvre for the first-order case, and the above argument in response to Williams, we should insist that these rules are to be read as open-ended. As such, any way of interpreting them on which second-order variables do not range over every available value of the domain should be ruled out as unintended.²³ So any interpretation of our (second-order) arithmetical language for which that language is not determinately and appropriately rendered by the standard model $(\mathbb{N}, <)$ ought similarly to be ruled out. This is, of course, without prejudice to our ability to study nonstandard models of arithmetic using the usual model theoretic techniques, an activity which has proven mathematically fruitful and illuminating. The point is one about the meaning of mathematical language.

It is easy to imagine an objection being made here. Acquisition of a predicate involves developing the capacity to apply that predicate's application conditions to some minimal standard of competency. These conditions must be decidable (we are concerned with our language, not God's). But now consider the power-set of an infinite set. Amongst the elements of this will be arbitrary infinite subsets to which no finitely storable condition corresponds. Doesn't this show that there is something awry with our confidence in the capacity of our second-order quantifiers to describe mathematical structures? For, given that no predicate letter could ever pick out S , where S is an arbitrary infinite subset, can we really say that we are quantifying over S ? Once again, the difference between internal and external viewpoints is crucial here. The situation as we have described it – whereby there is a subset of an infinite set which cannot be picked out by any predicate in the language – is viewed from an external perspective on the language. Being able to take such a perspective on, say, a second-order formal language is useful for the purposes a model theory. Nothing follows, however, about the ordinary languages in which we routinely think and talk about mathematics. In particular, for instance, we ought not to be worried that lurking forever beyond the shores of our language is some set quantification over which alone will serve to distinguish the genuine natural numbers from the non-standard impostors.

A way to think about the situation model theoretically is in terms of Hale's semantics for second-order logic, which interprets second-order variables as ranging, invariably, over every definable subset of the first-order domain [8, Ch. 8]. The worry about non-definable subsets is thus avoided (we might usefully think of the restriction to definable subsets as representing an 'internal' perspective on the language model-theoretically), yet we keep

²²Matters are somewhat more delicate with the existential quantifier, if one wants to capture the content of the comprehension schema within the natural deduction rules, rather than resting content with the obvious second-order equivalents of the first-order rules and treating comprehension axiomatically. This is especially the case if harmony is considered a constraint on introduction and elimination rules.

²³Feel free to replace 'subset' with 'plurality' or 'property', or to rephrase the entire sentence in a neutralist fashion. The operative point concerns maximality of interpretation, not any particular metaphysics of second-order logic.

categoricity for arithmetic and analysis, and the consequent failure of limitative meta-theorems. The situation regarding set-theory (although not mentioned by Hale) is more complicated. Quasi-categoricity will be preserved (since the interpretation forces constant width at each V_α), but there has to be a strong suspicion on the basis of set-theoretic practice that it is not the ‘real’ set-theoretic universe that is being tied down by second-order ZFC so interpreted. For it is implicit in that practice that we do *not* restrict our attention to definable sets, and moreover that definable sets are not all the sets there are. The rejection of $V=L$ as an axiom, documented well by Maddy, attests to this [11,].²⁴ There is much that could be said here, but crucially appeals to second-order logic in set-theory take different forms. The use of second-order logic in set theory in order to eliminate first-order schemata is one thing, the use to make appeal to quasi-categoricity (for example in suggesting that CH has a determinate truth-value [9]) is another entirely, and much more controversial. It is not clear how much would be lost if it turned out that the supposition that such an appeal is possible traded on a misunderstanding of mathematical language.

6

We have seen that a use-based theory of meaning is well placed to deflate concerns arising from Putnam’s model-theoretic argument. There is no need to be a sceptic about the meaning of quantified locutions, nor to fear that one’s universal quantifier somehow fails of the desired generality. It is instructive to compare this with Putnam’s own outlook in *Models and Reality*. Here Putnam poses a trilemma between three families of theory of meaning: Platonism, moderate realism and (what he calls) verificationism [13, 464]. The Platonist ‘posits nonnatural mental powers of directly “grasping” forms’, whilst the moderate realist ‘seeks to preserve the centrality of the classical notions of truth and reference without postulating nonnatural mental powers’. The *classical notions of truth and reference* here are represented faithfully by model theoretic semantics, and are compatible with sentences having verification-transcendent truth conditions. Finally verificationism, ‘replaces the classical notion of truth with the notion of verification or proof, at least when it comes to describing how the language is understood’.

Putnam, on the basis of his model theoretic argument, adopts the verificationist approach which, he argues, blocks the way to meaning scepticism,

The crucial question is this: do we think of the *understanding* of the language as consisting in the fact that speakers possess (collectively or individually) an evolving network or verification procedures, or as consisting in their possession of a set of “truth conditions”? If we choose the first alternative, the alternative of “nonrealist” semantics, then the “gap” between words and world, between our *use* of language and its objects, never appears. [13, 480]

Few have followed Putnam down this path, the dominant philosophical response being to see his argument as providing a puzzle for moderate realism, to which it is incumbent upon philosophers to respond in defence of that type of theory of meaning. The approach I have taken in the present paper takes off from McGee’s open-endedness response and ties it in to an inferentialist account of meaning. This is very naturally seen as verificationist in Putnam’s sense – although the use of that term is unfortunate, since for many it suggests something like logical positivism; Dummett’s ‘justificationism’ might have been a better label. For the inferentialist the truth-conditions for a quantified locution will not be stated in non-linguistic ‘worldly’ terms, but will be given by the assertability conditions attaching

²⁴Of course, defineability in the Hlean sense is distinct from that operative in the construction of L, but the point carries over. For details and finessing of Hale, addressing difficulties ignored here, see [] and [].

to the introduction rule for the quantifier. Moreover, my criticism of the use of model theoretic semantics in motivating scepticism as involving an external perspective on language resonates with the above quoted passage from Putnam.

Putnam's use of model theory to cause problems in the philosophy of language has fascinated philosophers for nearly four decades. If I am right, his solution to those problems deserves more of a hearing.

*

REFERENCES

- [1] Elizabeth Anscombe, *Existence and the Existential Quantifier*, Logic, Meaning, and Truth : Writings by G.E.M. Anscombe (Mary Geach and Luke Gormally, eds.), Imprint, Exeter, 2015, pp. 294–299.
- [2] Robert B. Brandom, *Making it Explicit : Reasoning, Representing, and Discursive Commitment*, Harvard University Press, Cambridge MA., 1994.
- [3] Michael Dummett, *What is a theory of meaning? (ii)*, Truth and Meaning: Essays in Semantics. (Gareth Evans and John McDowell, eds.), Oxford: Clarendon Press, 1976.
- [4] ———, *The Logical Basis of Metaphysics*, Harvard University Press, Cambridge, MA., 1991.
- [5] ———, *Origins of Analytical Philosophy*, Bloomsbury, London, 2014.
- [6] John Etchemendy, *The concept of logical consequence*, Harvard University Press, 1990.
- [7] Ian Hacking, *What is Logic?*, Journal of Philosophy **76** (1979), no. 6, 285–319.
- [8] Bob Hale, *Necessary Beings*, Oxford University Press, Oxford, 2013.
- [9] Georg Kreisel, *Informal rigour and completeness*, Problems in the Philosophy of Mathematics (I. Lakatos, ed.), North-Holland, Amsterdam, 1967, pp. 138–86.
- [10] Fraser MacBride, *Predicate Reference*, The Oxford Handbook of the Philosophy of Language (Barry C. Smith, ed.), Oxford University Press, Oxford, 2006, pp. 422–275.
- [11] Penelope Maddy, *Naturalism in Mathematics*, Oxford University Press, Oxford, 1997.
- [12] Vann McGee, *Everything*, Between Logic and Intuition : Essays in honour of Charles Parsons (Gila Sher and Richard Tieszen, eds.), Cambridge University Press, Cambridge, 2000, pp. 54–78.
- [13] Hilary Putnam, *Models and Reality*, The Journal of Symbolic Logic (1980), 464–482.
- [14] Agustin Rayo and Steve Yablo, *Nominalism through de-nominalization*, Nous **35** (2001), no. 1, 74–92, Pre-print available online at <http://web.mit.edu/arayo/www/>.
- [15] Agustin Rayo and Gabriel Uzquiano (eds.), *Absolute Generality*, Oxford University Press, Oxford, 2006.
- [16] Ian Rumfitt, 'Yes' and 'No', Mind **109** (2000), no. 436, 781–823.
- [17] ———, *The Boundary Stones of Thought : an essay on the philosophy of logic*, Oxford University Press, Oxford, 2015.
- [18] Bertrand Russell, *On Denoting*, The Philosophy of Language (A.P. Martinich, ed.), Oxford University Press, New York and Oxford, third ed., 1996, pp. 199–207.
- [19] Stewart Shapiro, *Foundations without Foundationalism*, Oxford University Press, Oxford, 1991, Oxford Logic Guides, Number 17.
- [20] Robert J. Williams, *Quantificational indeterminacy revisited*, 2013.
- [21] Ludwig Wittgenstein, *On Certainty*, Basil Blackwell, Oxford, 1975, Translated by G.E.M. Anscombe and Denis Paul.
- [22] Crispin Wright, *On quantifying into predicate position : Steps towards a new(tralist) perspective*, Mathematical Knowledge (Michael Potter, ed.), Oxford University Press, Oxford, 2007.

*