

3

The Existence (and Non-existence) of Abstract Objects

Richard G. Heck, Jr.

3.1 TWO PROBLEMS

As George Boolos (1998, 128–129) once remarked, much of our ordinary discourse seems to involve reference to abstract objects. It’s not just numbers and sets, though we do talk about them. We talk also of sentences: How many and what words they contain; how those words are spelled and pronounced; whether they were uttered on certain occasions. We talk of books, like *Die Grundlagen der Arithmetik*: We read them; talk about what sentences they contain; and argue about what is and is not said in them. One might almost be tempted to say that abstract objects are all around us, but for the fact that they aren’t, since they aren’t located in space.¹ It is this that gives rise to the ontological and epistemological problems that abstract objects pose. If abstract objects are not even spatial, they presumably cannot cause anything to happen. And for that reason, among others, we can have no perceptual contact with them. The notion of perception, however, seems to play a fundamental, grounding role in philosophical theories of reference to concrete, spatio-temporal objects, and in theories of knowledge about them. How, then, unless we follow Gödel (1990, 268) and suppose that we have “something like a perception” of sets and other abstract objects, can we explain our capacity to make reference to them or to have knowledge of them?² One does not have to commit a fallacy of hasty generalization,³ supposing oneself to have

¹At least, that’s what one normally supposes. But, as David Lewis once warned, one should be careful not to slide from “we don’t know where they are” to “we know they aren’t anywhere”.

²Charles Parsons’s (1980) notion of ‘quasi-concrete’ objects might of course be helpful here, and it is, in many ways, I think, continuous with the Fregean ideas to be discussed below.

³It is a standard reply to arguments for nominalism that proceed from premises about causal conditions on knowledge that they commit this fallacy. These arguments of course derive, in the contemporary literature, from Paul Benacerraf’s ‘Mathematical Truth’ (Benacerraf, 1973). Versions of the reply can be found in Wright (1983, §xii), in Lewis (1986, §2.4), and in Burgess and Rosen (1997, 35–37). For a recent defense of the claim that such concerns should still be taken seriously, see Øystein Linnebo’s ‘Epistemological Challenges to Mathematical Platonism’ (Linnebo, 2006).

proven that reference to abstracta is impossible, and that we can have no such knowledge, to find these questions pressing.

I am attracted to the following sort of view, which has its origins in Frege (1980, §§62ff). These problems will seem insoluble so long as we insist upon trying to explain directly, so to speak, how we come into cognitive contact with abstract objects. What we ought to do instead is to focus our attention on complete judgments we make about such objects. Rather than attempt to say directly how we are able to make reference to the number six, that is, we should begin by considering the meanings of complete sentences, such as “Six is not prime”. Frege’s thought, embodied in his famous context principle, is that, if we can explain the meanings of all sentences in which the word “six” occurs, there will be no further question what the meaning of that word is: If we can explain how we are able to understand sentences containing the word “six” as having the truth-conditions they do, there will be no further question how we are able to make reference to the number six.

Of course, simply saying that there will be no further question does not make it so. The problem facing Frege’s interpreters, and those who would follow him here, is that of saying *how* explaining the meanings of sentences containing the word “six” answers the question how we are able to refer to six. This is, in part, because of the gap between meaning and reference. As Dummett (1981, ch. 14) was the first (but hardly the last) to point out, the most natural reading of the context principle—which Frege states in terms of an undifferentiated notion of ‘content’—makes it one about *sense*: The sense of a term is exhausted by the contribution it makes to determining the senses of complete sentences in which it occurs. Understood as a thesis about reference, the context principle is much less obvious. That, however, is the form in which we need it. And it is just far from obvious why the fact that “six” contributes in some regular way to the senses of sentences in which it occurs should even seem to imply that it has a reference of its own. Russell (1905) insisted, famously, that descriptions⁴ have no meaning of their own, even though they do contribute regularly to the meanings of sentences in which they occur. Why should “six” be any different? Dummett pushes this sort of question quite hard.⁵

⁴And other quantificational phrases. This point seems to me often to be missed. But if one wants to understand what Russell means when he says that descriptions have no meaning of their own, then one ought first to focus on what he means when he says that “every man” has no meaning of its own, but makes a regular contribution to every sentence in which it occurs. Russell’s point is the now familiar but then revolutionary one that quantificational phrases are not names, and they do not refer to strange entities called *variables*. Russell’s enthusiasm for the point was no doubt connected with the fact that it cleaned up the mess that was his view of quantification in the *Principles* (Russell, 1903), on which (what we now call) quantificational phrases referred to a bewildering variety of variable entities.

⁵In some ways, the present chapter is my attempt to get at what was bothering Dummett, which is, I think, independent of the correctives issued in Wright’s (1983, §x) reply.

Frege’s proposal is most easily developed in the sort of case with which he introduces it, the case of what we may call *types*: Abstract objects that are essentially *of* certain other objects, the types’ *tokens*. Reference to such objects is to be thought of as made, most fundamentally, by means of expressions of the form “the direction of the line ℓ ” or “the word-type of the inscription w ”. The meanings of sentences containing such expressions are then to be explained in terms of a so-called *abstraction principle*. Consider, for example, terms referring to such objects as *Die Grundlagen der Arithmetik*. I shall call these objects *editions*, since the English word “book” can refer either to tokens or to types. Thus, *Die Grundlagen* is an edition, and the physical copy of *Die Grundlagen* presently sitting on my desk is what I shall call a *book*.

Frege suggests that the first task is to explain the meanings of identity statements involving such expressions.⁶ To do this, we must find an equivalence relation $R\xi\eta$ that holds between two books just in case they are copies of the same edition, whence the sentence “the edition of book a is the same as the edition of book b ” may be explained as being true just in case book a bears this relation to book b . That is:

$$\text{edn}(a) = \text{edn}(b) \equiv Rab$$

Similarly, the meanings of sentences such as “The edition of book a is short” are to be explained by finding a predicate “ $S\xi$ ” that is true of a book just in case its edition is short. Thus:

$$\text{short}(\text{edn}(a)) \equiv Sa$$

Similar clauses are needed for two- and more-place predicates.

A number of problems arise immediately. One might think, for example, that there is no guarantee that we will actually be able to find such an equivalence relation and such predicates. This is a problem, however, only if we require that they be definable in terms that do not involve the notion of an edition. But they need not be. That *there is* an equivalence relation that holds between any two books just in case they are copies of the same edition is obvious, namely: the edition of ξ is the same as the edition of η . Whether this can be defined in terms that do not involve editions is irrelevant.⁷ What matters is whether one can come to *understand* an expression denoting this relation without antecedently being able to refer to editions. And one might do so by

⁶Frege’s own attitude towards the proposal is actually quite unclear, since he goes on to raise an objection to it that he seems to regard as conclusive, the Caesar objection (Frege, 1980, §66). Then again, Frege seems to emphasize the import of the basic idea in his summary of the book’s results (Frege, 1980, §107).

⁷The older literature on nominalism takes this question to be central. That it isn’t was made clear to me by Wright (1983, §v). In the present context, the issue is not so critical, since both the neo-Fregean and the nominalist need such an equivalence relation: the neo-Fregean, to explain names of editions; the nominalist, to eliminate them. There is further discussion of this issue in ‘Syntactic Reductionism’ (Heck, 2000, 135ff).

the direct method, by ‘immersion’, rather than by having it defined in simpler terms. That is not to say that there are no interesting problems that arise in connection with the explanation of the relevant equivalence relation, of the ‘criterion of identity’ for the objects in question: It is simply to say that the apparent difficulty, or even impossibility, of defining the relation in terms that do not involve editions does not itself count as an objection to the Fregean approach.

Still, there are problems of principle facing the neo-Fregean project. It is not, for one thing, obvious how it differs from certain sorts of reductionist programs. Neo-Fregeans take the availability of such abstraction principles to show us how reference to abstract objects is possible. But reductionists might equally take the availability of such principles to show that terms purporting to refer to such objects can be eliminated, since we seem to have been told how we could ‘say the same things’ we ordinarily say using sentences containing names of editions without using any such names. One might wonder, however, why the possibility of eliminating names of editions should even seem relevant. The most it can imply is that we ‘need not’ make reference to editions for some purpose or other. But the challenge, at least initially, was to explain how we *can* make reference to abstract objects, not why we must.

But, as I have argued elsewhere (Heck, 2000, §4), abstraction principles do not, in fact, enable us to eliminate expressions that purport to refer to editions. We have not yet considered how quantification over editions is to be handled. And, while first-order quantifiers do not pose a problem, there are non-first-order quantifiers—such as “most”, “few”, and the like—that do. Consider, for example, the sentence “Most editions are short”. Obviously, the following will not do as an analysis:⁸

Most editions are short iff most books are short.

For it might well be the case that most books *are* short, there being many more copies of short editions than of long ones, although most editions are *not* short. To get the truth-conditions of such sentences correct without making explicit reference to editions, we need to consider a collection of *representative copies* of editions, one and only one for each edition:⁹ If most of these representatives are short, then most editions are short, and conversely. More generally, what is required is a functional expression $\varphi(\xi)$ which, given a book as argument, returns as value the representative copy that bears $R\xi\eta$ to it.¹⁰ We then have the following:¹¹

⁸Here, I am using “short” ambiguously, to refer to a property of editions and also to the corresponding property of books. Nothing, of course, turns on this.

⁹There are other ways to proceed, but the differences turn out to be superficial (Heck, 2000, 141ff).

¹⁰That is, φ must satisfy: $R(x, \varphi(x))$, which implies: $\varphi(x) = \varphi(y) \equiv Rxy$.

¹¹Here, “most” is a binary quantifier: “(Most x)($Fx; Gx$)” means: Most F s are G .

Most editions are short iff $(\text{Most } x)[\exists y(x = \varphi(y)); \text{short}(x)]$

Similar analyses will work for the other cases.

The attempt to eliminate apparent reference to editions therefore fails. The formal properties of the functional expression $\varphi(\xi)$ are all but indistinguishable from those of $\text{edn}(\xi)$: Hence, terms of the form $\varphi(a)$ are all but indistinguishable from those of the form $\text{edn}(a)$. In fact, the only relevant difference between $\varphi(\xi)$ and $\text{edn}(\xi)$ is that the range of the former consists of representative copies; that of the latter, apparently, of editions. Similarly, the only relevant difference between $\varphi(a)$ and $\text{edn}(a)$ is that the former refers to a representative copy; the latter, apparently, to an edition. The disagreement therefore concerns not whether names of editions are referring expressions—nor, irrelevantly, how those names are to be spelled—but *to what* they refer.

The fact that it is impossible to eliminate expressions that purport to refer to editions does not, therefore, imply that it is not possible to eliminate *reference* to editions: One could hold that names of editions denote, not abstract objects, but representative books. Indeed, that may not even be the best way to put it. Surely editions are just whatever names of editions denote. If so, then, on this view, editions just *are* books. We are thus left, quite generally, with the question whether editions are books; whether word-types are their tokens; whether either is instead an equivalence class; or what have you. Following Dummett, I call this problem the *problem of trans-sortal identification*.¹²

There is another, related problem. Equivalence relations are not hard to come by. Let $Q\xi\eta$ be an equivalence relation, chosen completely at random: It might, for example, have as one of its equivalence classes the set containing each of my shoes, my daughter Isobel, Brown University, and some other things. We can now introduce names purporting to stand for objects of a certain sort, call them *duds*, just as we introduced names of editions:

$$\text{dud}(a) = \text{dud}(b) \equiv Qab$$

But are we really to believe that there are such objects as duds? I, at least, have a reasonably strong intuition that there are no such things. This problem I call the *proliferation problem*.¹³ One might reply that the existence of duds follows trivially from the explanation given of names of duds: According to that explanation, $\text{dud}(a) = \text{dud}(b) \equiv Qab$; but $Q\xi\eta$ is an equivalence relation, so certainly Qaa ;¹⁴ so, trivially, $\text{dud}(a) = \text{dud}(a)$ and so by existential generalization, $\exists x(x = \text{dud}(a))$. Moreover, the argument recently rehearsed, that

¹²Note that the problem of trans-sortal identification is *not*, as I understand it, the Caesar problem. It is one aspect of the Caesar problem, but the Caesar problem has more heads than the hydra.

¹³This problem is mentioned at the end of ‘Syntactic Reductionism’ (Heck, 2000), but it was not developed there. It is mentioned by Hale and Wright in the postscript to *The Reason’s Proper Study* (Hale and Wright, 2001a, 423–424), where it is treated somewhat dismissively.

¹⁴Of course, $Q\xi\eta$ need not be a full equivalence relation. It need only be symmetric and transitive, and reflexive on its domain. But we can account for that easily enough.

expressions introduced by abstraction must be construed as referring expressions, applies here. So it is hard to see how we can avoid the claim that duds exist. And therefore, the reply might conclude, whatever intuition we might have that there are no such objects as duds ought just to be abandoned in the light of theory.

My worry, however, is that this seems to make the notion of an object, and the conception of what it is for an object to exist, so thin that the resulting view does not obviously deserve to be regarded as any form of realism. Editions, directions, word- and letter-types are all *said* to ‘exist’ all right, but they ‘exist’ in no more robust a sense than duds do, not to mention all the other objects whose names could be introduced in terms of random equivalence relations. In so far as we have an intuition that these latter objects simply *don’t* exist, any view that affirms that they do will seem to *deny* that word- and letter-types exist, even while it affirms our right to *say* that they do.¹⁵

The worry, then, is that it is becoming difficult to distinguish the neo-Fregean view from what one might call a ‘permissive’ nominalism. Permissive nominalists are perfectly happy to let us ‘speak of abstract objects’, so long as we do not take their existence seriously.¹⁶ And what makes this all the more troublesome is that permissive nominalism is not just the most resilient form of nominalism but one of the oldest. It is, for example, what one finds in Berkeley and Hume. Their view was not so much that we should not talk in ways that seem to make reference to abstracta, but that what passes for thinking about the abstract is really just abstract thinking about the concrete. Thus, Berkeley writes:

... [S]uppose a geometrician is demonstrating the method of cutting a line in two equal parts. He draws, for instance, a black line of an inch in length: this, which in itself is a particular line, is nevertheless with regard to its signification general, since, as it is there used, it represents all particular lines whatsoever; so that what is demonstrated of it is demonstrated of all lines, or, in other words, of a line in general. And, as that *particular* line becomes general by being made a sign, so the *name* “line”, which taken absolutely is particular, by being a sign is made general. And as the former owes its generality not to its being the sign of an abstract or general line, but of all particular right lines that may possibly exist, so the latter must be thought to derive its generality from the same cause, namely, the various particular lines which it indifferently denotes.

(Berkeley, 1930, 14–15, emphasis in original)

Berkeley would thus have been perfectly happy to let us talk of editions and even to say that there are such things as editions. The truth of such a claim, for him, is a more or less immediate consequence of facts about how ‘talk about editions’ is to be understood. It is, in particular, a more or less immediate consequence of the fact that there are books and a suitable equivalence relation on them.

Neo-Fregeans have tended to reply to this kind of remark by saying: Exactly! Their view is that there is no ‘metaphysical distance’ between, on the

¹⁵Compare Lewis (1999) on Meinongian views of existence.

¹⁶It is this sort of view that I take to be the one defended by Dummett (1981, ch. 14).

one hand, the fact that there are books and a certain equivalence relation on them and, on the other hand, the fact that there are editions; the former is supposed to be entirely sufficient for the latter. And if one wants to say that we now seem to be at stalemate, then neo-Fregeans have generally regarded that as good enough, insisting that the burden of proof is on their opponents. But is it really? Too much seems to depend on which case one takes to be exemplary. The neo-Fregean wants to say that we should take reference to duds as seriously as reference to editions, intuitions to the contrary be damned; the permissive nominalist wants to say that we should take reference to editions no more seriously than we take reference to duds. But the former perspective is the right one only if the neo-Fregean explanation of names of editions has accomplished what is claimed for it. Only if we assume that the explanation in terms of abstraction principles has succeeded in explaining reference to editions are we under any obligation to assume that the parallel explanation must have succeeded in explaining reference to duds. And that can't just be assumed, no matter who has the burden of proof. Indeed, the objection is precisely that the parallel explanation does *not* explain reference to duds, since there are no such things, whence the original explanation can't have explained reference to editions, either. To suggest that we should simply set the intuition that there are no such things as duds aside is to suggest that we should simply assume the correctness of the neo-Fregean account.

So what might a fan of the neo-Fregean approach do? The classical form of neo-Fregean view, developed by Bob Hale and Crispin Wright, and which I shall henceforth call *Naïve Platonism*, is that *all* terms introduced by abstraction refer, and refer to abstract objects. The view that such terms refer all right, but that none refer to abstract objects, is what I shall henceforth call *Semantic Reductionism*. And the view that such terms do not refer at all is what I shall call *Fictionalism*. What I want to explore is the possibility of a more discriminating view, according to which *some though not all* expressions introduced by abstraction refer to abstract objects.¹⁷

In the next section, I shall begin searching for such a discriminating view by exploring the differences between the cases of duds and editions. By the end of Section 3.2, I hope to have formulated a position naturally described as between Naïve Platonism and Semantic Reductionism. In Section 3.3, however, I shall offer various reasons for dissatisfaction with this position. These criticisms will lead us to a new position, one whose statement requires only a minor reformulation of the original one—almost just a shift of emphasis. This view is better described as between Naïve Platonism and Fictionalism.

If such a view can be found and properly motivated, then whatever serves to motivate it will also serve as the raw materials for an argument against all three of the alternatives I mentioned. Each of them will stand charged of ignoring the differences we will have identified: of treating all abstraction principles

¹⁷Some such view may in fact be that of Hale and Wright. It depends upon how exactly their response to the Caesar objection is ultimately developed.

the same, when they ought to be treated differently. I will not, however, be attempting to complete that argument here. We will have enough to do just to describe the middle ground.

3.2 SEMANTIC REDUCTIONISM AND PROJECTIBLE PREDICATES

According to Semantic Reductionism, names apparently of abstract objects do not refer to abstract objects, but to objects of some concrete sort. To evaluate this position, we need an account of what determines the sort of object to which a proper name refers. There is obviously no way that we are going to answer that question in full generality here. I shall limit attention to the case of terms purporting to denote types, expressions which are, according to the neo-Fregean approach, to be explained *via* abstraction principles.

Many who have discussed this question have attempted to characterize the sort of object to which a term refers in terms of the criterion of identity associated with that term, that is, in terms of the equivalence relation mentioned in the abstraction principle. The simplest such view is that the sort of object to which a name refers is wholly determined by the criterion of identity, whence no names with distinct criteria of identity can refer to objects of the same sort (Hale, 1988, 215). It is important here that the notion of a criterion of identity be intensional, in the sense that substitution of a co-extensive relation for the equivalence relation mentioned in the statement of the abstraction principle need not preserve its status as a correct specification of the criterion of identity for the objects in question. Otherwise, objects introduced by different abstraction principles might be of the same sort in one world—in which the equivalence relations were, purely accidentally, extensionally co-incident—but not in other worlds, and that seems bad.

This simple view would arguably imply that editions are neither books nor sets, but are *sui generis*. Identity statements of the form “ $\text{edn}(a) = \text{edn}(b)$ ” are true if, and only if, as I shall henceforth say, a copies b : The identity of sets, however, is determined by sameness of membership; the identity of books, by something else still. But consider these abstraction principles:¹⁸

$$\begin{aligned} \text{dir}(a) = \text{dir}(b) &\equiv a \parallel b \\ \text{dor}(a) = \text{dor}(b) &\equiv \exists x(a \perp x \wedge b \perp x) \end{aligned}$$

It seems to me that ‘dorections’ might well be our old friends directions. Perhaps they are distinct, but any principle that immediately entails that we cannot identify them is too strong.

The simple view can be weakened: Wright’s generalization of his condition N^d is an example of such a weakening. Suppose that Fx is a sortal concept,

¹⁸For simplicity, I assume here that we are dealing with two-dimensional Euclidean geometry.

and that names of F s have been explained by means of an abstraction principle formulated in terms of some equivalence relation $R\xi\eta$, which itself holds between objects of sort S . Then, Wright’s proposal is:

Gx is a sortal under which [F s] fall if and only if there are, or could be, terms, “ a ” and “ b ”, which recognisably purport to denote instances of Gx , such that the sense of the identity statement, “ $a = b$ ”, can be adequately explained by fixing its truth-conditions to be the same as those of a statement which asserts that the given equivalence relation [$R\xi\eta$] holds between a pair of objects [of sort S]. (Wright, 1983, 114)

To put it slightly differently: F s may be identified with G s if, but only if, identity statements concerning some G s may be explained in the same way that identity statements concerning F s are explained. In particular, the F s will be identifiable with objects of sort S if, and only if, identity statements containing names of some objects of sort S may be explained by means of the abstraction principle in terms of which names of F s are to be explained.¹⁹

This proposal probably resolves the problem of directions and dorections. Identity statements of the form “ $\text{dor}(a) = \text{dor}(b)$ ” plausibly can be explained in terms of the parallelism of the two lines a and b . But, in fact, that isn’t at all clear, in large part because the notion of explanation to which Wright appeals is not very clear. And that makes it hard to know how to apply Wright’s proposal in general. We may presumably take the names which “recognisably purport” to refer to sets to be those of the form “the set of books that copy ξ ”. Is it or is it not possible to ‘explain’ the senses of statements of the form “the set of books that copy $a =$ the set of books that copy b ” in terms of “ ξ copies η ”? Or consider terms of the form “the oldest extant copy of ξ ”. Does

the oldest extant copy of $a =$ the oldest extant copy of b iff a copies b

count as an explanation of identity statements involving these terms? Maybe one has intuitions about the matter, but one would like more than a brute appeal to intuition. For this reason, I shall offer a different sort of solution, leaving open the question to what extent it is compatible with Wright’s.²⁰

Consider the expression “the father of ξ ”. It seems obvious enough that the father of John is a person. And one might suppose that the fact that John and Jane have the same father if, and only if, the same male begat them is what determines that “the father of John” refers to a person. It is, no doubt,

¹⁹There seems no reason not to suppose that the converse must also be true: If some F s are G s, then some G s are F s, so we must, presumably, be able to explain identity statements concerning (some) F s in the same way we explain identity statements concerning (some) G s. Of course, since the condition, as formulated, applies only to sortals introduced by abstraction, we will not be able to apply it unless G is also such a sortal. But some such converse seems reasonable. And if so, then Wright’s proposal implies that it will be possible to explain the truth-conditions of *mixed* identity statements—such as “ $\text{dor}(a) = \text{dir}(b)$ ”—both in terms of the criterion of identity for directions and in terms of the criterion of identity for dorections. That seems *very* reasonable.

²⁰The modifications to this view made by Hale and Wright in ‘To Bury Caesar...’ (2001b, §6) leave untouched its near total reliance upon criteria of identity, though there are points of contact with the present view, for which see footnote 24.

of great importance that the father of a = the father of b if, and only if, the same male begat a and b . But this does not, by itself, entail that “the father of John” refers to a person. For consider the following expressions:²¹

- the set of persons who have the same father as John
- the oldest paternal half-sibling of John
- the singleton of the father of John
- the location of John’s oldest paternal half-sibling

Each of these expressions has the same ‘weak identity-conditions’ as “the father of John”: That is, the reference of any one of these expressions will remain unchanged when we substitute a new name for “John” if and only if the same male begat John and the person whose name is substituted for his. But not all of these expressions refer to objects of the same sort, and those which do refer to different objects.

A similar point applies to names of editions, as the following set of examples shows:

- the edition of book a
- the oldest extant copy of book a
- the set of books that copy book a

Again, substitution of the name of any book that copies a will leave the referent of each of these expressions unchanged; and *only* the substitution of such names will do so. Nonetheless, not all of these expressions refer to objects of the same sort: One refers to an edition; one, to a book; one, to a set.

The point I am illustrating with these examples can be stated precisely. Let $\varphi(\xi)$ be a function from objects of sort S to objects of sort T (not necessarily different from S). Then $\varphi(\xi)$ induces an equivalence relation $\Phi\xi\eta$ on objects of sort S , which we define as follows:

$$\Phi xy \stackrel{df}{=} \varphi(x) = \varphi(y)$$

Distinct functions from S to T induce the same equivalence relation, as do various functions from S to sorts T' distinct from T . There are thus many distinct functions on S that have the same weak identity-conditions.

There can thus be no objection to our introducing a functional-expression which will satisfy the abstraction principle by means of which names of editions are explained and whose range will consist of books, or sets, or objects of many other sorts. In principle, a Semantic Reductionist could hold that names

²¹This kind of point has the status of what mathematicians call ‘folklore’. I first heard a version of it made by Dummett in his 1989 Hilary Term lectures on *Die Grundlagen*. Dummett remarked that if (what I am calling) weak identity-conditions determine the sort of object to which a name refers, then it is philosophically confused to think that the eccentricity of an ellipse is a real number (Dummett, 1991, 162–163). I heard Lewis make points in the same vicinity at MIT a few years later; Warren Goldfarb once mentioned a similar example to me; and it surfaces in a paper by Sullivan and Potter (1997, 139ff). It is discussed by Hale and Wright in ‘To Bury Caesar...’ (2001b, 375ff).

of editions refer to just about anything: People, rocks, trees, books, or sets, so long as there are enough of them to go around. But the most principled such views are that names of editions refer to representative copies and that they refer to equivalence classes. The latter option is of course not one acceptable to a Nominalist (unless she has a Nominalistic treatment of set-theory waiting in the bushes).²² So I shall focus attention on the former proposal, that “ $\text{edn}(\xi)$ ” is a functional-expression whose range consists of actual, physical books. What we need to ask now, then, is what, if any, features of the use of this expression are sensitive to the sort of object in its range: We need to ask, that is, what difference it would make if expressions like “*Ulysses*” actually were treated as referring to books.

Suppose, then, that “*Ulysses*” really does refer to a particular, physical book. That book must have some physical location; one of its pages might be torn; someone probably owns it; perhaps someone is holding it at this very moment. If that book is the reference of “*Ulysses*”, then some such sentences as “*Ulysses* is in Texas”, “*Ulysses* has a torn page”, and so forth, must be true. That seems very odd. Some of the oddity can be avoided if one takes a slightly different view, one that actually seems implicit in the passage from Berkeley quoted earlier: Instead of saying that “*Ulysses*” denotes some particular one of its copies, we should say that it “indifferently” denotes each of them, a view naturally explained in terms of supervaluations.²³ The sentence “*Ulysses* is in Texas” would then be true only if *every* copy of it were in Texas; false, if none were; and neither true nor false, otherwise. Still, though, “*Ulysses* is on planet Earth” will probably come out true. And it is perfectly possible, though unlikely, that “*Ulysses* has a torn page” should also be true.

One does not usually think of such sentences as having even the remotest chance of being true. There is, indeed, a temptation to deny that these sentences so much as make sense, on the ground that they involve some kind of category mistake. But, as Frege in effect remarks, there is nothing to prevent us from saying such things as that *Ulysses* has a torn page, meaning by this that every copy of *Ulysses* has a torn page, if such a way of speaking should seem useful (Frege, 1980, §69).

Still, there is a felt difference between these sorts of claims and the claim that *Ulysses* contains the word “dog”. The intuition that there is such a differ-

²²Another option that might be popular in some circles would be to take a name of a type to be the fusion of its tokens. One might then justify the existence of types by appeal to unrestricted mereological comprehension. This view is of course vulnerable to the objections to be developed below, since it too treats all abstraction principles the same way, but there is a more specific worry about it, too, namely, that it is demonstrably false in some cases. Consider directions (in the plane again, for simplicity). The fusion of all lines parallel to any given line a is the whole plane—assuming, as seems reasonable (at least on the view we are considering), that lines and planes are both fusions of points. But then all lines have the same direction.

²³A view of this sort has been elaborated and defended by Hodes (1990). It is because of the availability of this view that I suspect the sort of modal differences explored by Uzquiano (2005) will not do all the work needed here.

ence is presumably what is behind the intuition that expressions like “*Ulysses*” do not refer to books. This difference has nothing to do with abstract objects as such. There is a similar difference between saying, of Frege, that he had blue eyes and saying, of Frege, that he had only blue-eyed children. The temptation here is to say that, in the former case, though not the latter, one is speaking of a *property* of Frege: It is a property of Frege that he had (or did not have) blue eyes; it is no property of him, in this strict sense, that he had (or did not have) only blue-eyed children. This distinction—between ‘real’ and ‘merely Cambridge’ properties—is infamously difficult to explain clearly. Present purposes, however, require only that it be explained for the case of types, and that is a good deal easier.

There is much that can be said about books: That they are dirty, that they have some mass, that they contain an inscription of some word, and so forth. There are, that is, many predicates that can sensibly be attached to names of books. Of these predicates, some play a special role in our talk of editions: Those that are satisfied by a given book if, and only if, they are satisfied by every book that copies it. More important still are predicates whose satisfaction, by a given book, can always be determined even if one does not know which other books copy it. Predicates like “contains an inscription of the word ‘dog’” are of this kind: One need not know which other books copy a given one to be able to determine whether that book satisfies this predicate; knowledge that a given book does satisfy it suffices for knowledge that every other copy also satisfies it. It is because there are such predicates that our ordinary practices involving editions are possible. Scholarly discussion could not be carried on as it is if one had to stop, every time one wanted to make a claim about what is said in *Die Grundlagen*, to determine what copies of it exist, what is said in them, and so forth. (For much the same reason, one can read an edition without reading any single copy of it, but rather parts of different books.) Borrowing a term, let us call such predicates as “contains an inscription of the word ‘dog’” *copy-Projectible* predicates of books.

Why think that predicates of types that are introduced in terms of Projectible predicates of tokens should express properties of the types?²⁴ The allusion to projectibility,²⁵ as that notion is employed in the philosophy of science, is intentional. Projectible predicates are ones whose satisfaction by a particular sample of a substance, or by a particular member of a species, in some sense implies its satisfaction by all samples and all members. It is natural, for that

²⁴Note that the idea that it matters whether there are Projectible predicates of the tokens subsumes the proposal made by Hale and Wright (2001a, 424 footnote 8) that it matters that the equivalence relation itself should in some sense be projectible: that we should know how to extend it to non-actual things. That is a clearly a weaker requirement than the one considered here, but it is very much in the same spirit. The idea that modal considerations might be relevant also surfaces in ‘To Bury Caesar...’ (Hale and Wright, 2001b, 357ff). Similar ideas have surfaced elsewhere, too.

²⁵I shall capitalize my term “Projectible” to remind the reader that I may or may not be using it in its usual sense.

reason, to think of these predicates as expressing properties of the substance or species itself. There is a clear analogy between such predicates and ones that are Projectible in my sense. Properties of the type, in this sense, do not depend upon what tokens happen to exist (just as properties of a substance should not depend upon what samples happen to exist):²⁶ The creation or destruction of particular tokens will not, on this analysis, affect what properties *Ulysses* has—though it might affect whether it has only copies with torn pages.

There is another, more technical reason to think that Projectible predicates of tokens are especially important here. To over-state the point slightly: Only if a predicate of types is explained in terms of a Projectible predicate of tokens is it possible to make genuinely informative applications of Leibniz’s Law involving that predicate of types. Consider sentences of the form “ $F(\varphi(a))$ ” and “ $F(\varphi(b))$ ”, where $R\xi\eta$ is the equivalence relation figuring in the abstraction principle for “ $\varphi(\xi)$ ”. If the predicate “ $f(\xi)$ ” in terms of which “ $F(\xi)$ ” is explained is not R -Projectible, then determining whether “ $F(\varphi(a))$ ” is true will, in general, require one to know which other objects bear $R\xi\eta$ to a and whether those objects satisfy “ $f(\xi)$ ”. One will, in the course of determining whether “ $F(\varphi(a))$ ” is true, therefore have to go through essentially the same procedure one would have to follow to determine whether “ $F(\varphi(b))$ ” is true: In particular, one will (normally) have to determine whether Rab . If $\varphi(a) = \varphi(b)$, then one will then have all the information necessary to determine whether “ $F(\varphi(b))$ ” is true—though, of course, one need not draw the conclusion explicitly and may not even realize as much. By contrast, if “ $f(\xi)$ ” is R -Projectible, one can determine whether “ $F(\varphi(a))$ ” is true simply by determining whether $f(a)$: In particular, one need not know whether Rab or whether $f(b)$. Hence, discovering that Rab —that is, that $\varphi(a) = \varphi(b)$ —can lead to genuinely new information.

Why should that matter? The possibility of informative applications of Leibniz’s Law seems to me to be bound up with the independence of an object from our ways of conceiving it. That we can know that the direction of a is such-and-such without being in a position to know whether the direction of b is such-and-such—even if, in fact, $\text{dir}(a) = \text{dir}(b)$ —suggests not just the familiar gap between sense and reference but an even more important gap between our ability to refer to (or think about) the object and the object itself. To the extent that informative applications of Leibniz’s Law are possible, to that extent the object will seem independent from our ways of conceiving it, and hence to that extent its existence will seem to be independent of the fact that we can conceive of it at all.

To sum up, then, in so far as we have an intuition that, say, “*Ulysses*” does not denote a particular copy of *Ulysses*, that intuition rests upon the thought

²⁶It would be really nice if we could also say: Properties of the type, in this sense, do not depend upon *whether* any tokens of the type happen to exist. If so, then the view we are discussing might also help with what Hale and Wright (2001a, 422–423) call the “problem of plenitude”. But more work needs to be done here.

that not everything that can be said about books can be said about *Ulysses*. But, as we saw, that can't be all there is to the intuition, since we could easily introduce conventions allowing us to say such things as “*Ulysses* has a torn page”, meaning thereby that all of its copies have a torn page (or that some of its copies do, or that one of its copies does, or whatever). The intuition is thus more refined: It is that, even if we did introduce such new predicates, they would not express *properties* of editions. So the problem became to say what a property of an edition is, and I proposed that a predicate expresses a property of editions if it is explained in terms of copy-Projectible predicates of tokens. And so, to the original question, what determines the sort of object to which expressions introduced by abstraction refer, my proposed answer is: It is determined by what predicates of its tokens are Projectible over the relevant equivalence relation.

In the case of duds, there seem to be almost no Projectible predicates of the tokens: The randomness of the equivalence relation $Q\xi\eta$ in terms of which names of duds were introduced essentially guarantees that there are no such predicates. That means that the introduction of names of duds would be largely without point. Any predicate introduced by means of abstraction must be introduced in terms of a predicate of tokens that is a congruence with respect to $Q\xi\eta$: That is, any such predicate must be true of all tokens of a given dud or else false of all of them. It is easy to produce such predicates. Most of these, however, will be similar to “ ξ has only copies which have a torn page”, for example: “ ξ is such that all objects to which it bears $Q\xi\eta$ weigh at least ten pounds”. Determining whether such a predicate is satisfied by a particular dud will, in general, require knowing to which other objects a given object bears $Q\xi\eta$. That there would be little point in talking about duds is less important, however, than *why* there would be little point: In the sense in which to speak of someone's eye-color is to speak about them, and to speak of the eye-color of their friends is not, there would be almost nothing to be said about duds; duds will have almost no ‘properties’ at all, other than being identical with or different from each other.

Assuming that names of duds refer, then, to what sort of object should they be taken to refer?²⁷ I would suggest that, in this case, we have as yet no reason to deny that they refer to equivalence classes under the relation $Q\xi\eta$: What it is true to say of an equivalence class is wholly determined by what is true of its members; what it is true to say of a given dud will, in general, be determined by what happens to be true of all, some, most, few, etc., of its tokens. Of course, that is not much of an argument for the identification. But we need not pursue the matter, as I shall shortly be rejecting the claim that names of duds refer at all.

²⁷If we were to say, as Hale and Wright suggest (see footnote 24), that the equivalence relation used in an abstraction principle must be projectible, then that would dispense with duds altogether. But we do not yet have any motivation for that proposal, and it does not solve the problem I will shortly introduce, about day-persons.

Nonetheless, we have made progress, since the present proposal does at least illustrate how we might catch what we are chasing: a view that does not treat all abstraction principles alike. I complained earlier that the going views seem to me problematic for precisely that reason, that they treat all types the same way: editions, just like duds; duds, just like editions. That this is true of the various versions of nominalism is clear enough; that it is true of Wright’s view, mentioned earlier, is perhaps less so. But Wright’s generalization of N^d , so far as I can tell, appeals only to broadly formal features of abstraction principles in attempting to specify the reference of expressions introduced by those principles. If so, then it is hard to see how duds and editions can come apart. On the other hand, in characterizing the sort of object denoted by names introduced by abstraction in terms of Projectible predicates of the tokens, the present proposal treats different abstraction principles differently and so offers us at least a hope of solving the proliferation problem.

All is not well, however. Consider an abstraction principle based upon the equivalence relation, “ ξ was born on the same day as η ”, by which I mean: within the same twenty-four hour period, Greenwich mean time. The abstraction principle is thus:

$$\text{day-person}(a) = \text{day-person}(b) \text{ iff } a \text{ was born on the same day as } b$$

What should we say about the reference of expressions introduced in terms of this relation? One might be tempted to say that there are going to be very few Projectible predicates of persons, and so that day-persons too are at best equivalence classes. But this would be incorrect. Any predicate of days, referenced to Greenwich mean time, will obviously be Projectible. The obvious thing to say, then, would be that names of day-persons in fact denote days. This is a satisfying result, in some ways. One might even wonder what reason there could be for holding that anything more is going on here than the introduction of an abbreviation for “the day on which ξ was born”. On the other hand, however, the conclusion that names of day-persons simply denote days is troubling. It isn’t very hard to explain the equivalence relation *born on the same GMT day* without making a direct appeal to anything about days. If not, then it seems something of a surprise that expressions introduced by abstraction on that relation should end up referring to days.

The sort of worry this example illustrates can be generalized. Let $R\xi\eta$ be the equivalence relation for a given abstraction principle with base sort B , and let S be an arbitrary sort of object of which there are at least as many as there are equivalence relations under R . Now define a function $\varphi(\xi)$ from B to S that respects R , in the sense that $\varphi(a) = \varphi(b) \equiv Rab$; assume further that we can do so in such a way that determining what $\varphi(x)$ is, for arbitrary x , does not require knowing to which objects x bears $R\xi\eta$. Now let F be any property of S s. Then $F(\varphi(\xi))$ is R -Projectible, in our sense. But, since this depends only upon how many S s there are (and our ability to define $\varphi(\xi)$)

appropriately), there are going to be lots and lots of R -projectible predicates that express all sorts of different properties, quite independent of what the base sort is and what the equivalence relation is. Surely, however, such oddly defined predicates cannot be guaranteed to express properties of the types.

As we characterized the notion above, an R -Projectible predicate was required to satisfy two conditions: First, that it be a congruence with respect to $R\xi\eta$; second, that it should be possible to determine whether the predicate is satisfied by a given token without knowing which, if any, other tokens bear $R\xi\eta$ to it. What we have omitted from the specification of R -Projectibility is another condition that was really implicit in the spirit of the proposal. The predicate, “On the day on which ξ was born, it was cloudy in London”, was not the sort of thing one had in mind as day-Projectible, the sort of thing that will express a property of the types. That, I take it, is because this predicate does not even express a property of the *tokens*. So we should add this condition to our account of R -Projectibility: An R -Projectible predicate must express a property of the tokens. Indeed, we might just speak not of R -Projectible predicates, but of R -Projectible properties. And what names of day-persons refer to will then depend upon what day-Projectible *properties* of people there are.

The resulting proposal is thus this: The sort of object to which names introduced by abstraction refer is determined by what R -Projectible properties of the tokens there are. Let us call this proposal the *Projectibility View*.²⁸ If the Projectibility View could be sustained, it would apparently imply that names of day-persons, like names of duds, at best denote equivalence classes, as there do not appear to be any properties of persons that are suitably Projectible. And yet, it would allow us to say that names of editions do not denote equivalence classes, since there are a large number of copy-Projectible properties of books; and it would allow us to say that names of editions do not denote books either, since not all properties of books are properties of editions. So it is, again the kind of view for which we are searching.

But, once again, all is not well. Are we really certain that there are no day-Projectible properties of persons? Whether there are is an empirical matter, not one which can be settled *a priori*. On the Projectibility View, then, it is an empirical question—and by no means a question belonging to linguistics—to what sort of object expressions of the form “day-person(t)” refer. And that is a consequence I simply find incredible. It is not that I think there are no interesting empirical questions in this area. Since it plainly is an empirical (or, more generally, a substantial) question whether any properties of persons are day-Projectible, there is certainly an empirical (or substantial) question about day-persons in the vicinity. But, on the Projectibility View, one could have as good an understanding of expressions denoting day-persons as it is possible to have and yet have *no idea* to what sorts of objects they refer, even whether those objects are abstract or concrete. That is the consequence that

²⁸This view was inspired by Bromberger’s discussion in ‘Types and Tokens in Linguistics’ (Bromberger, 1992b), and is intended to be a natural generalization of his view there.

I find simply incredible. So, while there’s an empirical question around here somewhere, the Projectibility View misidentifies it.

3.3 IDEOLOGY, EXISTENCE, AND ABSTRACT OBJECTS

Even if there are no day-Projectible predicates that express properties of persons, things might have been otherwise. Some years ago, there was a fad about what were called ‘bio-rhythms’. Persons born on the same day were supposed to share certain general features of their day-to-day mental and emotional states: degrees of awareness, laxity, happiness, and so forth. Persons were, that is, supposed to have the same *bio-rhythms* if they were born on the same day. Idealizing, let us suppose that there is a detailed theory, Bio-rhythm Theory, making more precise and enlarging upon this idea. It might have been true. Had it been true, there would have been a great many properties of persons that were day-Projectible. On the Projectibility View, therefore, expressions denoting day-persons would then have denoted a certain sort of abstract²⁹ object: They would denote structural features of a person’s mental and affective states, and these features would be shared by persons born on the same day, just as parallel lines share a direction and as books that copy one another share an edition.

According to the Projectibility View, then, as things are (or, at least, as we think they are), terms of the form “day-person(*a*)” denote equivalence classes; if Bio-rhythm Theory had been true, they would have denoted bio-rhythms. But, again, I just find it implausible that the sort of object such expressions denote should turn on the empirical question whether Bio-rhythm Theory is true. Indeed, it is not clear to me how a proponent of the Projectibility View could avoid saying more, namely, that, as things are, day-persons *are* equivalence classes and that, were things as just imagined, day-persons would have *been* bio-rhythms. And that, I take it, would be flatly incoherent: It would amount to supposing that what are in fact equivalence classes should have been bio-rhythms.³⁰ I have no fixed view about whether the Projectibility View can avoid this consequence. But it is suggestive.

²⁹At least, these *could* have been abstract. But of course one can imagine lots of ways in which Bio-rhythm Theory might have been true, and on some of them maybe bio-rhythms would be angels. The point does not actually matter to the argument, however, so far as I can see. There are lots and lots of variations on this theme, so if this one doesn’t work, another one will. And, for what it’s worth, I tend to agree with Lewis (1986, §1.7) that the abstract–concrete distinction is not very clear, anyway. Indeed, I think the attitude towards abstracta expressed here serves to blur the distinction even further: If I am right, then some sorts of abstracta only (fail to) exist contingently, which is a stronger claim than that individual abstracta may contingently (fail to) exist.

³⁰If that doesn’t seem so bad, then reflect upon the fact (to be noted below) that there are other theories, incompatible with Bio-rhythm Theory, whose truth would imply that quite different properties of persons were Projectible. Had one of those theories been true, then day-persons would have been some other sort of object. In any event, the view that all abstracta turn out to be equivalence classes is out of the spirit of Frege’s original proposal: That’s the view to which he retreats in the face of the Caesar problem.

Consider the following, slightly different version of the example. Suppose there were some people who mistakenly *believed* Bio-rhythm Theory to be true (as, indeed, many people more or less did). Our description of how the world would have been if Bio-rhythm Theory had been true is *eo ipso* a description of how those who believe that Bio-rhythm Theory is true believe the world actually to be. Thus, people who believe Bio-rhythm Theory to be true take there to be a variety of day-Projectible properties of persons and so, by the reasoning of the last paragraph, take themselves to refer, by means of expressions of the form “day-person(*a*)”, to bio-rhythms, and so not just to equivalence classes. How then should we describe to what they *do* refer? Should we say that, as a matter of fact, they refer to equivalence classes and not to bio-rhythms?

I do not think our hypothetical speakers would or should accept such a description of their linguistic practice. In speaking of bio-rhythms, they take themselves to be speaking of objects of a particular kind, ones that are shared by people born on the same day. If our hypothetical speakers were to become convinced that nothing like the properties they think are day-Projectible actually are day-Projectible—if they were to become convinced that people do not, in general, share anything like the properties they think they do—then what they would say is not that bio-rhythms have turned out to be equivalence classes, but that the objects to which they thought they were referring, bio-rhythms, *do not exist*. If this were not so, then it would be obscure why they were no longer prepared to accept the truth of a sentence such as “the day-person (that is, bio-rhythm) of *a* is lethargic”, even if everyone born on the same day as *a* was, as it happened, lethargic. It is not sufficient for the truth of that sentence, as they understand it, that everyone who happens to have been born on the same day as *a* just happens to be lethargic. The truth of such a sentence, as they understand it, depends upon the truth of Bio-rhythm Theory (or something like it). And that is because the existence of the objects to which they purport to be referring itself depends upon the truth of Bio-rhythm Theory (or something like it).

What I am suggesting is thus this. Instead of saying that our hypothetical speakers think they are referring to bio-rhythms but are in fact referring to equivalence classes, we should say that they think they are referring to bio-rhythms and are in fact referring to *nothing*. I will return shortly to the question why this description of the situation should be preferred. First, we need to develop the view further.

The Projectibility View, the reader will recall, is that the sort of object to which an expression introduced by abstraction *in fact* refers is determined by which properties of the tokens are *in fact* Projectible. I have been arguing for the last several paragraphs that this cannot be right. But it is almost right. The right view is: The sort of object to which expressions introduced by abstraction *purport* to refer is determined by what properties of the tokens users of those expressions *presume* to be Projectible.

Presume in what sense? Contrary to what is usually supposed, an understanding of terms introduced by abstraction is not exhausted by a grasp of the criterion of identity for those terms. That is, abstraction principles are not, by themselves, (always) adequate to introduce a class of expressions. We must recognize a second component, which, borrowing a term from Quine, we might call the *Ideology* associated with the abstraction principle. The Ideology is not a theory about the objects whose names the abstraction principle characterizes, like Bio-rhythm Theory, but it is closely related to such a theory. The Ideology is something like a collection of properties of the tokens that use of the expressions introduced by the abstraction principle presupposes³¹ are Projectible. So, as said, it is not that the sort of object such expressions in fact denote is determined by what properties are in fact Projectible, but rather that the sort of object to which these expressions purport to refer is determined by which properties users of these expressions presuppose are Projectible.

One reason to recognize the independence of the Ideology from the abstraction principle with which it is associated is that there is no particular reason that the example discussed above in connection with day-persons had to take the form it did: Bio-rhythm Theory is only one of many theories which, had they been true, would have assured us of the existence of day-Projectible properties of persons. We might, instead, have considered a theory asserting that all persons born on the same day have common physiological features, be these gross anatomical ones or such properties as heart-rate, blood-pressure, and the like. Had such a theory been true, it seems to me that certain sorts of abstract objects would then have existed. On the other hand, I should want to deny that *bio-rhythms* should then have existed, if for no other reason than that it seems to me that both sorts of objects might have existed together. What *kinds* of objects day-persons are supposed to be is thus connected with the sort of properties they are supposed to have, and so with the Ideology that is associated with them. Only if one knows the associated Ideology can one know what kinds of objects are supposed to be denoted by the terms in question.

I said earlier that the Projectibility View wrongly supposes that it might be a substantial question to what sort of objects expressions of the form “day-person(*t*)” refer. I also said that there was a substantial question in the vicinity that had been misidentified. We can now see which question that is. On the view being developed, the Projectibility of the properties contained in

³¹I think we really do want the notion of presupposition here, but I am less sure which notion of presupposition we want. The general idea that understanding an expression of a certain type involves accepting certain presuppositions is nowadays fairly common. The advantage to using it here is that we do not need to modify the abstraction principles themselves. In particular, we can still have

$$\varphi(a) = \varphi(b) \equiv Rab$$

and allow that this might fail to be true because use of the expressions on the left-hand presupposes something that is not true.

the Ideology is a necessary condition of the truth of any sentence containing such expressions.³² There being such objects as those the introduced terms purport to denote thus depends upon the (actual) Projectibility of the properties contained in the Ideology. So the substantial question toward which the Projectibility View was groping is just that one: Whether *there are* any such objects as the ones the expressions in question purport to denote.

It is hard to see how it could be an intelligible question whether names of abstracta denote what they purport to denote. (I speak from my own experience.) Why? It is, I think, because we have tended to suppose that the only constraint on the assignment of reference to expressions introduced by abstraction is given by the abstraction principle itself. If so, then the only constraint on the *existence* of those objects is also given by the abstraction principle, and if there is an assignment that verifies the abstraction principle, then that is enough to show that the types in question exist. If so, however, then we are stuck with the proliferation problem, because it is hard to see how any one abstraction principle could be preferred over any other. But that, in turn, is because we miss the importance of the associated Ideology. Once it is in place, there is an additional constraint on the sort of object to which terms introduced by abstraction purport to refer, and there is therefore an additional constraint on the existence of referents for those terms. In short: The objects to which the terms refer must be ones that would not exist unless the properties in the Ideology were in fact Projectible.

Another reason the presence of the Ideology is easily overlooked is that it is often obvious from context what it is supposed to be. Consider, for example, Frege’s definition of names of what he calls “orientations”: The orientation of a plane *a* is the same as the orientation of a plane *b* if, and only if, *a* is parallel to *b* (Frege, 1980, §64). Immediately upon encountering this definition, one immediately knows precisely what Frege means to be talking about: It is obvious what sort of object an orientation is supposed to be. By contrast, few will have discerned a conception of ‘bio-rhythms’ in the abstraction principle introducing names of day-persons. Whence the difference? It lies in the fact that, in the case of orientations, it is obvious what the Ideology is intended to be. As is clear from the context of Frege’s discussion, the theory of orientations is to be a geometrical theory; so orientations are geometrical objects.³³

What, then, of duds, which haven’t been mentioned for a while? Does the view outlined support the intuition that there are no such objects? From the perspective of the present view, we can only say that the question is ill-posed: No Ideology associated with names of duds has been offered, so we have no idea what kind of object it is on whose existence we are being asked

³²This is too strong, as I have occasionally indicated parenthetically. We surely want to allow for near misses and second chances, so it is really only something like the Projectibility of enough predicates sufficiently like the ones contained in the Ideology that is required. We shall return to this point.

³³Special thanks to George Boolos here.

to pronounce. The intuition that there are no such objects comes, I would suggest, from the thought that, whatever duds are supposed to be, they are supposed to stand to their tokens in the same sort of relation that editions, words, and the like, stand to theirs. And what I have been suggesting is simply that there is more to there being a relation “of that sort” between a type and its tokens than what is captured in an abstraction principle. Bearing *that* kind of relationship to one’s tokens amounts to having properties determined by *R*-Projectible properties of the tokens. And, confronted with such a peculiar example as that of duds, we are essentially unable to imagine any sensible Ideology that might be associated with it. And if there is no sensible Ideology, then there are no duds, *whatever* sort of object they might be supposed to be.

So, to sum up, again, the view I’m suggesting has two parts. First, the sort of object to which names introduced by abstraction purport to refer is determined by the Ideology associated with those names; a specification of the Ideology is as fundamental a part of the explanation of those names as the specification of the abstraction principle itself. Second, there being any such objects as those to which the names so introduced purport to refer depends upon whether the properties contained in the Ideology are in fact Projectible properties of the tokens.

Why should one accept this view? We need to distinguish two sorts of issues here. The ontological issue concerns what we *ought* to say about the various cases at which we have been looking: Whether we should say there are no such objects as duds, but that there are editions; and, if so, whether we should say that editions are abstract and *sui generis* or, instead, are just books, equivalence classes, or what have you. There is also, however, a purely descriptive problem: To give some account of the felt difference between names of duds and day-persons, on the one hand, and names of editions, on the other. However seriously it should be taken, there is an intuition to say, in the former cases, either that there are no such objects or that they are ‘just’ something else. And one might want to understand the source of this intuition, understand to what features of the (imagined) use of these names we are responding, whatever one’s view about the ontological question.

So far as the descriptive problem is concerned, it seems to me that the view I’ve described does get a lot of the intuitions right, in all their varied and confusing forms. Moreover, in my experience, many people’s first response to the suggestion that editions might be books is that, if they were, there would be all sorts of things one could say about editions which just don’t seem to make much sense. The view on offer explains the relevance of this thought. It isn’t that no sense can be made of the claim that *Die Grundlagen* weighs twelve ounces, but rather that regarding such a claim as intelligible involves committing oneself to the copy-Projectibility of properties like *weighs twelve ounces*, contrary to obvious fact.³⁴

³⁴This sort of point seems to me relevant to certain sorts of examples Chomsky often mentions when deriding semantics, examples like: John wrote a book that weighs over a pound. What is

As concerns the ontological problem, the view’s getting these intuitions right constitutes at least some evidence in its favor. My primary goal here has been, as I said at the outset, not so much to argue for but simply to characterize a view that responds, in some principled and discriminating manner, both to the problem of trans-sortal identification and to the proliferation problem. That the view does so is another consideration in its favor.

But once one has started to take the significance of the Ideology associated with a given abstraction principle seriously, it is natural to start to wonder if it is not the Ideology that is really fundamental, and so if the concentration on abstraction principles, common to Naïve Platonism and Semantic Reductionism, is misplaced. Consider, for example, the case of linguistic types, such as sentences. Discussion of linguistic objects in the literature on ontology generally supposes that linguistic tokens can unproblematically be identified as physical entities, bits of ink or chalk, say, or disturbances in the air.³⁵ The problem then seems to be to identify the abstraction principle that underlies our use of names of types. In fact, however, the question what sorts of things sentences are is an empirical question. We have good reason to believe, for example, that there are *two* sentences both of which are written “Flying airplanes can be dangerous”, not one sentence that sometimes means that it can be dangerous to fly airplanes and other times means that airplanes in flight may pose a hazard. If so, then it is hard to see how tokens *could* be just blotches of ink or chalk, for there is no difference between the blotches that are tokens of the one sentence and the ones that are tokens of the other. Or perhaps the right thing to say is just that the relation *is of the same type as* is not going to be explicable in terms of anything like *having the same shape*, as philosophers since Frege have often assumed, nor in terms of any other purely physical properties.³⁶

The evidence that there are two sentences written “Flying airplanes can be dangerous” lies deep in linguistic theory: in phenomena whose explanation, given other principles that seem well supported, seems to demand such a conception of what sentences are. What informs our conception of when we have two sentences and when we have one is thus *the kinds of things we can say about sentences*. The Ideology here is thus not, as our earlier examples may have led one to suppose, simply independent of the abstraction principle. On the contrary, the Ideology *shapes* the abstraction principle, whose discovery is thus an empirical enterprise, not one of *a priori* conceptual analysis.³⁷

puzzling about this example is that the book John wrote is an edition; the one that weighs over a pound is a copy; and yet “that” seems to refer to the latter but to have as antecedent the former.

³⁵Unhappiness with that conception surfaces from time to time, however, for example, as in Kaplan’s (1990) paper ‘Words’.

³⁶There are interesting discussions of these issues, to which I am greatly indebted, in the paper by Bromberger (1992) mentioned earlier, and in another that he wrote with the phonologist Morris Halle (Bromberger and Halle, 1992).

³⁷Cognate points concerning concrete objects are made by Wiggins in *Sameness and Substance* (Wiggins, 1980, 2001). Indeed, the notion of an Ideology plays a role in the present view not

This conclusion does not really conflict with the neo-Fregean approach to abstracta, though it does cast it in a different light. One might have thought it did conflict with it, on the ground that Wright and Hale tend to emphasize the idea that abstraction principles can be freely stipulated, in the sense that someone who makes such a stipulation incurs no obligation to guarantee that there really are objects of the sort required. Well, the view being elaborated here is, of course, in tension with part of this view, since I am claiming that we do not have any *a priori* guarantee that objects of the right sort exist. But that observation is really intended as a friendly amendment. Wright often emphasizes that, in laying down an abstraction principle, we do not thereby stipulate the existence of objects but simply the truth-conditions of certain sentences (Wright, 2001a, 162; 2001b, 311). But it is hard to take this suggestion seriously when the truth-condition of “ $\varphi(a) = \varphi(a)$ ” ends up being: *Raa*, which is itself guaranteed to be true in virtue of the fact that $R\xi\eta$ is an equivalence relation. The view I am developing here makes space for the sort of distinction on which Wright correctly wants to insist.

Moreover, the emphasis on free stipulation of abstraction principles is dispensable. It is an artifact of idealizations Wright and Hale make, ones that allow us to abstract away from messy questions about what the abstraction principle for linguistic types (say) really is. To think that this idealization was important to Wright and Hale would be to overlook a point made earlier: The question whether the *same type* relation can be defined in other terms, so central to earlier discussions of nominalism, has turned out to be a distraction. Frege’s great insight was that our understanding of what sentence-types are is intimately bound up with our appreciation of when we have one sentence-type and when we have two. It does not matter whether *same sentence* can be defined in a way that would be comprehensible to someone ignorant of sentences, or of linguistic-types generally, and it does not matter either whether our grasp of “same sentence” is or is not independent of our grasp of what can be said about sentences.

What is true, however, on my view, is that an understanding of what a sentence is does not issue simply from a grasp of when we have one sentence and when we have two but also requires an appreciation of what sorts of things can be said about sentences. But that, though certainly in conflict with Naïve Platonism, is not, so far as I can see, in any conflict at all with the spirit of Frege’s position. Or so I am about to argue.

3.4 THE JULIUS CAESAR PROBLEM

As I have just been saying, the idea of accounting for our capacity to make reference to abstract objects by means of abstraction principles has its origins in Frege. But, in the end, Frege rejects the view that names of abstracta can

unlike that played by a ‘principle of activity’ in Wiggins’s view. I very much wish I could make the relation between these two notions precise.

be explained in this way, his reason being that such a view does not resolve the Caesar problem: This explanation does not determine the truth-values of such sentences as “the edition of a is Julius Caesar” nor, for that matter, the truth-values of any sentences of the form “the edition of a is t ”, unless t is itself of the form “the edition of b ”. There has been little agreement, however, about why this is supposed to be a problem, and, in closing, I would like to explain what I take the problem to be and how the foregoing might allow us to address it.

Frege raises the Caesar objection against a proposed answer to the famous question of *Die Grundlagen* §62: “How, then, are numbers to be given to us, if we cannot have any ideas or intuitions of them?” The proposal he considers is that we may explain the senses of identity statements in which number-words occur by means of an abstraction principle, namely, what I call ‘HP’: The number of F s is the same as the number of G s just in case the F s and G s are in one-one correspondence. So the view against which the Caesar objection is offered is this: We recognize numbers as the referents of expressions of the form “the number of F s”, and our understanding of these expressions consists (in large part) in our grasp of HP. Frege’s objection to this view is that HP “will not, for instance, decide for us whether [Caesar] is the same as the [number zero]...” (Frege, 1980, §66).³⁸

It is generally supposed that Frege is here raising an instance of the problem of trans-sortal identification, and in some sense that must surely be true. What is not so widely noted, however, is that Frege takes for granted we *do* recognize that Caesar is not a number. His objection is *not* that HP does not decide the truth-values of *all* ‘mixed’ identity statements. If that were the objection, then our intuitions about whether Caesar is a number would be irrelevant. But the problem Frege raises is not, say, that HP does not decide whether zero is the singleton of the null set—which, on Frege’s explicit definition, it happens to be. Rather, the problem is that HP does not decide a question about which he takes us to have strong intuitions: Whatever numbers may be, Caesar is not among them. If so, then one might suppose that any complete account of our apprehension of numbers as objects must include an account of how we distinguish people from numbers, and Frege’s objection to HP, regarded as constituting a complete explanation of how we apprehend numbers as objects, is that it alone yields no such explanation. That is why Frege writes: “Naturally, no one is going to confuse England with the direction of the Earth’s axis [or Caesar with the number of non-self identicals]; but that is no thanks to our definition of direction [or of Number]” (Frege, 1980, §66). So our ‘definition’ of direction or number must include more than just an abstraction principle.

I have suggested here that a full account of our understanding of expressions introduced by abstraction must include an account of the Ideology associated with those expressions, and that the wanted explanation of why people

³⁸Note that the objection is not so much that the abstraction principle fails to decide the truth-value of this sentence, but that it fails to give any sense to it at all.

are not numbers can be given in terms of the Ideology. My own view, however, is that the Caesar problem is not just about trans-sortal identification. Rather, Frege uses the Caesar problem to raise another, more semantical sort of issue.³⁹ HP, and other abstraction principles, are supposed to feature identity statements on their left-hand sides, and they are supposed to explain a range of *terms*: expressions that are to be treated, semantically, as purporting to refer to objects. Only if expressions of the form “the number of *F*s” are terms in this sense can our capacity to refer to numbers be explained in the way Frege considers.

Now, if expressions of the form “the number of *F*s” are to be understood as referring to objects, and if a statement of the form “the number of *F*s is the same as the number of *G*s” really is an identity statement, then a complex predicate such as “ ξ is the number of *G*s” must itself be a predicate that is true or false of objects. And if HP is truly sufficient to ground an understanding of a class of terms purporting to refer to objects, then it must also issue in an understanding of such predicates as being true or false of objects. To put the point differently, if “the number of *F*s” is truly a semantic constituent of “the number of *F*s is the number of *G*s”, then it must be replaceable by a variable: It must be an intelligible question whether the open sentence “*x* is the number of *G*s” is true or false of any particular object, independently of how that object is given to us. This is the point of Frege’s remarks about sentences of the form “*q* is the direction of *a*”: The definition of directions gives us no purchase whatsoever on the question whether this open sentence is true or false of England (Frege, 1980, §66). And, indeed, it gives us no more purchase on the question whether it is true or false *of* the direction of the Earth’s axis, independently of how that object is given to us. It is only when we imagine the object *given as* the direction of the Earth’s axis that we understand the question (Frege, 1980, §67).

The reason this matters so much to Frege is that it implies that HP does not suffice to explain the concept of cardinal number. The local point is that the abstraction principle for directions does not suffice to explain the concept of direction. If it did, Frege tells us, then “*q* is the direction of *a*” could be explained by distinguishing cases: If *q* is not a direction, then the proposition is false; if it is a direction, then it is the direction of some line, and the abstraction principle will take over from there (Frege, 1980, §66). Another route to the same point proceeds from the observation that the complex predicate “ $\exists x(\xi = \text{dir}(x))$ ” clearly defines “ ξ is a direction”. So we understand this concept only if we understand the complex predicate in question. Indeed, one might reasonably suppose that, if the abstraction principle issues in a grasp of the concept of direction, then it does so *via* an understanding of “ $\exists x(\xi = \text{dir}(x))$ ”. But this complex predicate just embeds the very context “ $\xi = \text{dir}(x)$ ” that is causing all the trouble.

³⁹To the best of my knowledge, this point was first made by Parsons in ‘Frege’s Theory of Number’ (Parsons, 1995).

So we have, on the one hand, something that the abstraction principle does not give us: an understanding of the question whether “ $\xi = \text{dir}(a)$ ” is true or false of an object, independently of how that object is given to us. And we have, on the other hand, the fact that the abstraction principle does give us an ability to understand the question whether the result of substituting a certain sort of expression—one of the form “ $\text{dir}(t)$ ”—for the placeholder is true. The contrast is thus between what we might call an ‘objectual’ understanding of such predicates and a merely ‘substitutional’ understanding. But this issue, as should now be apparent, is itself intimately connected with our understanding of quantification over directions. Consider, for example, the claim that the direction of a exists:

$$\exists y(y = \text{dir}(a))$$

This too embeds the very sort of context we have been discussing. And the question whether we are equipped with an objectual understanding of the quantifier thus becomes the question whether we understand this sentence as saying that there is an *object* that is the direction of a or merely as saying that there is an expression—*not* necessarily a term, in any real sense—whose substitution for “ y ” would yield a truth.⁴⁰ Among the many things at issue here is thus the question whether abstraction principles can provide for an understanding of objectual quantification over the range of objects whose names they are supposed to introduce.

What I want to suggest, in closing, is that this problem—that of securing an objectual interpretation of quantification—may actually be a form of the proliferation problem. For note that the substitutional reading of the quantifier is as available in the case of day-persons or duds as it is in the case of directions, numbers, or editions. Indeed, it is often suggested that quantification over types should be explained in terms of quantification over tokens, so that “ $\exists y(y = \text{dir}(a))$ ” would just become “ $\exists l(\text{dir}(l) = \text{dir}(a))$ ” and so “ $\exists l(l \parallel a)$ ”. But then, of course, by the same reasoning “ $\exists y(y = \text{dud}(a))$ ” becomes “ $\exists d(\text{dud}(d) = \text{dud}(a))$ ” and so “ $\exists d(Qda)$ ”, and we are back in the soup. We need to keep “ $\exists x(x = \text{dir}(b))$ ” from just reducing to “ $\exists l(l \parallel b)$ ”, but it will be impossible to do so as long as the truth of “ $a \parallel b$ ” is supposed to be wholly adequate for the truth of “ $\text{dir}(a) = \text{dir}(b)$ ”. Giving some teeth to the notion of existence is thus of a piece with securing an objectual interpretation of the quantifiers: In both cases, the issue is one of making “there is a direction...” have some force not had by “there is a line...”.

Defenders of the neo-Fregean approach to abstracta often say that the core of their view is that, in introducing a class of expressions by abstraction, we assume no new epistemological burden as regards the existence of referents for those expressions: The epistemology of direction-talk is supposed to be wholly reduced to the epistemology of line-talk and knowledge of the abstraction

⁴⁰It’s important to remember here that the substitution class for a substitutional quantifier can consist of almost anything, e.g., parentheses or suffixes (Kripke, 1976).

principle that links the two sorts of discourse. I am, clearly, denying this. But my view is simply not vulnerable to the strongest objection typically brought against such denials, namely: The question whether there are directions makes sense only if we know what the existence of directions involves; but the neo-Fregean account of what it is for there to be directions has been rejected, and no alternative has been offered (Wright, 2001a). I have said as precisely as I can what the existence of directions and other types involves. And, in so far as my view imposes an epistemic obligation on those who claim the existence of directions, sentences, editions, or numbers, it is not a philosophical obligation but a (broadly) scientific one, and it is one that can clearly be met, at least in some cases.⁴¹

REFERENCES

- Benacerraf, Paul (1973). ‘Mathematical truth’, *Journal of Philosophy* 70:661–680.
- Berkeley, George (1930). *A Treatise Concerning the Principles of Human Knowledge*. London: Open Court.
- Boolos, George (1998). ‘Must we believe in set-theory?’, in Richard Jeffrey (ed.), *Logic, Logic, and Logic*. Cambridge, Mass.: Harvard University Press, pages 120–132.
- Bromberger, Sylvain (1992a). *On What We Know We Don’t Know: Explanation, Theory, Linguistics, and How Questions Shape Them*. Chicago: University of Chicago Press.

⁴¹Thanks to George Boolos, Sylvain Bromberger, Kate Elgin, Michael Dummett, Michael Glanzberg, Steve Gross, Bob Hale, Jim Higginbotham, Thomas Kuhn, Charles Parsons, Hilary Putnam, Gideon Rosen, Bob Stalnaker, Jason Stanley, and Crispin Wright for discussions that contributed enormously to this paper. The comments of two referees, and the editors, also helped improve the paper, as did the contributions of the members of a seminar I taught at Harvard in 1994.

As may already be clear, this paper has been a long time coming. The central ideas first appeared in the last of the three papers comprising my dissertation (Heck, 1991). I had intended to prepare those papers for publication shortly after defending. But it was at about this time that George Boolos asked me a question that led me to start studying Frege’s *Grundgesetze der Arithmetik*, and by the next fall my research energies were directed quite firmly elsewhere.

It was not until the Spring of 1994 that I would return to the topics of my dissertation. That semester, I taught a graduate seminar on ontology and produced revised versions of the two papers that had concerned abstract objects. (The other (Heck, 1998) was on vague objects.) I submitted the first of these, “Syntactic Reductionism”, for publication that fall, but it was rejected on the ground that my argument for nominalism was full of holes. Depressed by that outcome, I again set the papers aside, since the second required the first. I entertained ideas of expanding them into a short book but before long was back at work on Frege’s philosophy of mathematics.

I might not have published either paper except for an invitation from Stewart Shapiro to contribute to a special issue of *Philosophia Mathematica* devoted to the neo-Fregean program; it was there that ‘Syntactic Reductionism’ finally appeared (Heck, 2000). Yet again, though, the desire to do something new took precedence over the desire to revise something old, and the present paper was set aside until I was invited to contribute something to the present volume. I am grateful to the editors for the opportunity, finally, to see this paper published.

- Bromberger, Sylvain (1992b). ‘Types and tokens in linguistics’, in Bromberger (1992a), pages 170–208.
- Bromberger, Sylvain and Morris Halle (1992). ‘The ontology of phonology’, in Bromberger (1992a), pages 209–228.
- Burgess, John P. and Gideon Rosen (1997). *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. Oxford: Oxford University Press.
- Dummett, Michael (1981). *Frege: Philosophy of Language*, 2d edition. Cambridge, Mass.: Harvard University Press.
- Dummett, Michael (1991). *Frege: Philosophy of Mathematics*. Cambridge, Mass.: Harvard University Press.
- Frege, Gottlob (1980). *The Foundations of Arithmetic*, 2d revised edition, tr. by J. L. Austin. Evanston, Ill.: Northwestern University Press.
- Gödel, Kurt (1990). ‘What is Cantor’s continuum problem?’, in Solomon Feferman, *et al.* (eds.), *Collected Works*, volume 2, 3d edition. Oxford: Oxford University Press, pages 254–270.
- Hale, Bob (1988). *Abstract Objects*. Oxford: Blackwell.
- Hale, Bob and Crispin Wright (2001a). *The Reason’s Proper Study*. Oxford: Clarendon Press.
- Hale, Bob and Crispin Wright (2001b). ‘To bury Caesar...’, in Hale and Wright (2001a), pages 335–396.
- Heck, Jr., Richard G. (1991). *Logic, Semantics, Ontology*. PhD thesis, Massachusetts Institute of Technology.
- Heck, Jr., Richard G. (1998). ‘That there might be vague objects (so far as concerns logic)’, *The Monist* 81:277–299.
- Heck, Jr., Richard G. (2000). ‘Syntactic reductionism’, *Philosophia Mathematica* 8:124–149.
- Hodes, Harold T. (1990). ‘Ontological commitment: Thick and thin’, in George Boolos (ed.), *Meaning and Method*. Cambridge: Cambridge University Press, pages 235–260.
- Kaplan, David (1990). ‘Words’, *Proceedings of the Aristotelian Society*, sup. vol. 64:93–119.
- Kripke, Saul (1976). ‘Is there a problem about substitutional quantification?’, in Gareth Evans and John McDowell (eds.), *Truth and Meaning: Essays in Semantics*. Oxford: Oxford University Press, pages 325–419.
- Lewis, David (1986). *On the Plurality of Worlds*. Cambridge, Mass.: Blackwell.
- Lewis, David (1999). ‘Noneism or allism?’, in *Papers in Metaphysics and Epistemology*. New York: Cambridge University Press, pages 152–163.
- Linnebo, Øystein (2006). ‘Epistemological challenges to mathematical platonism’, *Philosophical Studies* 129:545–574.
- Parsons, Charles (1980). ‘Mathematical intuition’, *Proceedings of the Aristotelian Society* N.S. 80:145–168.
- Parsons, Charles (1995). ‘Frege’s theory of number’, in William Demopoulos

- (ed.), *Frege's Philosophy of Mathematics*. Cambridge, Mass.: Harvard University Press, pages 182–210.
- Russell, Bertrand (1903). *The Principles of Mathematics*. Cambridge: Cambridge University Press.
- Russell, Bertrand (1905). ‘On denoting’, *Mind* 14:479–493.
- Sullivan, Peter and Michael Potter (1997). ‘Hale on Caesar’, *Philosophia Mathematica* 3:135–153.
- Uzquiano, Gabriel (2005). ‘Semantic nominalism’, *Dialectica* 59:265–282.
- Wiggins, David (1980). *Sameness and Substance*. Cambridge, Mass.: Harvard University Press.
- Wiggins, David (2001). *Sameness and Substance Renewed*. New York: Cambridge University Press.
- Wright, Crispin (1983). *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.
- Wright, Crispin (2001a). ‘Field and Fregean platonism’, in Hale and Wright (2001a), pages 153–168.
- Wright, Crispin (2001b). ‘Is Hume's principle analytic?’, in Hale and Wright (2001a), pages 307–332.