

The *Metaontology* of Abstraction

Bob Hale and Crispin Wright

§1 We can be pretty brisk with the basics. Paul Benacerraf famously wondered¹ how any satisfactory account of mathematical knowledge could combine a face-value semantic construal of classical mathematical theories, such as arithmetic, analysis and set-theory—one which takes seriously the apparent singular terms and quantifiers in the standard formulations—with a sensibly *naturalistic* conception of our knowledge-acquisitive capacities as essentially operative within and subject to the domain of causal law. The problem, very simply, is that the entities apparently called for by the face-value construal—finite cardinals, reals and sets—do not, seemingly, participate in the causal swim. A refinement of the problem, due to Field, challenges us to explain what reason there is to suppose that our basic beliefs about such entities, encoded in standard axioms, could possibly be formed reliably by means solely of what are presumably naturalistic belief-forming mechanisms. These problems have cast a long shadow over recent thought about the epistemology of mathematics.

Although ultimately Fregean in inspiration, Abstractionism—often termed ‘neo-Fregeanism’—was developed with the goal of responding to them firmly in view. The response is organised under the aegis of a kind of linguistic—better, propositional—‘turn’ which some interpreters, including the present authors, find it helpful to see as part of the content of Frege’s Context principle. The turn is this. It is not that, *before* we can understand how knowledge is possible of statements referring to or quantifying over the abstract objects of mathematics, we need to understand how such objects can be given to us as objects of acquaintance or how some other belief-forming mechanisms might be sensitive to them and their characteristics. Rather we need to tackle directly the question how propositional thought about such objects is possible and how it can be knowledgeable. And this must be answered by reference to an account of how meaning is conferred upon the ordinary statements that concern such objects, an account which at the same time must be fashioned to cast light on how the satisfaction of the truth-conditions it associates with them is something that is accessible, in standard cases, to human cognitive powers.²

Abstraction principles are the key device in the epistemological project so conceived. Standardly, an abstraction principle is formulated as a universally quantified biconditional—schematically:

$$(\forall a)(\forall b)(\Sigma(a) = \Sigma(b) \leftrightarrow E(a,b)),$$

where a and b are variables of a given type (typically first- or second-order), ‘ Σ ’ is a term-forming operator, denoting a function from items of the given type to objects in the range of the first-order variables, and E is an equivalence relation over items of the given type.³ What

¹ in Benacerraf (1973)

² For some of our own efforts to develop and defend this approach, see Wright (1983), ch.1; Hale (1987), chs.1,7; Hale & Wright (2001), Introduction sect.3.1, Essays 5,6; Hale & Wright (2002)

³ More complex forms of abstraction are possible—see, for example, Hale (2000), p.107, where positive real numbers are identified with ratios of quantities, these being defined by abstraction over a four-term relation. One could replace this by a regular equivalence relation on ordered pairs of quantities, but this is not necessary—it is straightforward to extend the usual notion of an equivalence relation to such cases. It is also

is crucial from the abstractionist point of view is an epistemological perspective which sees these principles as, in effect, stipulative implicit definitions of the Σ -operator and thereby of the new kind of term formed by means of it and of a corresponding sortal concept. For this purpose it is assumed that the equivalence relation, E , is already understood and that the kind of entities that constitute its range are familiar—that each relevant instance of the right hand side of the abstraction, $E(a,b)$, has truth-conditions which are grasped and which in a suitably wide range of cases can be known to be satisfied or not in ways that, for the purposes of the Benacerrafian concern, count as unproblematic. In sum: the abstraction principle explains the truth conditions of Σ -identities as coincident with those of a kind of statement we already understand and know how to know. So, the master thought is, we can now exploit this prior ability in such a way as to get to know of identities and distinctions among the referents of the Σ -terms—entities whose existence is assured by the truth of suitable such identity statements. And these knowledge possibilities are assured without any barrier being posed by the nature—in particular, the abstractness—of the objects in question (though of course what pressure there may be to conceive of the referents of terms introduced by abstraction *as* abstract, and whether just on that account or for other reasons, is something to be explored independently⁴.)

§2 There are very many issues raised by this proposal. One might wonder, to begin with, whether, even if no other objection to it is made, it could possibly be of much interest merely to recover the means to understand and know the truth value of suitable examples of the schematised type of identity statement, bearing in mind the ideological richness displayed by the targeted mathematical theories of cardinals, real numbers and sets. The answer is that abstraction principles, austere as they may seem, do—in a deployment that exploits the collateral resources of second-order logic and suitable additional definitions—provide the resources to recover these riches—or at least, to recover theories which stand interpretation as containing them.⁵ There then are the various misgivings—for example, about “Bad Company” (differentiating acceptable abstraction principles from various kinds of unacceptable ones), about Julius Caesar (in effect, whether abstraction principles provide for a sufficient range of uses of the defined terms to count as properly explaining their semantic contribution, or justifying the attribution of reference to them), about impredicativity in the key (second-order) abstractions that underwrite the development of arithmetic and analysis, and about the status of the underlying (second-order) logic—with which the secondary literature over the last twenty-five years has mostly been occupied. For the purposes of the

possible—and possibly philosophically advantageous, insofar as it encourages linking the epistemological issues surrounding abstraction principles with those concerning basic logical rules—to formulate abstractions as pairs of schematic introduction- and elimination-rules for the relevant operator, corresponding respectively to the transitions right-to left and left-to right across instances of the more normal quantified biconditional formulation.

⁴ See Hale & Wright (2001), Essay 14, sect.4, for discussion of an argument aimed at showing that abstracts introduced by first-order abstraction principles such as Frege’s Direction Equivalence cannot be identified with contingently existing concrete objects.

⁵ At least, they do so for arithmetic and analysis. So much is the burden of Frege’s Theorem, so called, and the works of Hale and, separately, Shapiro. For arithmetic, see Wright (1983), ch.4; Boolos (1990) and (1998), pp.138-41; Hale & Wright (2001), pp.4-6; and for analysis, Hale (2000) and Shapiro (2000). The prospects for an abstractionist recovery of a decently strong set theory remain unclear.

present discussion, we assume all these matters to have been resolved.⁶ Even so, another major issue may seem to remain. There has been a comparative dearth of head-on discussion of the abstractionist's central *ontological* idea: that it is permissible to fix the truth-conditions of one kind of statement as coinciding with those of another—"kind" here referring to something like logical form—in such a way that the overt *existential* implications of the former exceed those of the latter, although the *epistemological* status of the latter, as conceived in advance, is inherited by the former. Recently however there have been signs of increasing interest in this proposal among analytical metaphysicians. A number of writers have taken up the issue of how to "make sense" of the abstractionist view of the ontology of abstraction principles, with a variety of proposals being canvassed as providing the 'metaontology' abstractionists need, or to which they are committed.⁷ We will here summarily review what we take to be the two leading such proposals—Quantifier-Variance and Maximalism—so far proposed to make sense of, or justify, the neo-Fregean use of abstraction principles. As is to be expected, each draws on background ideas and theses that require much fuller critical assessment than we can provide in the present space. But we can indicate briefly why we find neither direction of theorizing inviting, much less irresistible.⁸

§3 *Quantifier-Variance*⁹ is the doctrine that there are alternative, equally legitimate meanings one can attach to the quantifiers—so that in one perfectly good meaning of 'there exists', I may say something true when I assert 'there exists something which is composed of this pencil and your left ear', and in another, you may say something true when you assert 'there is nothing which is composed of that pencil and my left ear'. And on one view—

⁶ Since the 'noise' from the entrenched debates about Bad Company, Impredicativity, etc., is considerable, it may help in what follows for the reader to think in terms of a context in which a first order abstraction is being proposed—say Frege's well known example of the Direction Principle:

Direction (a) = Direction (b) iff. a and b are parallel

in which range of 'a' and 'b' is restricted to concrete straight lines—actual inscriptions, for example—and of the listed concerns, only the Caesar problem remains. The pure ontological problems about abstraction—if indeed they are problems—arise here in a perfectly clean form.

Previous discussions of ours of the more purely ontological issues are to be found in Wright (1983), chs.1-3; Hale (1987); Hale & Wright (2001), Essays 1-9 and 14.

⁷ In particular, Eklund (2006), Sider (2007), Hawley (2007), and Cameron (2007) all discuss the neo-Fregean abstractionist's (alleged) need for a suitable metaontology. As the italics in our title might forewarn, it is not, in our view, as clear as one could wish what 'metaontology' is supposed to be. One might naturally take it to apply to any general view about the character of (first-order) ontological claims or disagreements, or about how certain key terms (e.g. 'object', 'property', etc.) figuring in such claims or disputes are to be understood. But some recent writers seem to have had in mind something going significantly beyond this—roughly, some very general thesis about the metaphysical nature of the World which can be seen as underlying and somehow underwriting more specific ontological claims. It is beyond dispute that meta-ontology of the first sort is often useful and needed, and plausible that that there is call for a metaontology of abstraction in this sense. Certainly much of what needs to be said (including much of what we shall be saying in the sequel), if the character of abstractionist ontology is not to be misconstrued, could reasonably be regarded as metaontology of this sort. As will become clear as we proceed, however, we are sceptical about the demand for a metaontology of the second kind.

⁸ For a little more critical discussion of these views, and of the arguments for the claim that neo-Fregeans need to embrace one or other of them, see Hale (forthcoming)—fuller critical assessments are in preparation.

⁹ The name, but not the doctrine, comes from Eli Hirsch (Hirsch (2002). Hirsch finds the doctrine itself in various writings by Hilary Putnam.

perhaps not the only possible one—the general significance of this variation in quantifier meanings lies in its deflationary impact on ostensibly head-on disagreements about what kinds of objects the world contains: such conflicts may be less straightforward than they appear, and more a matter of their protagonists choosing to use their quantifiers (and other associated vocabulary, such as ‘object’) to mean different things—so that in a sense they simply go past each other. Its special interest for us lies in its application to the abstractionist use of Hume’s Principle. In particular, Ted Sider claims that neo-Fregeans need, or are well advised, to invoke quantifier-variance to make sense of the metaphor of “content-recarving”—specifically, the idea of the left-hand-sides of instances of abstraction principles as reconceptualisations of the right-hand-sides—and to block otherwise awkward demands for justification of the existential presuppositions he takes to attach to Hume’s Principle:

There are many equally good things one can mean by the quantifiers. If on one ‘there are numbers’ comes out false, there is another on which ‘there are numbers’ comes out true. ... ‘Reconceptualization’ means selecting a meaning for the quantifiers on which Hume’s Principle comes out true. (Sider (2007), p.207)

If there were a single distinguished quantificational meaning, then it would be an open possibility that numbers, directions, and other abstract are simply *missing* from existence in the distinguished sense of ‘existence’, even though we speak in a perfectly consistent way about them ... But if quantifier variance is true, then this is not an open possibility. (ibid, p.229)

This strikes us as a paradigm case of *ad obscurum per obscurius*—of explaining the (allegedly) obscure by appeal to what is (quite certainly) more obscure. Just what are the postulated variant quantifier meanings supposed to be? Of course, one can introduce any number of *restricted* quantifiers, but these clearly cannot be what the quantifier-variantist has in mind, since they just *aren’t* all equally good, when it comes to ontological disagreements. If, when you assert ‘there are no snakes’, you restrict your quantifier to creatures to be found in Ireland, you secure truth for what you say only by ignoring the existence of snakes elsewhere.

The quantifier-variantist owes us two things: he needs to explain why the allegedly different quantifiers which can all be expressed by the words ‘there are’ are all *quantifiers*; and he needs also to tell us how they *differ* in meaning. The first requires him to identify a common core of meaning for the quantifier-variants; the second requires him to tell us, in general terms, what the variable component is—what the dimension of meaning-variation is.

An obvious answer to the first is: they all share the same inferential behaviour—are subject to the same inference rules.¹⁰ As regards the second, it remains very difficult to see how the relevant dimension of variation could be other than the range of the bound variables (or their natural language counterparts)—so that (relevantly) different quantifier meanings differ just by being associated with different domains. But while this answer seems unavoidable, it seems in equal measure unfit for the intended purpose. For, on the one hand, we’ve already seen that the quantifier variantist’s allegedly different quantifiers can’t differ by being different *restrictions* of some other, perhaps unrestricted, quantifier—for then they wouldn’t all be ‘equally good’. But on the other, it is no good claiming that domain variation

¹⁰ See Hirsch (2002), p. 53, and Sider (2007), p.208, fn.12 (where Sider mentions this answer, but does not explicitly endorse it)

comes about through *expansion*, unless one can explain how that is supposed to work. The only obvious suggestion—that by introducing concepts of new kinds of objects (e.g. *mereological sum*, or *number*) we somehow enlarge the domain—is, in so far as it's clear, clearly hopeless. We cannot expand the range of our existing quantifiers by saying (or thinking) to ourselves: 'Henceforth, anything (any object) is to belong to the domain of our first-order quantifiers if it is an *F* (e.g. a mereological sum)'. For if *F*s do not already lie within the range of the initial quantifier 'anything', no expansion can result, since the stipulation does not apply to them; while if they do, then again, no expansion can result, since they are *already* in the domain.

Accordingly, it seems that the quantifier variantist faces a critical dilemma—either he proposes to explain how variant quantifiers differ in meaning in terms of domain variation, or he does not. If not, it is completely unclear what *other* kind of explanation he can plausibly give, since whether or not the domain includes *F*s is what, intuitively, precisely and exclusively determines the truth-value of 'there are *F*s'. But if the theorist goes for domain variation, he either breaks faith with his claim that the variants are equally good (if variation is explained in terms of restriction), or lapses into apparent incoherence (if it is explained in terms of expansion).

In fact, the situation is even worse, if the following simple train of thought succeeds. We've thus far left unscrutinized the suggestion that the shared meaning of variant quantifiers—say different versions of the existential quantifier—can consist in their being governed by the same inference rules, consistently with the distinctive quantifier variantist claim that the same quantificational sentences (syntactically individuated) embedded in the same language (again, demarcated purely syntactically) can be true when read with one quantifier meaning but false when read with another. Let us represent our two variant existential quantifiers as $\exists^1 x \dots x \dots$ and $\exists^2 x \dots x \dots$. Suppose $\exists^1 x A(x)$. Assume $A(t)$ for some choice of '*t*' satisfying the usual restrictions. Then by the introduction rule for \exists^2 , we have $\exists^2 x A(x)$ on our second assumption and so, by the elimination rule for \exists^1 , can infer $\exists^2 x A(x)$ discharging that assumption in favour of the first. We can similarly derive $\exists^1 x A(x)$ from $\exists^2 x A(x)$ ¹¹. Yet by hypothesis, one of the two is true, the other false. It follows that either the inference rules for \exists^1 , or those for \exists^2 , are *unsound*—and hence that that one set of rules or the other must fail to reflect the meaning of the quantifier it governs. The claim that the common core of quantifier meaning can be captured by shared inferential role is therefore unsustainable. It is quite unclear what better account the quantifier variantist can offer. In the absence of one, the very coherence of the view must be reckoned questionable.¹²

§4 We shall take *Maximalism* to be the thesis that whatever *can* exist *does*. If we restrict our attention to objects, it is the thesis that, for any sort or kind of objects *F*, if it is *possible* that *F*s should exist, they *do*.¹³ Matti Eklund, to whom the name 'maximalism' is due, claims that

¹¹ We here assume that both pairs of inference rules are harmonious—if both introduction rules are stronger than necessary in order for the corresponding elimination rules to be justified (say, because they are, bizarrely, subjected to the same restrictions as the usual universal quantifier introduction rules), the derivation suggested will break down. But this hardly offers a way out of the difficulty!

¹² We don't, of course, claim that this settles the issue. There are various moves a determined quantifier variantist might make—we can't chase them down in this paper, and can here only record our view (which we hope to defend more fully elsewhere) that none of them provides a satisfactory way around the problem.

¹³ Eklund gives a more complicated formulation (cf. Eklund (2006), p.102), but admits (p.117, note 23) it is not without problems. We think that is a massive understatement, and follow Sider (2007) and Hawley (2007) in

neo-Fregeans are actually committed to Maximalism, because it simply generalizes a principle which he labels *priority* and takes to underpin our argument for accepting the existence of numbers as objects. Since he gives no clear and explicit formulation either of *priority* or of the argument which is supposed to lead from it to maximalism, this claim is difficult to assess. On the face of it, it is straightforwardly false. The only relevant priority thesis to which we are committed¹⁴ (cf. Wright (1983), p.13-15); Hale (1987), pp.10-14) asserts the priority of truth and logical form over reference of sub-sentential expressions—it says, roughly, that it is sufficient for expressions functioning as singular terms to have reference to objects that they be embedded in suitable true statements. Since *actual*—not just *possible*—truth of the host statements is required, it is hard to see how this priority thesis—which is already completely general¹⁵—could possibly entail maximalism.

Others (Hawley (2007) and Sider (2007)) have considered whether maximalism, though not entailed by anything neo-Fregeans assert, is something they should embrace, as the best way to justify stipulating Hume's Principle as an implicit definition, given that its truth demands the existence of an infinity of numbers (or at least an ω -sequence of some sort). We shall explain later why we do not think we need a justification in this sense. Here our point is that even if we did, there would be ample reason not to look for it in this direction. Most obviously, maximalism denies the possibility of contingent non-existence, to which there are obvious objections: surely there could have been a £20 note in my wallet, even though there isn't? Attempts to mitigate the implausibility of the thesis by appeal to a distinction between existence in a logical sense (being something) and existing as a concrete object (being concrete) are vain, since there surely could have been abstract objects answering to certain descriptions even though no objects in fact do so—there surely could, for instance, have been a 63rd piano sonata by Haydn, even though in fact he wrote only (!

adopting the simpler formulation in the text. This formulation certainly doesn't put friends of maximalism at any disadvantage, as far as the points made here are concerned.

¹⁴ Contrary to what Eklund supposes (op.cit., p.100), there is certainly no commitment to what Hartry Field labelled (in Field (1984)) the 'strong priority thesis' that 'what is true according to ordinary criteria really is true, and any doubts that this is so are vacuous'. As Wright (1990, sect. 2) points out, this rests on a simple misreading of his earlier statement:

... when it has been established, by the sort of syntactic criteria sketched, that a given class of terms are functioning as singular terms, and when it has been verified that certain appropriate sentences containing them are, by ordinary criteria, true, then it follows that those terms do genuinely refer. (Wright (1983), p.14)

The intended sense was that the relevant sentences must be found to be true. The point of the addition 'by ordinary criteria' was just to observe that in the arithmetical case, operating in accordance with the ordinary criteria for appraising such statements will not lead us astray. There was no claim that in general, going by our ordinary criteria cannot but lead to truth; nor was there any relaxation of the requirement that the relevant embedding statements be actually true. Hale (1987), p.11, is completely explicit on the point. In any case, if we had endorsed the (obviously unacceptable) 'strong priority thesis', it would be a complete mystery why we should take various kinds of scepticism about abstracta (including Field's own version of nominalism) to pose a significant challenge to our position (as we both do—see, for example, Wright (1983), ch.2, Hale (1987) chs. 4-6, and Hale (1994), and Hale & Wright (1994))—we could simply have dismissed them out of hand as merely vacuous doubts!

¹⁵ In the sense that it is not restricted to numbers, or even to abstract objects, but applies—as each of us emphasizes—to all objects of whatever kind. Eklund gives the impression that we failed to recognize the generality of the underlying principle. We didn't. Of course, we don't accept that it should be generalized in the way Eklund proposes.

62 of them. Neo-Fregeanism does best to avoid commitment to such an extravagant thesis if it can; and it can. In the remaining part of the paper, we will attempt to explain how.¹⁶

§5 The way abstractionism wants to look at abstraction principles makes two *semantic* presuppositions. The first is that the statements schematised on the left hand side are to be taken as having the syntactic form they seem to have—that of genuine identity statements configuring (complex) singular terms. In the case of Hume’s Principle, this is clearly a precondition of the proposed implicit definition working as intended—if what is to be defined is a term-forming operator, the context *must* be one in which terms formed by its means occur, and this means that we must take ‘=’ seriously as the identity predicate. The second is that, when they are so taken, their counterparts on the right hand side may legitimately be regarded as coinciding in their truth conditions. Thus what it takes for ‘ $\Sigma(a) = \Sigma(b)$ ’ to be true is exactly what it takes for *a* to stand in the *E*-relation to *b*, no more, no less—which of course is quite unproblematic until we add that the syntax of the former is indeed, as it appears, that of an identity statement, at which point the abstractionist may seem to have committed to the dubiously coherent idea that statements whose logical forms so differ that their existential commitments differ may nevertheless be (necessarily) equivalent.

There are just two foreseeable ways of avoiding the dubiously coherent idea. One is to drop the assumption that the explained identity-statements are to be construed in such a way that their truth requires that their ingredient terms refer. Identity is indeed sometimes so read that, for example, “Pegasus is Pegasus” expresses a truth, the non-existence of any winged horse notwithstanding. Since that is not the way the abstractionist proposes to

¹⁶ One other recent metaphysical foray on behalf of abstractionism deserves mention. Ross Cameron (2007) claims to offer a third way to ‘make sense’ of neo-Fregeanism: we should reject Quine’s well-known criterion of ontological commitment in favour of one based on ‘truth-maker theory’. His general idea is that ‘the ontological commitments of a theory are just those things that must exist to *make true* the sentences of that theory’; on his preferred version of truth-maker theory, the things that must exist to make a statement true can be a proper subset of the things over which it quantifies or to which it involves singular reference. So, for instance, he claims that ‘the (mereological) sum of *a* and *b* exists’ is made true by just *a* and *b*—i.e. in asserting this sentence to be true (or asserting its disquotation), we are ontologically committed to just the objects *a* and *b* (and not to their mereological sum). Like quantifier variance, Cameron’s proposal is intended to deflate ontological disputes—we can both assert the existence of mereological sums and yet be ontologically committed only to the things of which they are ultimately composed. As with both quantifier-variance and maximalism, we have space only to indicate the targets of our two principal misgivings.

First, then, it is crucially unclear how Cameron’s replacement criterion is supposed to be applied. How, in particular, are we to determine when fewer things are needed to make a statement true than it asserts, or implies, exist? It is a consequence of the account that a statement’s truth-value together with its logical form is at best a guide to what ‘exists’, not to the statement’s underlying ‘ontological commitments’. Yet we are given not the slightest clue how we are supposed to determine the latter. In rejecting Quine’s criterion, Cameron opens up a gap between a statement’s logical form and what would make it true. Since what makes a statement true presumably ensures that its truth-condition is met, logical form must be insufficient to fix truth-conditions. Even if this is coherent, it remains a complete mystery how, and by what, truth-conditions *are* fixed.

Second, since our ontological commitments, as normally understood, are to exactly those things our theories require us to believe to exist, Cameron’s proposal invites the objection that it simply changes the subject. Hoping, perhaps, to outflank this objection, he invokes a contrast between what exists and what *really* exists. But in the absence of any clear account of what’s required for *real* existence, this makes no progress and merely invites a reformulation of the objection. Further, it may lead one to doubt that—for all his protestations to the contrary—Cameron’s third way really is a *third* way at all, rather than a misleadingly presented version of quantifier variantism. To be sure, he makes no (overt) claim about variant *quantifier* meanings; but we are, in effect, being invited to multiply meanings of ‘exist’, which comes to near enough the same thing.

construe identity statements, nor anything germane to the project more generally of reconciling a face-value ontology of mathematics with plausible epistemological constraints, we set it aside. The other is the abstractionist's actual view: the existential commitments of the statements which the abstraction pairs together *are* indeed the same—and hence the right-hand side statements, no less than the Σ -identities, implicate the existence of Σ -abstracts while containing no overt reference to them.

Now, this is not *per se* a problematic notion. That it is not is easily seen from two nearby cases:

- (i) The parents of A are the same as the parents of B iff A is a sibling of B
- (ii) A's MP is identical to B's MP iff A and B are co-constituents¹⁷

In each of these, the truth-conditions of a type of statement configuring a certain kind of complex term coincide with those of a type of statement which does not. And in each of them, as in abstraction principles proper, the latter type of statement affirms an equivalence relation on entities of a certain kind while the former affirms a related identity. Thus these biconditionals schematise a range of statements where the truth of the right-hand sides suffices for the truth of the left-hand sides, but where the former involve no overt reference to the denotata of complex terms occurring in the latter even though the truth of the type of statement schematised on the left hand side does involve successful reference to such entities. This phenomenon, then, is not peculiar to abstraction nor, as far as it goes, should it give rise to concern.

However, there are of course disanalogies. Two, related and very immediate, are these: in (i) and (ii) there is no question of using a prior understanding of the right-hand side to impart an understanding of the concept of the kind of thing (parents, MPs) denoted by the distinctive terms on the left. On the contrary one who understands the right-hand sides must *already* have that concept: you don't understand what siblings are—how two creatures have to be related in order to be to be siblings—unless you know what parents are.¹⁸ And you don't understand what it is for two people to live in the same constituency unless you know that constituencies are the areas MPs represent and what MPs are. So there is no analogue of the right-to-left epistemological priority claimed by abstractionism for actual abstraction principles. Second, the range of entities that constitute the domain of reference for the terms occurring in the left hand sides of instances of (i) and (ii) goes no further than the field of the equivalence relation on the right hand side: parents and MPs are *people*, and it is people who constitute the (relevant) field of the siblinghood and co-constituency relations. These principles deploy means of reference not to a novel kind of thing, but back into a prior ontology.

So, while (i) and (ii), suitably understood, show that there need be no problem about combining the two semantic presuppositions that abstractionism needs—a face-value, existentially committal reading of the terms occurring on the left-hand sides together with sameness of truth-condition across the biconditional—there are at least two salient differences between (i) and (ii) and abstraction principles proper. The former, but not the

¹⁷ The example is due to Sullivan and Potter (1997). In order to vouchsafe coincidence in truth-conditions, one could rule that people can be co-constituents only if they are somebody's constituents. This would still allow us to speak of unrepresented constituencies—but their inhabitants would be only potential, not actual, constituents.

¹⁸ We prescind from any complications occasioned by the theoretical possibility of laboratory synthesised and fissioned gametes.

latter, are both referentially and conceptually conservative. Even if the two semantic presuppositions are unproblematic under those two conditions, it is accordingly another question whether they remain so when the two conditions lapse. And the lapse of these conditions is just what is most distinctive about the process of abstraction: it is of the essence of the abstractionist proposal that abstraction principles be both conceptually and referentially non-conservative.

§6 Why might someone think that there *is* a special problem about stipulative identity of truth-conditions in the case of conceptually and referentially non-conservative principles of the relevant general form? Since sceptics about abstraction have not, to our knowledge, articulated their dissatisfaction in a manner responsive to exactly this stage-setting, we have to speculate. But it seems that such dissatisfaction might come in weaker and stronger forms.

The stronger form would rest upon the assumption of a certain *transparency* in the relationship between the understanding of a certain kind of statement and the nature of the states of affairs—the relevant kind of *truth-conferrer*¹⁹—whose obtaining suffices for the truth of a statement of that kind. This transparency would involve that there will not be more, so to speak, to such truth-conferrers than is manifest in the conception of their form and content that is part and parcel of the ordinary understanding of the statements concerned. Since the conceptual non-conservativeness of abstraction principles precisely involves that someone can possess a full, normal understanding of the right-hand side type of statement without any inkling of the sortal concept the abstraction aims to introduce, let alone a recognition of the entailment of the existence of instances of that sortal putatively carried by such a right-hand side statement, transparency is violated.

The transparency principle so characterised is, however, surely unacceptable. At any rate, it is inconsistent with acknowledging any form of distinction between the conception of its form and content that is available to someone who possesses a normal, theoretically unrefined understanding of a statement, and the conception of its form and content that would feature in a theoretically adequate account of its deep semantic structure (logical form.) Whatever the pressures—considerable, of course—to admit such a notion of logical form are therefore reasons to reject transparency as formulated. But this doesn't really address the stronger dissatisfaction. Abstractionism, after all, is not saying that the overt syntactic structure of the right-hand side statements masks their real logical form, which is better portrayed by the left-hand side statements. There is nothing in abstractionism that is intended to war with the idea that the overt syntactic structure of the right-hand side statements is a fully adequate reflection of their logical form. Suppose we therefore refine the statement of the transparency principle to something like this: there will not be more, so to speak, to the truth-conferrers for a given kind of statement than is manifest in the conception of their form and content that is conditioned by an appreciation of the deep semantic structure (logical form) of the statements concerned. Then the tension remains.

The question, though, is why even the refined principle should seem compelling. Logical form is, plausibly, theory-determined—how best to think of the deep semantic structure/logical form of a given type of statement is a matter of what structure is assigned to it by best semantic theory. Such theory is subject, familiarly, to all the constraints to which empirical theorising in general is subject, and then some more, peculiar to its special

¹⁹ We deliberately avoid the term “truth-maker” to avoid any unwanted implicature of assumptions from that literature.

project—it should for instance be capable of explaining speakers’ competence with the parsing of novel utterances and consistent with the learnability of the language under study, and it should explain the inferences that speakers take to be immediately admissible. But what there is not is any constraint of making a match between assigned semantic structures in general and the (structural) nature of the relevant truth-conferrers. If there were, how on Earth would we know how to set about complying with it? More generally, what reason is there to think that semantic theories which count as best by the constraints recognised by semanticists and linguists will thereby also satisfy the—as it seems—additional and independent *metaphysical* constraint of assigning logical forms to statements in the target language that somehow mirror the structure and ontology of the associated truth-conferrers? The revised transparency principle seems to be drawing on something akin to the spirit of a Tractarian ontology of structured facts or states of affairs, which get to count as truth-conferrers for statements by, as it were, matching—being isomorphic to—the semantic structure of those statements. Such a metaphysics of linguistic representation and truth may well jar with abstractionism, or force it into implausible claims—for instance, that right hand side statements do indeed misrepresent the ontology of the associated states of affairs (for how in that case did those states of affairs get to be associated with those statements in the first place?) But if the objection to abstractionism is that it is incompatible with a Tractarian theory of meaning, that seems more interesting than damaging.

It might be suggested, though, that there is no need for the transparency principle to inflate into commitment to a Tractarian metaphysics of truth and content. The thought can be more simple: that absent any reason to draw a distinction between the overt and deep semantic structure of a kind of sentence, there can be no justification for ascribing a kind of ontological or ideological commitment to them which exceeds what is manifest in their overt structure. This principle certainly seems well suited to clash with abstractionism at minimal metaphysical expense. But how plausible is it in turn? Let E be an equivalence relation. Then if $E(a,b)$, it follows that $(\forall x)(E(x,a) \leftrightarrow E(x,b))$, and conversely. Yet the latter involves both ideological and, arguably, ontological commitments that go unreflected in the surface structure of the former—in particular, to the concept of universal quantification, and to the operation that constitutes it. If there is a well-motivated transparency principle—a principle insisting on the transparency of the relation between the logical form of a sentence and its ideological and ontological commitments—that is uncompromised by this example but not by the relationship claimed by abstractionism between ‘ $E(a,b)$ ’ and ‘ $\Sigma(a) = \Sigma(b)$ ’, it is by no means clear what it may be. In general, a priori necessarily equivalent statements may deploy differing conceptual resources without there being any well-motivated suggestion that either or both involve a mismatch between overt and deep semantic structure. It is hard to see how the proposed transparency principle can survive this observation.

§7 The stronger reservation, based on some form of transparency constraint, was announced above as contrasting with a possible weaker one. The weaker reservation is to the effect not that abstractionism violates some basic metaphysical principle about representation but merely that there are some questions, metaphysical and epistemological, that need answering before abstractionism should be considered to be a competitive option. An example of such a metaphysical question would, be:

(M) What does the world have to be like in order for (the best examples of) abstraction to work?

And an associated epistemological question would be:

(E) How do we know—what reason have we to think—that the transition, right to left, across the biconditional in instances of (the best examples of) abstraction is truth preserving?²⁰

Before we proceed further, it is worth pausing to register the point that it is a substantial issue to which questions exactly, arising in this vicinity abstractionism owes developed, satisfactory answers before it has any claim to credibility. The proposal is that we may implicitly define the meanings of abstraction operators by laying down abstraction principles — that is, by stipulatively identifying the truth-conditions of instances of their left-hand sides with those, as already conceived, of instances of their right-hand sides. One very broad class of issues concern implicit definition of this general character — the stipulative association of the truth-conditions of two syntactically differing sentence-types, one of which (but not the other) configures novel vocabulary— rather than abstraction specifically. Can such a ploy (ever) succeed in attaching meaning to the novel vocabulary? Can the (biconditional) vehicle of the implicit definition (ever) be understood and known to be true (a priori) just on the basis of an intelligent reception of the stipulation? We have argued elsewhere²¹ for positive answers to these questions. There are of course a number of qualifications that need to be entered since such implicit definitions, like any explanations, may go wrong. In the discussion just cited, a variety of conditions are proposed —including forms of conservativeness, harmony, and generality—as necessary and (tentatively) sufficient for an implicit definition of this general character to be both meaning-conferring and knowledge-underwriting. Our position, however, is that, in any particular case, the satisfaction of these conditions is a matter of *entitlement*²². It is not for the would-be user to show that his implicit definition is in good standing by the lights of these, or related, conditions before he is justified in putting the implicit definition in question to work in knowledge-acquisitive projects— any more than he needs to show that his perceptual apparatus is functioning properly before he is justified in using it to acquire knowledge about his local perceptible environment. Implicit definition is *default* legitimate practice — although, again, subject to defeat in particular cases—and particular such principles proposed, together with our claims to knowledge of their deductive progeny, are to be regarded as in good standing until shown to be otherwise.

On this view, abstraction principles, once taken as legitimate instances of this genre of implicit definition, don't stand in need of justification. If the thrust of question (E) is simply an instance of the general form of question: what reason do we have to think that the vehicle of a proposed implicit definition is true (and therefore meaning-conferring), then our answer is that no answer is owing—though of course one may still, as a theorist, interest oneself in the satisfaction of the relevant conditions in the particular case.

However, that need not be the thrust of question (E). In insisting that something needs to be said 'up front' to make out an abstraction's right to asylum, as it were, the critic's focus of concern may be, not with implicit definition in general, but with the credentials of abstraction principles in particular to be classed as such and so to inherit the benefits of that status. We should not, on this suggestion, propose such principles—even in cases which there is no reason to suppose will trip up over other constraints—before the very practice of

²⁰ We will hence generally omit the parenthetical qualification “the best examples of”. But except where stated otherwise, it is to be understood.

²¹ In Hale & Wright (2000).

²² In the sense of Wright (2004a) and (2004b)

abstraction as a legitimate form of implicit definition has been authenticated. And it is this, so it may be suggested, that requires the development of satisfactory answers to (M) and (E) and related questions. If it *looks* as if the truth of abstraction principles may turn on substantial metaphysical hostages, or as if there are special problems about knowing that they are true, or can be stipulated to be true, this appearance needs to be disarmed before the abstractionist can expect much sympathy for his proposals.

We are here content to defer to this concern. Certain of the special features of abstraction principles—in particular their role in the introduction of a conceptually novel ontology—do suggest that some special considerations need to be marshalled, not to show that particular cases are in good standing, but to shore up their assimilation to the general run of implicit definitions for abstractionist purposes. Still, there is an important qualification to enter here concerning what exactly it is that we are agreeing to try to do—for very different conceptions are possible of what it is to give a satisfactory answer to question (E) in particular; that is, to justify the thought that a good abstraction is truth-preserving, right-to-left. One such conception which we reject is, we venture, implicit in maximalism. This conception has it, in effect, that it is, in some sense, *possible*²³—something we have initially no dialectical right to discount—for any abstraction to fail right-to-left unless some relevant kind of collateral assistance is forthcoming from the metaphysical nature of the world. There are, that is to say, possible situations—in some relevant sense of ‘possible’—in which an abstraction which actually succeeds would fail, even though conceptually, at the level of explanation and the understanding thereby imparted, everything is as it is in the successful scenario. Hence in order to make good that the right-to-left transition of an otherwise good abstraction is truth-preserving, argument is needed that some relevant form of metaphysical assistance is indeed provided. This is, seemingly, the way those who have advocated maximalism as neo-Fregeanism’s best course are thinking about the issue. The ‘possible’ scenario would be one in which not everything that could exist does exist—in particular, the denoted abstracts do not exist. And the requisite collateral consideration would be that this ‘possibility’ is not a genuine possibility—because maximalism is true (and is so, presumably, as a matter of metaphysical necessity.) Although the idea is by no means as clear as one would like, we reject this felt need for some kind of collateral metaphysical assistance. The kind of justification which we acknowledge *is* called for is precisely justification for the thought that no such collateral assistance is necessary. There is no hostage to redeem. A (good) abstraction *itself* has the resources to close off the alleged (epistemic metaphysical) possibility. The justification needed is to enable—clear the obstacles away from—the recognition that the truth of the right-hand side of an instance of a good abstraction is *conceptually* sufficient for the truth of the left. There is no gap for metaphysics to plug, and in that sense no ‘metaontology’ to supply. This view of the matter is of course implicit in the very metaphor of *content* recarving. It is of the essence of abstractionism, as we understand it—but, interestingly, if we have the proposal right, it is essential to the quantifier-variantist ‘rescue’ of abstractionism as well.²⁴

²³—perhaps this modality is: *epistemically* [*metaphysically possible*]!

²⁴—since on the quantifier variantist line here, or so we take it, the conservation in truth conditions, right-to-left, across a good abstraction is ensured purely by so understanding the quantification in the three possible existential generalisations of the left-hand side that the right-hand side suffices for their truth at a purely conceptual level, without collateral metaphysical assumption. It is a substantial thesis that it is possible to do this. But it is a thesis about what meanings—concepts—there are, not about the World of the metaphysician.

§8 Question (M) was: What does the world have to be like in order for (the best examples of) abstraction to work? A short answer is that it is at least necessary that the world be such as to verify their *Ramsey sentences*: the results of existential generalisation into the places occupied by tokens of the new operators. So for any particular abstraction,

$$(\forall a)(\forall b)(\Sigma(a) = \Sigma(b) \leftrightarrow E(a,b))$$

the requirement is that this be true:

$$(\exists f)(\forall a)(\forall b)(f(a) = f(b) \leftrightarrow E(a,b))$$

More generally, the minimum requirement is that each equivalence relation suitable to contribute to an otherwise good abstraction be associated with at least one function on the members of its field that takes any two of them to the same object as value just in case they stand in the relation in question.

A world in which abstraction works, then—a world in which the truth values of the left-and right-hand sides of the instances of abstraction principles are always the same—will be a world that displays a certain ontological richness with respect to functions. Notice that there is no additional requirement of the existence of values for these functions. For if ‘ Σ ’ is undefined for any element, c , in the field of E , then the instance of the abstraction in question, $\Sigma(c) = \Sigma(c) \leftrightarrow E(c, c)$, will fail right-to-left. This brings us sharply to the second question, (E). To know that the transition right to left across an otherwise good abstraction principle is truth-preserving, we need to know that the equivalence relation in question is indeed associated with a suitable function. Here is George Boolos worrying about the latter question in connection with Hume’s principle (“octothorpe” is a name of the symbol, ‘#’, which Boolos uses to denote the cardinality operator, “the number of...”):

...what guarantee have we that there is such a function from concepts to objects as [Hume’s Principle] and its existential quantification [Ramsey sentence] take there to be?

I want to suggest that [Hume’s Principle] is to be likened to “the present king of France is a royal” in that we have no analytic guarantee that for every value of “F”, there is an object that the open definite description²⁵ “the number belonging to F” denotes....

Our present difficulty is this: just how do we know, what kind of guarantee do we have, why should we believe, that there is a function that maps concepts to objects in the way that the denotation of octothorpe does if [Hume’s Principle] is true? If there is such a function then it is quite reasonable to think that whichever function octothorpe denotes, it maps non-equinumerous concepts to different objects and equinumerous ones to the same object, and this moreover because of the meaning of octothorpe, the number-of-sign, or the phrase “the number of.” But do we have any analytic guarantee that there is a function which works in the appropriate manner?

Which function octothorpe denotes and what the resolution is of the mystery how octothorpe gets to denote some one particular definite function that works as described are questions we would never dream of trying to answer.²⁶

Boolos undoubtedly demands too much when he asks for “analytic guarantees” in this area. But the spirit of his question demands an answer that at least discloses some reason to believe

²⁵ The reader should note Boolos’ ready assimilation of “the number belonging to F” to a definite description – of course, it *looks* like one. But the question whether it is one depends on whether it has the right *kind* of semantic complexity. The matter is important, and we will return to it below.

²⁶ Boolos (1997), p. 306

in the existence of a function of the relevant kind. So: what, in general, is it to have reason to believe in the existence of a function of a certain sort?

If, as theorists often do, we think of functions as *sets*—sets of pairs of argument-tuples, and values—then standard existence postulates in set theory can be expected to provide an answer to Boolos's question in a wide range of cases: there is whatever reason to believe in the existence of the functions required by abstraction principles as there is to believe in the existence of the relevant sets. But that is, doubly, not the right kind of way to address the issue for the purposes of abstractionism. For one thing, abstractionism's epistemological objectives require that the credibility of abstraction principles be *self-standing*. They are not to (need to) be shored up by appeal to independent ontological commitments—and if the abstractionist harbours any ambition for a recovery of set-theory, especially not by appeal to a prior ontology of sets. However there is a deeper point. Abstraction principles purport to introduce *fundamental* means of reference to a range of objects, to which there is accordingly no presumption that we have any prior or independent means of reference. Our conception of the epistemological issues such principles raise, and our approach to those issues, need to be fashioned by the assumption that we may have—indeed there may be possible—no prior, independent way of conceiving of the objects in question other than as the values of the relevant function. So when Boolos asks, what reason do we have to think that there is any function of the kind an abstraction principle calls for, it is to skew the issues to think of the question as requiring to be addressed by the adduction of some kind of evidence for the existence of a function with the right properties that takes elements from the field of the abstractive relation as arguments and objects of some independently available and conceptualisable kind as values. If the best we can do, in order to assure ourselves of the existence of a relevant function or, relatedly, of the existence of a suitable range of objects to constitute its values, is to appeal to our independent ontological preconceptions—our ideas about the kinds of things we take to exist in any case—then our answer provides a kind of assurance which is both insufficient and unnecessary to address the germane concerns: insufficient, since independent ontological assurance precisely sheds no light on the real issue—viz. how we can have reason to believe in the existence of the function purportedly defined by an abstraction principle, and accordingly of the objects that constitute its range of values, when proper room is left for the abstraction to be fundamental and innovative; unnecessary since, if an abstraction can succeed when taken as fundamental and innovative, it doesn't need corroboration by an independent ontology.

§9 Let us therefore refashion question (E) as follows:

(E') How do we know—what reason have we to think—that the transition, right to left, across the biconditional instances of abstraction principles is truth preserving, once it is allowed that the means of reference it introduces to the (putative) values of the (putatively) defined function may be fundamental, and that no antecedently available such means may exist?

An answer to (E') in any particular case must disclose a kind of reason to believe in the existence of a suitable function which originates simply in resources provided by the abstraction principle itself, and independent of collateral ontological preconceptions. Those resources must pertain to what an abstraction can accomplish as an implicit definition of its definiendum—the new term forming operator. Allow, at least *pro tem*, that an abstraction principle, laid down as an implicit definition of its abstraction operator, may at least succeed

in conferring on it a *sense*. So much is tacitly granted by Boolos when he writes in the passage quoted above:

If there is such a function then it is quite reasonable to think that whichever function octothorpe denotes, it maps non-equinumerous concepts to different objects and equinumerous ones to the same object, and this moreover because of the meaning of octothorpe But do we have any analytic guarantee that there is a function which works in the appropriate manner?

For it is, after all, by its stipulated role in the relevant version of Hume's principle that the meaning of octothorpe is fixed. So the question is: what, for functional expressions—one standard practice calls them *functors*—needs to be in place in order for possession of sense to justify ascription of reference?

For Frege, functors are to be conceived as an instance of the more general category of *incomplete* expressions: expressions whose 'saturation' by a singular term results in a further complex, object-denoting term. So let's ask in the first instance: is there something general to be said about what justifies the ascription of reference to an incomplete expression? And what, in particular, is the role played by sense? We are not, in posing this question, taking it as uncontroversial that incomplete expressions as a class should be credited with a potential for reference as well as sense. The question is rather: for a theorist not already inclined—because of nominalist scruple or whatever reason—to deny reference to incomplete expressions across the board, what should justify the ascription of reference in any particular case?

Let's try the case of simple predicates. Take it that in order to assign a sense to a predicate, it suffices to associate it with a sufficiently determinate satisfaction-condition: to fix under what circumstances it may truly, or falsely, be applied to an item in some appropriate assigned range. And take it that the question whether it has a reference amounts to whether we have thereby succeeded in associating it with a genuine property. Then there is a contrast between two broad ways of taking the question. On one way of taking it, the relevant notion of genuine property is akin to that in play when we conceive it as a non-trivial question whether any pair of things which both exemplify a certain set of surface qualities—think, for example, of a list of the reference-fixers for 'gold' given in a way independent of any understanding of that term or an equivalent—have *a property in common*. When the question is so conceived, the answer may be unobvious and negative: there may be 'fool's' instances of a putative natural kind, or there may even just be no common kind underlying even normal cases of presentation of the qualities in question. Theorists who think of *all* properties in this way—sometimes termed "sparse" theorists—will recognise a gap between a predicate's being in good standing—its association with well-understood, feasible satisfaction conditions—and its hitting off a *real worldly property*. However this conception stands in contrast with that of the more "abundant" theorist, for whom the good standing, in that sense, of a predicate is *already* trivially sufficient to ensure the existence of an associated property, a (perhaps complex) *way of being* which the predicate serves to express.²⁷ For a theorist of the latter spirit, predicate sense will suffice, more or less,²⁸ for

²⁷ The terminology of abundant and sparse properties originates in Lewis (1986). The general distinction is in Armstrong (1979). See also Bealer (1982) and Swoyer (1996) For a useful overview see Mellor and Oliver (1997).

predicate reference. The sparse theorist, by contrast, will view the relationship as very much akin to that which obtains in the case of complex singular terms: the sense of—a satisfaction condition of—a predicate will aim at an underlying property fit to underwrite in some appropriate manner the capacity of an object to meet that satisfaction condition, and the predicate will have reference only insofar as there is indeed such a property provided by the world. Whether that is so will then depend in turn on one's metaphysics of worldly properties.²⁹

It is clear enough that the two conceptions of property need not be in competition: it is perfectly coherent to work with both simultaneously. What do compete, however, are the two associated views of predicate reference since no-one inclined to admit both conceptions of property is going to wish to maintain, presumably, that in the case when a predicate is associated with properties of both kinds, it somehow divides its reference over them both, or something of the sort. The natural compatibilising view will be, rather, that it is for the abundant properties to play the role of *bedeutungen* in semantic theory, and the sparse ones to address certain metaphysical concerns.³⁰

For predicates at least, then, there is a good conception of reference such that to confer a sense is, more or less, to confer a reference. Nor, arguably, is the point restricted to predicates. Consider the category of sentential connectives. And suppose that we conceive, much in the spirit of abstractionism, that we may fix the sense of a connective by stipulatively associating formulae in which it is the principal operator with certain natural deductive introduction and elimination rules. Once again, there are of course, as Prior's classic example of 'tonk' shows, ways in which this process can go wrong. But suppose, as we are doing in the case of abstraction principles, that we are concerned with the best kind of case, where no triviality, inconsistency or other form of disharmony intrudes and the resulting inferential practice runs smoothly and without singularity. Won't we feel we understand the connective in question in such circumstances? And won't the resulting plausibility of the contention that a sense has been fixed for it go hand in hand with the belief that there is an operation in good standing that it serves to express? Here too, then, conferral of sense seems, *ceteris paribus*, sufficient for conferral of reference. We can transpose Boolos's question—

If there is such an operation then it is quite reasonable to think that whichever operation [the relevant connective] denotes, it is an operation which complies with the specified introduction and elimination rules and this moreover because of the meaning of [the relevant connective]... But do we have any analytic guarantee that there is an operation which works in the appropriate manner?

—to this context as indicated. But with connectives, as with predicates, there seems clear room for an abundant view whereby to fashion a straightforward answer: that there is a statement-forming operation associated with any connective of which it is possible to succeed in imparting a satisfactory understanding by natural deductive characterisation of its

²⁸ "More or less" because the abundant theorist may still want to deny reference to certain significant predicates—for instance, those associated with inconsistent satisfaction conditions, or which embed empty terms ("That car is *my dog's favourite colour*").

²⁹ For example, versions of both Aristotelian and Platonic conceptions of property are consistent with sparseness. For discussion of varieties of sparseness see Schaffer (2004).

³⁰ Cf. Schaffer *op. cit.*

inferential role, and that this operation may be conceived as the reference of the expression in question.

Do these ideas suggest a way of responding to Boolos's question, and thence to question (E'), for the target case: the term-forming operators introduced on the left hand side of instances of abstraction principles? Well, connectives, like abstraction operators, denote functions of a certain kind: functions, we can suppose, from statements to statements. So for one in sympathy with the ideas just canvassed, they provide a precedent for the thought that the existence of a function may be settled just by conferring a sense upon a functor. Still, there remain evident disanalogies with the case of predicates and an abundant conception of properties. Any predicate associated with a (sufficiently) determinate satisfaction condition is, *ceteris paribus*, assured of reference to an abundant property. But it seems there should be room for a would-be functor to have sufficient sense to be associated with a determinate condition on any function that is to qualify as presented by it and yet fail to present one. Setting aside any issue about the existence of a range of suitable *arguments* for the purported function in question—as we may in the case of abstraction principles—there are two ways this can happen. One is if a relation can meet the condition in question and yet not be functional—not *unique*. And the other is precisely if there are no objects suitable to constitute *values* for the purported function in question.

There is, notably, no problem with either of these conditions in the case of the connectives. The conferral of sense upon a connective precisely ensures that there will be a statement formed whenever the connective is applied to an appropriate n-tuple of sentences each of which possesses a prior sense. So here sense alone ensures the existence of a value for every suitable n-tuple of arguments. And uniqueness is ensured by functionality of semantic composition: the principle that the content of a semantically complex expression is a function of that of its semantically relevant constituents and mode of composition. The case is, however, special. Clearly, the point doesn't carry for functional expressions as a class. With connectives, both the arguments and the values of the operation/function for which the connective stands are intensional entities (statements, propositional contents, or some such). This is why composition guarantees both existence and uniqueness. This contrasts with the general run of functions, whose arguments and values are typically non-intensional entities. The sense assigned to a putative functor may precisely carry sufficient information to enable us to show that the associated relation is not many-one (nor one-one) or that it fails to correlate the intended range of arguments with anything at all. Functors generally may have sense yet fail to present any function—so fail to have reference—if these conditions, of uniqueness and existence, are not met.

The question, accordingly, is whether a significant doubt is possible about whether they are met in the case of the functors introduced by (the best) abstractions. Might uniqueness be open to reasonable doubt in such a case? Here is a consideration that strongly suggests not. In order to entertain such a doubt, one needs to associate the relevant functor—' Σ '—with an underlying *relation* and then to think of ' $\Sigma(a)$ ' as purporting to denote what is the only object so related to *a*. Uniqueness fails just when there more than one such object. But is there in general any conception of such a relation somehow conveyed as part of the sense that is attached to an abstraction operator by its implicit definition via the relevant abstraction principle? Take the case of Hume's principle and the associated cardinality operator, glossed as "the number of". In order to raise a meaningful doubt about uniqueness, we need to identify an associated relation such that the sense of "the number of Fs" may be conceived of as grasped *compositionally*, via grasping this relation plus the presumption of

uniqueness incorporated in the article. The issue of uniqueness will be the issue of the many-oneness of this relation,—something which might ideally admit of proof. It is very doubtful however whether there is any good reason to think of the sense assigned to the cardinality operator by Hume’s principle as compositional in this particular way³¹. And if not—if the operator is best conceived as semantically atomic—then there is no scope for a significant doubt about uniqueness of reference, since there is no associated condition which more than one item might satisfy.³²

It is, on the other hand, by no means as evident that there is no room for a significant doubt about existence.³³ The abstraction operator refers (to a function) only if the singular

³¹ The issue is not uncontroversial. In MacFarlane (forthcoming), John MacFarlane, like Boolos above, canvasses the view that numerical terms having the surface form ‘the number of Fs’ are Russellian definite descriptions, presumed constructed using an underlying relational expression ‘x numbers the Fs’—so that a sentential context ‘A(the number of Fs)’, with the definite description having wide scope, gets paraphrased as ‘ $\exists!x(x \text{ numbers the Fs} \wedge Ax)$ ’. On this view, at least as MacFarlane presents it, numerical terms are not genuine singular terms at all but a kind of quantifier. One could still enquire whether the postulated numbering relation is functional—i.e. whether, for any F, there always exists a unique x which numbers F. This would now be a substantial question, both as regards existence and uniqueness. This is not the place for detailed criticism of MacFarlane’s proposal, to which we respond in our reply to his paper (Hale & Wright (forthcoming)). But it is worth briefly separating some issues. One, obviously, is whether MacFarlane’s proposal is viable at all. If Hume’s Principle works as an implicit definition in the way we propose, it defines a certain functor—the number operator—*directly*. There simply is no underlying relational expression, from whose sense that of the functor is composed. One can of course define a relational expression, ‘x numbers F’, to mean ‘ $x = Ny:Fy$ ’—but this relational expression is evidently compositionally *posterior* to the number operator. The question, for the viability of MacFarlane’s proposal, must therefore be whether ‘x numbers the Fs’ can be defined independently, without presupposing prior understanding of numerical terms. It is certainly not obvious that it can be. But even if it can be, the more important issue for present purposes is not whether one could introduce the number operator on the basis of an underlying relation, but whether one can, as we contend, introduce it as semantically atomic—if so, then there is, for the reasons noted in the main text, no scope for a significant doubt about *uniqueness* of reference for terms formed by its means.

³² Lest there be any misunderstanding, this concern needs sharply distinguishing from the concern about uniqueness raised by Harold Hodes in Hodes (1984). Hodes’ concern is based on the fact that one can, consistently with the truth of Hume’s Principle, permute the references of terms formed by means of the number operator, provided one makes compensating adjustments elsewhere (e.g. to the extension of the <-relation). Thus besides the ‘standard numberer’ which takes empty concepts to 0 as value, singly-instantiated concepts to 1, doubly-instantiated concepts to 2, and so on, there are many non-standard numberers—e.g. one which coincides with the standard numberer except in its values for empty and singly-instantiated concepts (1 and 0, respectively), compensating with a non-standard <-l relation which coincides with standard < except that we have $1 < 0$. Hodes grants, at least for the sake of argument, that the number operator, as introduced by Hume’s Principle, will denote a function—the trouble, he thinks, is that there is no unique, privileged such function that it can succeed in defining; rather, there are infinitely many such functions, between which it is powerless to discriminate. The problem is not that it is open whether “the number of” succeeds in picking out any operation whose values are, as required by functionality, unique but that it is unsettled whether it succeeds in picking out any unique such operation. This kind of doubt is not at issue in the text, and demands a quite different response. The crux is whether Hodes succeeds, as he claims, in demonstrating that a special, distinctively recalcitrant type of indeterminacy afflicts numerical terms as introduced by Hume’s Principle—i.e. that we have something worse than the kind of permutational indeterminacy that can be engineered for expressions of any type, and is not confined to those purporting reference to abstracta. See Hale (1987), pp.220-4 for some further discussion.

³³ To be sure, one kind of doubt about existence *is* pre-empted by the same point. There can be no doubt whether certain items stand in a relevant underlying relation to anything if there is no relevant underlying relation—if there is no prior relation R such that ‘the Φ of A’ is constrained to stand, if for anything, then for the unique B such that $R(A,B)$. But those anxious about the existential consequences of abstraction principles will probably not be quickly persuaded that any proper doubt about existence here has to assume this pattern.

terms it enables us to form refer (to objects.) What reason is there to think that (any of) these terms so refer?

To fix ideas, think of the routine ways in which one might satisfy oneself that *any* singular term refers. Suppose, for instance, you take it into your head to try to show that “Bin Laden” is the name of a real man, rather than, say, the focal point of an elaborate fiction, promulgated by the CIA. There are various courses of action you might undertake to try to settle the matter, at least to your own satisfaction. But ultimately, what you need to do is gather evidence which is arguably sufficient for the truth of an *identity statement*, ‘ $q = \text{Bin Laden}$ ’, for some ‘ q ’ whose reference to a real man is not in question. In this, ‘ q ’ might be a compendious definite description of the words and actions (“the man who said and did all of these things:....”) of an unquestioned real man; or it might be a token demonstrative for the robed, bearded figure standing before you at the entrance to a cave in the Tora-Bora mountain range and revealed only after many days blind-folded travelling on the back of a donkey. The point generally is that verification of the existence of a referent for a term N is verification of a statement of the form: $(\exists x)(x = N)$. And the premium method for doing that is to verify an identity, $q = N$, where the existence of a referent for ‘ q ’ is not in doubt.

But this model exactly presupposes, of course, that the term in question is not fundamental. What of the case when N is a term purporting to stand for a new kind of object for which it is understood that no anterior means of reference need exist in the language—so that it is a given that there need be no suitable ‘ q ’? The latter condition is but rarely satisfied, of course—at least if we assume the language to contain *demonstrative* means of reference—since it excludes that N refers to any kind of object capable of anchoring the attention well enough to attract demonstration (even if the user has only a partial grasp of the kind of object that is being demonstrated.) In these circumstances verifying that N refers cannot be a matter of verifying that it co-refers with any expression, even a demonstrative, whose reference is not in doubt. So what can it be?

The only possible answer appears to be that such a feat of verification must consist in verifying—if not an identity statement linking the term in question with another whose reference is assured—then *some* form or forms of statement embedding the term in question whose truth *requires* that it refer: a statement, or range of statements, in which the term in question occupies a *reference-demanding* position. Such will be afforded by provision of the means to verify some form of *atomic* statement configuring such terms. Identity contexts are one kind of atomic statement. So abstraction itself—as a characterisation of putatively canonical grounds for the verification of such identity contexts—supplies a paradigm means, indeed an example it seems of the only foreseeable broad kind of means, for accomplishing the assurance required.

Consider the resulting dialectical position. The challenge posed by (E’) was that before there can be sufficient reason to accept an abstraction as true, grounds are owing to think that a function suitable to witness its Ramsey sentence exists. Reason to believe in the existence of such a function depends in turn on reason to believe that the characteristic singular terms formed by means of the associated operator refer. Reason to believe that this is so has to consist in reason to take some associated range of statements that embed them in reference-demanding ways as true. And reason to do that presupposes a conception of a type of ground or grounds that would mandate so regarding members of that range of statements. The question is therefore: what conception of that kind is being presupposed when the demands of question (E’) are taken in a metaphysically anxious spirit? What will the anxious

metaphysician accept as grounds to regard members of some relevant such range of statements as true?

Well, we know what will *not* be accepted. A relevant type of ground is easily identified if we take it that abstraction represents a legitimate means of fixing the truth-conditions of one relevant kind of statement! In that case, there is no difficulty in returning a positive answer to the question, what grounds can we have to think that the new singular terms generated by an abstraction principle refer. So the metaphysically anxious question presupposes that abstraction *per se* cannot be taken to represent a legitimate such means: that if the transition right-to-left across the instances of an abstraction principle is indeed truth-preserving, it will not be so purely in virtue of the fact that the two halves have, stipulatively, the same truth-condition but will be courtesy of an, as it were, collateral fact that each element in the abstractive domain does indeed *have* an associated abstract of the appropriate kind—that each abstracted term refers—about which independent reassurance is therefore needed; in short, exactly the conception of the matter that we argue above is implicated in the maximalist response to the issues.

If, however, the anxious metaphysician wishes no truck with maximalism, it becomes extremely doubtful whether there is available to him any lucid conception of what such “independent reassurance” might consist in. For all it could consist in, it seems, is the identification of another kind of ground for accepting some range of statements—perhaps identity statements, perhaps others—involving the relevant abstracted terms in reference-demanding ways, but presupposing no other means of reference to abstracts of the relevant kind. In order not to beg the question, such a ground must allow of characterisation in a manner free of occurrences of the relevant class of terms. And now there are just two cases. (i) The ground may follow the broad example of abstraction itself—that of proposing stipulative conceptual equivalences between statements configuring the relevant abstract terms in reference-demanding ways and others. In that case it will raise exactly the same issues as abstraction, and ought to provoke the same anxieties, if they are justified at all. But (ii) if the ground involves, rather, the presentation of what is claimed as *defeasible* evidence for the truth of statements of the relevant kind, the claim will be false in any case where the abstraction in question is *conservative*³⁴—where it has, roughly, no differential consequences for which statements free of reference to or quantification over the relevant abstracts are true or not. For in that case, all the (defeasible) evidence will be exactly as it would be if the abstraction were untrue. To press the demand for independent reassurance in the case of any

³⁴ The optimum characterisation of the relevant notion of conservativeness has proved controversial. (See Weir (2003)) Here is one formulation previously offered in Wright (1999). Let $(\forall \alpha_i)(\forall \alpha_j) (\Sigma(\alpha_i) = \Sigma(\alpha_j) \leftrightarrow \alpha_i \approx \alpha_j)$, be any abstraction. Introduce a predicate, Sx , to be true of exactly the referents of the Σ -terms and no other objects. Define the Σ -restriction of a sentence, T , to be the result of restricting the range of each objectual quantifier in T to non- S items, — thus each sub-formula of T of the form, $(\forall x)Ax$, is replaced by one of the form, $(\forall x)(-Sx \rightarrow Ax)$, and each sub-formula of the form, $(\exists x)Ax$, is replaced by one of the form, $(\exists x)(-Sx \& Ax)$. The Σ -restriction of a theory, θ , is correspondingly the theory containing just the Σ -restrictions of the theses of θ . Let θ be any theory with which Σ -abstraction is consistent. Then Σ -abstraction is conservative with respect to θ just in case, for any T expressible in the language of θ , the theory consisting of the union of (Σ) with the Σ -restriction of θ entails the Σ -restriction of T only if θ entails T . The requirement on acceptable abstractions is, then, that they be conservative with respect to any theory with which they are consistent.

As noted earlier, we regard conservativeness is a prime desideratum if abstractions are to rank as good implicit definitions.

conservative abstraction is thus, we contend, to pose a challenge for which no clear model — maximalism apart—can be given of how it might be answered.

It seems fair, accordingly, to characterise such a challenge as *Sceptical*. So to characterise it is not to answer it, of course, or give a reason for not taking it seriously. But it may mitigate a tendency to sympathise with it. We can extend the parallel a little further. Imagine a situation in which we have only one means of reference to material objects—demonstratives, say, perhaps qualified by a sortal predicate: “that man”, “this tree”, and so on (*material demonstratives*). And suppose we are challenged to produce a reason to think that *any* uses of such expressions succeed in referring. Again, any such reason would have to be a reason to think that certain statements—“that man is running”, “that tree is tall”—embedding material demonstratives in reference-demanding ways are, in their context of use, true. And that in turn will demand a conception of what justifies taking such a statement to be true. Such a conception, so says the Sceptic, will be that of the occurrence of a certain pattern of experience—a pattern which might be fully described in terms of appearances, without commitment to entities of the kind in question. Since the evidence may be so described, independent assurance is wanted that successful referential use of the relevant expressions is possible in the actual world—a fortiori that there are middle sized physical objects out there to be referred to at all—before we may justifiably take such evidence to establish the truth of the appropriate type of statements.

Responses to this kind of scepticism about material objects are of course various. They include denying the ‘neutralist’ (Lockean) conception of experience it exploits, and allowing that conception but denying that any need is thereby entailed for *independent* corroboration of a material world ontology before experience can carry the evidential significance customarily accorded it. Abstractionism, in so far as it reads an ontology of abstracta into the commitments of the right-hand-sides of abstractions, stands comparison with the former (direct realist!) line. But the question we would press on the anxious metaphysician is this: if one is not content to acquiesce in a sceptical view of the referential aspirations of material demonstratives, how is it relevantly different with the terms introduced by abstraction?

§10 Although we take some satisfaction in the dialectical situation as it has just emerged, it is actually very much *not* where we want—or promised—to end up. If the best that can be done with an obdurate doubt about the truth-preservingness of the transitions right to left across the instances of an abstraction is to make good an analogy with the relation between experience and material world claims as viewed by a Sceptic, then we have precisely not made good on what we characterised as of the essence of abstraction: the contention of the conceptual sufficiency of the truth of the right-hand sides for the truth of the left. The whole point was to be that there *is* no metaphysical hostage in the transition, no need for an ‘assist’ from the World, and therefore no scope for doubt, even Sceptical doubt, that the requisite assistance is to hand. The best response to (E’), therefore—at the least, the response to which we are committed—cannot rest upon a comparison between doubt about the inference, right to left, across an instance of an abstraction principle and scepticism about the reality of ordinary material objects. Rather, it has to be to make out a perspective from which abstraction actually involves nothing akin to the element of epistemological *risk* which scepticism finds in our purported cognitive commerce with the external world.

Let’s step back. To ask, with Boolos, how we know that there is any function—hence, any objects to constitute its range of values—that behave as an abstraction principle demands is, in effect, to view the principle as proposed in a spirit of *reference fixing*: as imposing a

condition, viz. association with the elements in the field of the abstractive relation in a fashion isomorphic to the partition into equivalence classes which it effects, which it is then up to the world to produce a range of objects to satisfy. This is the conception of the matter articulated in the following passage:

What did Locke realise about 'gold'? Effectively, that there is an element of blind pointing in our use of such a term, so that our aim outstrips our vision. Our conception fixes what (if anything) we are pointing at but cannot settle its nature: that is a matter of what's out there. One image of the way [Hume's Principle] is to secure a reference for its terms shares a great deal with this picture.³⁵

On this conception, we 'point blindly', using the sortal concept and terms explained by an abstraction principle, in the hope of *hitting off* reference to a range of entities qualified to play the role that the principle defines, and it is accordingly readily intelligible how the process might fail—it goes with the model that it must be at least initially intelligible that a principle proposed in this spirit fails to hit off reference to *anything*. It cannot just be a given that reference is secured, even if it is—let alone that it is secured to entities of which the principle states a necessary truth. Rather, this is something which needs to be verified as a by-product of our, so to say, *finding* a range of objects 'out there' to which the conception embodied in the principle is (necessarily) faithful. And of course if that is to be possible, the objects in question must first be given to us under some *other* mode of presentation.

It is pointless to deny that it is possible to regard abstraction principles in this fashion. One can always ask, with respect to any particular domain of objects, whether there are any that are so related to the elements of the abstractive domain that identity and distinctness among them is tracked by the obtaining, or non-obtaining, of the relevant equivalence relation on pairs from that domain. It may be that in a particular case, the answer is not only affirmative but necessarily so—and in that case, the abstraction principle too will state a necessary truth, even when understood in the reference-fixing spirit. But this spirit—necessary for the 'anxious metaphysical' stance—is simply in flat tension with the abstractionist conception of the matter; indeed, it is to view abstraction principles in a manner inconsistent with their capacity to serve the process of abstraction itself. Properly viewed, the very stipulative equivalence of the two sides of an instance of an abstraction principle is enough to ensure both that it is not to be seen as proposed as part of a project of reference-fixing and that there *is no significant risk of reference failure*.

How can there be no such risk? In order to understand this, we need to be mindful again of the distinction between sparse and abundant properties and the role it can play in the semantics of predicates. For in general terms, the abstractionist metaphysics of abstract objects, and of reference to them—sometimes called *minimalism*³⁶—stands to the conception of the matter that underwrites the reference-fixing model as an abundant conception of properties stands to a sparse one. The analogy admittedly needs some care. On the most generous version of 'abundance' theory, there is for predicates, as remarked, no gap between sense and reference: the association of a predicate with a sense—a determinate satisfaction-condition, even if a necessarily unsatisfiable one—is enough to ensure the existence of a property—a way of being—to play the role of the reference of the predicate. It is not, by contrast, part of the minimalist view of the reference of singular terms introduced by

³⁵ Sullivan and Potter [1997], pp.145-46, quoted in Potter & Sullivan (2005)

³⁶ No cousin, of course, of maximalism in the sense discussed in this paper!

abstraction to conceive of reference as bestowed purely by sense. But nor, according to the minimalist view, is reference secured by the abstraction's merely serving to introduce a conception of a kind of object whose exemplification requires a form of worldly co-operation going beyond anything that can be assured by the laying down of an abstraction principle which is in good standing by normal criteria—and so in particular features a *bona fide* equivalence relation. Anyone should agree that a justification for regarding a singular term as having objectual reference is provided just as soon as one has justification for regarding as true certain atomic statements in which it functions as a singular term. According to the abundant—“neo-Fregean”—metaphysics of objects and singular reference, such a justification is provided by the very manner in which sense is bestowed upon abstract singular terms, which immediately ties the truth conditions of self-identities featuring such terms to the reflexivity of the relevant relation. As with the abundant conception of properties, there is no additional gap to cross which requires “hitting off” something on the other side by virtue of its fit with relevant specified conditions, as the property of being composed of the element with atomic number 79 is hit off (or so let's suppose) by the combination of conditions that control our unsophisticated use of ‘gold’. But nor is it the case that reference is bestowed by the possession of sense alone. The latter view, for singular terms, is Meinongianism. The abstractionist view agrees with the reference-fixing conception that it takes, over and above the possession of sense, the truth of relevant contexts to ensure reference. But it diverges from the reference-fixing conception in what it holds has to be accomplished before those contexts may justifiably be taken as true, and in how straightforward it views the accomplishment as being.

Can we make this clearer? On the abundant view of properties, predicate sense suffices for reference. But it is not the abstractionist view of singular terms that sense suffices for reference—the view is that the truth of atomic contexts suffices for reference. However everyone agrees with that. The controversial point is what it takes to be in position reasonably to take such contexts to be true. The point of analogy with the abundant view is that this is not, by minimalism, conceived as a matter of hitting off, Locke-style, some ‘further’ range of objects. We can perfect the analogy if we consider not simple abundance but the view that results from a marriage of abundance with Aristotelianism. Now the possession of sense by a predicate no longer suffices, more or less, for reference. There is the additional requirement that the predicate be true of something, and hence that some atomic statement in which it occurs predicatively is true. That is a precise analogue of the requirement on singular terms that some atomic statement in which they occur referentially be true. And abstractionist minimalism with respect to objects and singular reference is the exact counterpart of Aristotelian abundance with respect to properties and predicate reference. The Lockean conception, by contrast, is to be compared to the position of the ‘sparse’ opponent of the abundant Aristotelian who construes the relevant range of predicates as purporting reference to sparse properties. On that view there is scope for a doubt whether a relevant predication is true, even when the subject meets the working satisfaction-conditions assigned to the predicate — for there may be no genuine property associated with meeting those conditions. Likewise on the Lockean view, there is scope for a doubt whether an abstract-identity is true even though the appropriate equivalence relation holds between the relevant elements in its field—for there may be no, as it were, ‘sparse’—metaphysical Worldly—objects suitable to serve as the referents of the relevant abstract terms. The abstractionist conception of the truth of the right-hand sides of instances of good abstractions as conceptually sufficient for the truth of the left-hand sides precisely takes the terms in

question out of the market for ‘hitting off’ reference to things whose metaphysical nature is broadly comparable to that of sparse properties, and assigns to them instead a referential role relevantly comparable to that of predicates as viewed by the abundant Aristotelian.

Let us begin to draw things together. Aside from the earlier, rather obvious remarks about the requirement of the truth of the corresponding Ramsey sentences, we have been rather neglecting question (M):

What does the world have to be like in order for (the best examples of) abstraction to work?

What, in the light of the foregoing discussion, should now be said in answer? First, for each equivalence relation which is to underpin an abstraction—for all we have said, indeed, for *every* equivalence relation—there has to be an associated function taking each of the elements which are equivalent under the relation to a common object and no two inequivalent elements to the same such object. Second, the existence of such a function will of course require the existence of a properly behaved range of values. The anxious metaphysician and the abstractionist can agree thus far. Their disagreement concerns what it takes for that to be so. The anxious metaphysician thinks of the issue on the analogy of the existence of a sparse property: just as a predicate's being semantically well-behaved and even featuring in true atomic predications is no assurance that it refers to one of the real properties characteristic of the divisions in the metaphysical World, so the fact that the terms introduced by an abstraction behave as singular terms should and feature in what, if the abstraction is accepted, are well understood and often verified contexts, is no assurance that they refer to any of the real objects in the metaphysical World. One who subscribes to this way of thinking then has to take a decision about whether they refer at all, with the minimalist conception of objects and singular reference on offer to play a role in a positive answer counterpart to that of abundant Aristotelian conceptions of property and predication. If the offer is spurned, the metaphysician will have to deny that abstractions can ever be simply stipulatively true. For the abstractionist, by contrast, there is no well-conceived objection to the unqualified stipulation of (the best) abstractions— if it seems otherwise, it is only because one is trying to combine their stipulative character with a reference-fixing conception of them—and the abundance of the entities thus recognised is simply the objectual counterpart of the abundance of abundant properties.

These remarks are not a *defence* of minimalism but merely a reminder—since it seems that one may be needed—of the kind of background thinking about objects and ontological commitment which undergirds the abstractionist view. Perhaps this background thinking constitutes a ‘metaontology’. If so, then there is much more to say about the spirit of this metaontology —especially about the sense, if any, in which it is happily described as ‘platonist’. But if it is accepted, the answer to question (M) could not be simpler: a world in which abstraction works is a world in which there are equivalence relations with non-empty fields.

University of Sheffield;

Arché Research Centre University of St. Andrews, and New York University

Bibliography

- David Armstrong (1979) *A Theory of Universals*, Cambridge: Cambridge University Press
- George Bealer (1982) *Quality and Concept*, Oxford: Oxford University Press
- Paul Benacerraf (1973) 'Mathematical Truth' *Journal of Philosophy* 70, pp.661-80
- George Boolos (1990) 'The Standard of Equality of Numbers' in George Boolos ed. *Meaning, and Method: Essays in Honor of Hilary Putnam* Cambridge: Cambridge University Press 1990, pp.261-77; reprinted in Boolos (1998), pp.202-19
- (1997) 'Is Hume's Principle Analytic?' in Richard G. Heck, Jr., ed. *Logic, Language, and Thought* Oxford: Oxford University Press 1997; pages references are to the reprint in Boolos (1998), pp.301-14.
- (1998) *Logic, Logic, and Logic* Cambridge, Mass.: Harvard University Press, 1998
- Ross Cameron (2007) 'Truthmakers and Ontological commitment: Or how to deal with complex objects and mathematical ontology without getting into trouble' <http://www.personal.leeds.ac.uk/%7Ephlrpc/research.htm>
- Matti Eklund (2006) 'Neo-Fregean Ontology' *Philosophical Perspectives* 20, pp.95-121
- Bob Hale (1987) *Abstract Objects* Oxford: Basil Blackwell
- (1994) 'Is platonism epistemologically bankrupt?' *Philosophical Review*, Vol 103, No 2, pp.299-325
- (2000) 'Reals by Abstraction' *Philosophia Mathematica* (3) Vol.8, pp.100-23, reprinted in Hale & Wright (2001)
- (forthcoming) 'Neo-Fregeanism and Quantifier-Variance: Chairman's Remarks' (forthcoming in *Proceedings of the Aristotelian Society*)
- Bob Hale & Crispin Wright (1994) 'A reductio ad surdum? Field on the contingency of mathematical objects' *Mind* Vol 103, No 410, pp.169-184
- (2001) *The Reason's Proper Study: Essays towards a Neo-Fregean Philosophy of Mathematics* Oxford: Oxford University Press
- (2002) 'Benacerraf's Dilemma Revisited' *European Journal of Philosophy* Vol.10, No.1, pp.101-129
- (forthcoming) "Focus Restored", reply to MacFarlane (forthcoming) in Øystein Linnebo, ed. Special Issue of *Synthese* on the Bad Company Problem
- Katherine Hawley (2007) 'Neo-Fregeanism and Quantifier Variance' *Aristotelian Society Supplementary Volume LXXXI*, pp.233-49
- Eli Hirsch (2002) 'Quantifier Variance and Realism' *Philosophical Issues* 12, *Realism and Relativism*, pp.51-73
- Harold Hodes (1984) 'Logicism and the ontological Commitments of Arithmetic', *Journal of Philosophy* Vol.81, pp.123-49
- David Lewis (1986) *On the Plurality of Worlds*, Oxford: Basil Blackwell
- John MacFarlane (forthcoming) 'Double Vision: Two Questions about the Neo-Fregean Programme' in

Øystein Linnebo, ed. Special Issue of *Synthese* on the Bad Company Problem

- D. H. Mellor & Alex Oliver (1997) 'Introduction', in *Properties*, Oxford: Oxford University Press, pp. 1–33
- Michael Potter & Peter Sullivan (2005) 'What Is Wrong With Abstraction?', *Philosophia Mathematica* 13, pp.187-93
- Chris Swoyer (1996) 'Theories of Properties: From Plenitude to Paucity', *Philosophical Perspectives* 10, pp. 243–264
- Jonathan Schaffer (2004) 'Two Conceptions of Sparse Properties', *Pacific Philosophical Quarterly* 85, pp. 92–102
- Stewart Shapiro (2000) 'Frege meets Dedekind: A neo-logicist treatment of real analysis' *Notre Dame Journal of Formal Logic* Vol.41, pp.335-64
- Theodore Sider (2007) 'Neo-Fregeanism and Quantifier Variance' *Aristotelian Society Supplementary Volume* LXXXI, pp.201-32
- Peter Sullivan and Michael Potter (1997) 'Hale on Caesar', *Philosophia Mathematica* 5, pp. 135-152
- Alan Weir (2003) 'Neo-Fregeanism: An Embarrassment of Riches' *Notre Dame Journal of Formal Logic*, 44 pp. 13-48
- Crispin Wright (1983) *Frege's Conception of Numbers as Objects*, Aberdeen: Aberdeen University Press
- (1990) 'Field and Fregean Platonism' in A.D.Irvine, ed. *Physicalism in Mathematics* Dordrecht/Boston/London: Kluwer, reprinted as Essay 6 in Hale & Wright (2001)
- (1999) 'Is Hume's Principle Analytic?' *Notre Dame Journal of Formal Logic*, 40(1), pp.6-30; reprinted in Hale & Wright (2001)