

Mereological Sums and Singular Terms

KATHRIN KOSLICKI

1. Introduction

One prominent answer to the question of what the correct theory of part-hood and composition is for material objects, which has been embraced by three-dimensionalists and four-dimensionalists alike,¹ is that these objects are best viewed as ‘mereological sums’, ‘fusions’, or ‘aggregates’, according to a particular, standard, conception of mereology, viz., the family of systems formally analogous to Henry Leonard’s and Nelson Goodman’s ‘Calculus of Individuals’ (Leonard and Goodman 1940) and referred to in Simons (1987) as ‘Classical Extensional Mereology’ or (CEM). (A brief synopsis of the basic principles of standard mereology is given in section 2.) Mereological sums, according to this standard conception, are like sets in that their existence and identity depend on nothing more than the existence and identity of the parts that compose them; no structural requirements are set on the manner in which these parts must be arranged.

¹ Three-dimensionalism (also known as ‘endurantism’) and four-dimensionalism (also known as ‘perdurantism’ or ‘the doctrine of temporal parts’) are competing theories concerning the persistence of ordinary material objects over time. According to the four-dimensionalist, objects persist over time by perduring, i.e., by having temporal parts, in addition to their ordinary spatial parts, at all those times at which they exist. The three-dimensionalist, on the other hand, holds that ordinary material objects persist by enduring, i.e., by being, as they say, ‘wholly present’ at each time at which the object exists. For detailed discussion and references, see Sider (2001). From a three-dimensionalist perspective, perhaps the most well-known defense of the thesis that material objects are standard mereological sums can be found in Thomson (1983); see also Thomson (1998). For prominent representatives of this approach within the four-dimensionalist tradition, see Lewis (1986b), (1991); Sider (2001).

The relative merits of standard mereology in the context of giving an analysis of material objects have received quite a bit of attention from metaphysicians in recent years.² A question that has not been pursued to the same degree, however, is what sort of *semantic* repercussions a commitment to mereological sums in the standard sense might have in particular on the predicted behavior of singular terms and our practices of using such terms to refer to objects. The basic problem in this connection for those who believe that standard mereology has a place in a metaphysical analysis of material objects, it seems to me, is that our practice of using singular terms to refer to objects, at least on the face of it, pretty obviously does not track mereological sums: the objects we are interested in and care about, and hence which are represented in our discourse by means of singular terms, generally do not exhibit the sorts of persistence conditions we would expect them to have if these objects were in fact correctly analyzed as mereological sums in the standard sense. We are much more likely to talk about individual trouts and turkeys than we are to talk about trout-turkeys, to use David Lewis's well-known example.³

The apparent mismatch between our actual referential practices and the persistence conditions attributed to material objects by the supporters of standard mereology puts these philosophers, other things being equal, at a disadvantage compared to those whose ontology matches more closely the observed behavior of singular terms, as they are commonly used in ordinary discourse. In particular, as will become apparent below, the proponents of standard mereology suffer from self-imposed handicaps with respect to the sorts of constraints on reference to which they can appeal in the context of offering an account of why our singular terms have the content-determining powers they appear to have. If they believe, as David Lewis does (see especially Lewis 1983, 1984), that some constraints on reference originate from the side of the *object* referred to (as opposed to the speaker or the speaker's linguistic community) and that the constraints in question are not exclusively *causal* in nature, then additional machinery is needed to make up somehow for the fact that ordinary language-users speak as though there are trouts and turkeys but not trout-turkeys. To this end,

² See for example Fine (1982), (1999); Harte (2002); Johnston (2002); Koslicki (2008); van Inwagen (1981), (1987), (1990a), (1990b), (1993), (1994), (2002).

³ A trout-turkey is a mereological sum consisting of, say, the still-attached upper half of a trout and the still-attached lower half of a turkey.

David Lewis leans heavily on his distinction between *natural* and *non-natural* properties: a highly eligible referent, according to Lewis, tends to be one whose boundaries are well-demarcated with respect to the highly natural properties; the distinction between natural and non-natural properties, in his view, ultimately derives from fundamental physics.

While I am sympathetic to Lewis's position that a credible account of reference will need to appeal to the presence of well-demarcated boundaries among the objects referred to, I am skeptical that his already heavily burdened natural/non-natural distinction can bear the additional weight of drawing boundaries in the right places around the referents of our singular terms: for conglomerations of undetached rabbit-parts, for example, will in general contrast with their environment just as much as single rabbits do, when it comes to such considerations as, say, the density of matter, the relative abundance of chemical properties, favored loci of causal chains, and the like, in the relevant regions of spacetime. Thus, Lewis's natural/non-natural distinction among properties, as I will argue below, is not sufficient to avoid Quinean indeterminacy for singular terms. Those who are in the business of giving an analysis of constructions involving full-fledged predication, as opposed to constructions that presuppose merely spatial overlap among denotations (see Evans 1975), will thus want to go in for an ontology that places more stringent structural constraints on the referents of singular terms than are supplied by standard mereology.

2. Standard Mereology

(CEM) is a very simple, elegant and surprisingly powerful theory.⁴ It requires only a single primitive notion in terms of which the remainder of the mereological concepts utilized by the theory can be defined. In its standard formulations, (CEM) consists of a mere three axioms; all other statements of the theory follow as theorems from the definitions and axioms of the system. The single primitive can be chosen to be proper parthood, proper-or-improper parthood, overlap, disjointness, or sum; the other notions are definable in terms

⁴ The first formulation of (CEM) appears to have been given by Stanislaw Leśniewski, informally in Leśniewski (1916) and formally in Leśniewski (1927–1930). The classical statement of (CEM) in English is Henry Leonard and Nelson Goodman's 'Calculus of Individuals' (Leonard and Goodman 1940).

of whichever one is taken as primitive. Identity is either assumed as given or (more controversially) as definable in terms of the primitive mereological notion. Algebraically speaking, while parthood is a mere partial ordering, (CEM) has the strength of a complete Boolean algebra, with the zero element deleted.

In terms of parthood and overlap, the notion of a mereological sum (aggregate/fusion) for example can be defined as follows:

Mereological Sum: S is a mereological sum of some objects, $x_1 \dots x_n$, just in case S has all of $x_1 \dots x_n$ as parts and has no part that does not overlap any of $x_1 \dots x_n$.

A very accessible formulation of (CEM), which is slightly different from, but formally equivalent to, that of Leonard and Goodman (1940), is given in Lewis (1991), where the three basic axioms of standard mereology are stated informally as follows:

Axiom 1 (Unrestricted Composition): Whenever there are some objects, then there exists a mereological sum of those objects.

Axiom 2 (Uniqueness of Composition): It never happens that the same objects have two different mereological sums.

Axiom 3 (Transitivity): If x is part of some part of y, then x is part of y.

Given Unrestricted Composition and the Uniqueness of Composition, mereological sums can now be seen to be analogous to sets in certain respects: by Unrestricted Composition, there is a sum whenever there is a plurality of objects (any plurality), just as there is a set for any plurality of objects (modulo the paradoxes); and by the Uniqueness of Composition, the identity of a sum depends on nothing more than the identity of its parts, just as (by the Axiom of Extensionality) the identity of a set depends on nothing more than the identity of its members. In other respects, however, mereological sums and sets are not alike: in particular, standard mereology makes no room for a distinction analogous to that between subset and membership. For given that all the entities quantified over in standard mereology are of the same ontological type, viz., the type, individual, standard mereology only recognizes a single sum that can be formed from any given plurality of individuals.

3. The Utility of Proper Names

If we take our common practice of using singular terms, in particular proper names, at face value, it would appear that such expressions often serve as names for certain commonly recognized kinds of things: among the proper names we typically employ, for example, are names for human beings, fictional or real ('Hamlet', 'Bismarck');⁵ names for pets ('Fido'); names for places ('Paris'); names for significant historical events ('the Civil War'), periods ('the Renaissance'), or important agents within history ('the Holy Roman Empire');⁶ names for artistic creations ('Star Wars'); brand names ('Duracell') or names for products ('iPod'); and so on.⁷

P.F. Strawson writes that a proper name is called for among a circle of language-users, when there is frequent need or occasion to refer to a particular; when there is interest in the continuing identity of the particular from occasion to occasion; and when there is no short description or title for that particular available to members of the linguistic circle in question (Strawson 1974, p. 42). In such a situation, Strawson remarks, the use of a proper name for the object in question is appropriate:

So it is convenient to have in circulation in such groups a tag, a designation, which does not depend for its referential or identifying force upon any particular such position or relation, which preserves the same referential force through its object's changes of position or relation and has the same referential force for communicators who know the object in different connections and for whom quite different descriptions would be uppermost. (Strawson 1974, pp. 45–46)

⁵ For the purposes of this discussion, I take names like 'Hamlet' or 'Bismarck' to be names of human beings, rather than persons, only in order to sidestep complicated issues concerning personal identity which are tangential to my main concerns.

⁶ I take it that such expressions as 'the Civil War', 'the Renaissance', or 'the Holy Roman Empire', are names, even though they are prefixed by the definite article; this issue is, however, also not central to the current discussion.

⁷ By pointing out that names are typically used as names for certain kinds of objects, I do not intend to commit myself to a descriptive view of names, i.e., one according to which a designation of the kind of object referred to in some way becomes part of the meaning of a name (e.g., as though 'human being' somehow makes it into the semantic content of 'Bismarck'). Speakers can, and certainly have been, wrong about what kind of object a given name refers to: witness, for example, 'Venus', 'Hesperus', and 'Phosphorus', which were at first intended to be used as names for a star; it was only found out later that these names were in fact names for a planet. In this way, the case of proper names is parallel to that of natural kind terms, as they were analyzed by Kripke (1971), (1980) and Putnam (1975): the extension of a term like 'water', on this approach, is determined by means of a same-kind relation of some sort which instances of the liquid bear to certain initial samples; but speakers can certainly be either ignorant or wrong about the nature of the same-kind relation which in fact determines the extension of their term.

It is convenient in such situations to be able to take recourse to a proper name, Strawson proposes, because there is in such situations “no one particular fact about any member, no particular relationship in which [the object referred to] stands, such that it is on account of that fact or relationship that [it] is a standing object of reference for members of this group” (Strawson 1974, p. 46). Contrast a situation of this type, for example, with one in which it is sufficient to single out an object of interest by referring to it as ‘your cigar’ or ‘the yellow one in the middle’; clearly, designations of this sort, while they might be perfectly sufficient to achieve one’s communicative purposes at a particular moment, are of a much more fleeting nature and would not be able to preserve continuity of reference through, say, changes of ownership, color, or position. Thus, the purpose of the availability of a proper name for a particular group of language-users, if Strawson’s observations are correct, is typically to offer the members of a language-group a means of referring to a particular object of interest that is both (i) *stable* and (ii) *shared*: (i) it is stable in the sense that continuity of reference is preserved in the face of changes in the object referred to or among the members of the language-group; (ii) it is a shared means of referring to the object in question in the sense that the availability of the proper name may bridge differences in the information that the members of the group may possess regarding the object to which they are referring.⁸

If a certain expression is commonly used as a name for a human being, pet, place, etc., then it may be employed as a means not only of communicating

⁸ A very different, substitution-theoretic, account of what distinguishes singular terms from other types of expressions is given in Brandom (1994). In his account of the distinction between singular terms and general terms, Brandom focuses on a contrast he perceives between the inferential roles played by these two types of expressions. Singular terms, in Brandom’s view, give rise to *symmetric* substitutional inferences: for example, from ‘Mark Twain authored *Huckleberry Finn*’ we may infer ‘Samuel Clemens authored *Huckleberry Finn*’; moreover, so Brandom notes, this inference is *reversible*. In contrast, substitution-inferences which turn on the predicates contained in statements, in Brandom’s view, tend not to be symmetric in the same manner: thus, from ‘Benjamin Franklin walked’, we may infer ‘Benjamin Franklin moved’; but this inference is not reversible. As it stands, however, the contrast in question is more of an artifact of how the particular examples Brandom considers are picked than an indication of any genuine semantic difference between singular terms and general terms. For the object referred to by ‘Mark Twain’ is numerically identical to the object referred to by ‘Samuel Clemens’, whereas the denotation of ‘moves’ neither coincides with nor is included in the denotation of ‘walks’; rather, the denotation of ‘walks’ is included in the denotation of ‘moves’. Inferences which turn on predicates whose denotations do coincide (e.g., perhaps ‘is a doctor’ and ‘is a physician’) will be similarly reversible. Of course, Brandom cannot help himself to this explanation of the contrast between the two inferences in question, since for him the metaphysical notion of an object is supposed to be derivative of the semantic distinction between singular terms and general terms. Thus, he cannot take for granted that identities among the denotations of terms are already given to us, as a datum to which the semantics may refer in order to draw a distinction between singular and general terms.

information of interest about the object in question at a single time (as when we say, for example, that Bismarck unified Germany in 1871); it may also be used to track the same human being, pet, place, etc., over time (as when we say, for example, that Bismarck was Chancellor of Germany from 1871 to 1890).⁹ Taking our common practice of using proper names again at face value, these expressions certainly seem to preserve stability of reference, among other things, in the face of changes in the parts of the object referred to. Since we commonly regard human beings as the sorts of things that can persist through the gaining and losing of many of their parts, we similarly take the name, 'Bismarck', given that this expression is standardly employed as the name for a human being, to preserve its reference despite the fact that the object it refers to can be expected to undergo changes with respect to its parts. On the other hand, given that the parts of a human being must be arranged in a particular characteristic manner in order to result in a functioning specimen of its kind, there are other features about the parts of a human being, in particular their manner of arrangement, which we do expect to remain relatively constant in order for a single human being to persist from one time to another; this stability, again, transfers straightforwardly to names for human beings. Thus, it would seem to be a bit of a stretch for example to employ the name, 'Bismarck', to refer to the scattered bits of soil into which this man's remains have decomposed after Bismarck's death: since these scattered bits of soil currently fail to compose a human being, it would be decidedly odd to apply the same name to these scattered bits of soil as to the human being from whose corpse they originated. A name like 'Bismarck', therefore, at least on the face of it, seems to refer to the sort of object that belongs to a commonly recognized kind (e.g., the kind, *human being*), that can survive gaining and losing parts over time, and whose parts at each time at which the object exists must be arranged in the particular manner characteristic of the kind of object at issue.

⁹ This point is often expressed by saying that proper names (or common nouns, for that matter) are associated not only with criteria of application, but also with criteria of identity. When an expression functions as a proper name, speakers who have mastered the use of this expression must be able to recognize the object referred to by the name as the same again on different occasions. According to Dummett (1973), this mastery on the part of the speaker also involves being able to group objects referred to by the use of names into *categories*: for every proper name, he writes, must belong to some *range* of names, all of which are associated with the *same* criterion of identity, viz., the criteria of identity governing the particular category of object to which the referent of the name in question belongs; to recognize that the name in question belongs to this range is also, for Dummett, part of what must be mastered by a speaker in order to grasp the sense of a name (Dummett 1973, p. 76; see also Frege 1884, §62).

If, on the other hand, material objects are construed as mereological sums and those singular terms that purport to refer to material objects are construed as referring to mereological sums, then we would expect the referents of these singular terms to have the persistence conditions associated with mereological sums. But the existence and identity of mereological sums, as we have seen in the previous section, depends on nothing more than the existence and identity of their parts: thus, a mereological sum, S_1 , and a mereological sum, S_2 , are numerically identical over time or at a time just in case the parts that compose S_1 are numerically identical to the parts that compose S_2 ; no requirements are set on the manner in which these parts are arranged. If the supporter of standard mereology takes the route, then, of viewing singular terms, as they are standardly employed in ordinary discourse, as referring to mereological sums, there appears to be, at least on the face of it, a clash between the persistence conditions that are attributed by the supporter of standard mereology to the objects referred to by our singular terms and the way in which speakers in fact employ these singular terms to track objects over time.

To be sure, this is not the end of the story and there is additional machinery which the supporter of standard mereology can invoke to try to capture the behavior of singular terms, as they are commonly used in ordinary discourse.¹⁰ However, when evaluating the relative merits of standard mereology as the correct theory of parthood and composition for material objects, we should take note of the fact that there is at least a *prima facie* mismatch between the persistence conditions apparently exhibited by the referents of singular terms, as they are standardly used in ordinary discourse, and the persistence conditions we would expect these objects to have if ordinary singular terms in fact referred to mereological sums. And this apparent mismatch requires the supporter of standard mereology to introduce some

¹⁰ For example, a natural way to go for the supporter of standard mereology would be to adopt the view that singular terms, as they are standardly employed in ordinary discourse, refer only to certain *phases* in the life of a mereological sum or to certain *collections* of mereological sums, different ones at different times. The first line is congenial to four-dimensionalist supporters of standard mereology like Lewis and Sider; the second to three-dimensionalists like Thomson. According to the first approach, a singular term like 'Bismarck' might be construed for example as referring only to those phases in the career of a particular four-dimensional space-time worm during which the ordinary spatial parts of each temporal slice of this worm are arranged in such a way that they sustain a living human organism. According to the second approach, a singular term like 'Bismarck' would be construed as referring to different mereological sums at different times (corresponding to the changes in parts which Bismarck undergoes during his lifetime), all of which are composed of parts arranged at those times in such a way that they can sustain a living human organism.

further *semantic* machinery, in addition to the metaphysical apparatus to which he has already committed himself (e.g., counterpart theory) in order to overcome the apparent divergence in persistence conditions between mereological sums and material objects as we standardly conceive of them.

4. Constraints on Reference

The solution to this apparent difficulty concerning the observed behavior of our ordinary singular terms which is favored by David Lewis relies prominently on his distinction between natural and non-natural properties (see especially Lewis 1983, 1984).¹¹ Constraints on reference, in Lewis's view, fall into two kinds: (i) those originating from facts about *use*; and (ii) and those originating from facts concerning the *intrinsic eligibility* of what is referred to. The first constraint locates the source of at least some of the content-determining powers of the language in the speaker or the speaker's linguistic community: this constraint holds that content is to be determined at least in part by maximizing fit with how speakers or their linguistic communities tend to use the expressions and statements they utter. The second, and more controversial, constraint on reference traces some of the content-determining powers of the language to the intrinsic nature of the referent itself. Following this constraint on reference, we are to choose among two different rival interpretations of a language based on how well each does in connecting the content of what we say to a special group of elite properties, the natural properties. The distinction between natural and non-natural properties, in Lewis's view, derives ultimately from fundamental physics, at least if the latter proves to be successful.¹²

¹¹ Page numbers from Lewis (1983) and (1984) in what follows come from the reprinted versions of these papers in Lewis (1999).

¹² Lewis considers and rejects two other approaches to reference: global descriptivism and the causal theory of reference (see Lewis 1983, pp. 45–47; as well as Lewis 1984 for a more detailed discussion). According to global descriptivism, an intended interpretation of our language is one which satisfies a certain purely descriptive condition, viz., it must be compatible with ideal theory (an idealized descendant of our current body of theory). This sort of approach to reference, in Lewis's view, feeds straight into Putnam's model-theoretic argument (Putnam 1978, 1980) and hence leads to radical indeterminacy of reference, as the purely descriptive constraints imposed by this approach to reference are powerless to distinguish among the countless available candidate interpretations for a given language. Alternatively, a purely causal approach, in Lewis's view, only raises the further question of what it is about the referent of an expression that selects it as the object that satisfies the causal constraints in question. Lewis's natural/non-natural distinction and the concomitant notion of intrinsic eligibility is intended as a response to this latter question.

Naturalness, as conceived of by Lewis, is a matter of degree. Thus, the perfectly natural properties are those singled out by fundamental physics; others, such as those which go into the primitive vocabulary of chemistry or biology, for example, are less than perfectly natural, but still quite a bit more natural than, say, blueness. Blueness in turn is quite a bit more natural than, say, grueness, the (alleged) property objects have if they are green, when examined before some fixed time t , or blue, when not examined before t .

How do we compare the degree of naturalness exhibited by the less than perfectly natural properties? Here, Lewis appeals to two assumptions: (a) that the less than perfectly natural properties will be definable in terms of the perfectly natural properties; and (b) that we can measure a property's degree of naturalness by means of the complexity of the definition that can be given of it in terms of the perfectly natural properties. Both (a) and (b) strike me as quite implausible; however, since the issue of definability has already been extensively discussed by others, I will for present purposes bypass the question of how exactly degree of naturalness is to be measured.¹³

Naturalness is initially conceived of by Lewis as first and foremost relevant to the determination of predicate-meaning. Thus, a set of objects from the domain is more or less eligible to count as the extension of a one-place predicate from the language under consideration depending on whether and to what degree its members share a natural property; similarly for many-place predicates and relations. An interpretation for the language consequently is eligible depending on whether and to what degree it assigns eligible extensions to the predicates of the language.

Even though naturalness and eligibility are conceived of by Lewis as applying most directly to properties and hence to the determination of predicate-meaning, he does suggest a way in which these notions tend to have a sort of ripple-effect that extends to singular terms and their referents as well:

Naturalness of properties makes for differences of eligibility not only among the properties themselves, but also among things. Compare Bruce with the

¹³ For critical discussion of Lewis's account of relative naturalness as ease of definability in perfectly natural terms, see for example Hawthorne (2006), (2007); Schaffer (2004); Sider (1995), (1996); Williams (2007).

cat-shaped chunk of miscellaneous and ever-changing matter that follows him around, always a few steps behind. The former is a highly eligible referent, the latter is not. (I haven't succeeded in referring to it, for I didn't say just which such chunk 'it' was to be.) That is because Bruce, unlike the cat-shaped chunk, has a boundary well demarcated by differences in highly natural properties. Where Bruce ends, there the density of matter, the relative abundance of the chemical elements,...abruptly change. Not so for the chunk. Bruce is also much more of a locus of causal chains than is the chunk; this too traces back to natural properties, by the considerations of the previous section [cf., the role of natural properties with respect to causation and laws]. Thus naturalness of properties sets up distinctions among things. The reverse happens also. Once we are away from the perfectly natural properties, one thing that makes for naturalness of a property is that it is a property belonging exclusively to well-demarcated things. (Lewis 1983, pp. 48–49)

Lewis suggests in this passage that highly eligible referents are those whose boundaries are well-demarcated with respect to highly natural properties and which act as the loci of causal chains more so than their more gerrymandered cousins. Thus, a highly eligible referent, such as Bruce the cat, in Lewis's view, will distinguish itself from a less highly eligible referent, such as any one of the presumably indefinitely many cat-shaped chunks of miscellaneous and ever-changing matter that inhabit spacetime regions in the vicinity of Bruce the cat, with respect to for example the density of matter, the relative abundance of chemical elements, and the like, that can be found in these regions of spacetime. Moreover, Bruce the cat will more often be identifiable as the locus of causal chains than the more gerrymandered cat-shaped chunks of matter in the vicinity. Lewis's suggested extension of the notions of eligibility and naturalness from properties and predicates to objects and singular terms, if successful, is thus intended to yield a distinction between objects whose boundaries are well-demarcated with respect to natural properties and those whose boundaries are not well-demarcated in this way; the former tend to be more highly eligible referents than the latter.

5. Evans on Predication

That the presence of objects with well-demarcated boundaries is vital to an analysis of the distinction between singular terms and general terms is also the upshot of an argument provided by Gareth Evans in a completely

different setting (Evans 1975).¹⁴ Evans's main target in this context is Quine's inscrutability of reference thesis: Evans wants to establish, contrary to Quine, that only certain ones among the stimulus-synonymous hypotheses concerning the nature of the entities ranged over by the expressions of a particular language in fact do justice to the behavior of these expressions with respect to predication and negation.¹⁵

Suppose for example that we are confronted with a language containing expressions, $G_1, G_2, G_3, \dots, G_n$, which, when presented as one-word questions, are assented to by speakers of this language when material objects of various sorts are present; these expressions are, when presented as one-word sentences, stimulus-synonymous with the English expressions 'A rabbit!', 'A man!', 'A hut!', and so on. In addition, this language contains expressions, $F_1, F_2, F_3, \dots, F_n$, which, when presented as one-word questions, are assented to when certain general features are present in the environment; these expressions are, when presented as one-word sentences, stimulus-synonymous with the English expressions 'White!', 'Warm!', 'Furry!', 'Bloodstained!', and so on.

According to Quine's inscrutability of reference thesis, the expressions of the G -type, in being stimulus-synonymous with English one-word sentences signifying the presence of material objects, as in 'A rabbit!', are equally stimulus-synonymous with English sentences signifying the presence of undetached parts of material objects, or temporal stages of material objects, as in 'Undetached rabbit-parts!' or 'Rabbit-stages!':

For, consider 'gavagai'. Who knows but what the objects to which this term applies are not rabbits after all, but mere stages, or brief temporal segments, of rabbits? In either event the stimulus situations that prompt assent to 'Gavagai' would be the same as for 'Rabbit'. Or perhaps the objects to which 'gavagai' applies are all and sundry undetached parts of rabbits; again the stimulus

¹⁴ Page numbers from Evans (1975) in what follows are taken from the reprinted version of this paper in Evans (1985). Evans, of course, should not be read as agreeing with Lewis's claim that the boundaries in question which are relevant to an analysis of the distinction between singular terms and general terms originate from the division between natural and non-natural properties; in fact, since he is sympathetic to an ontology consisting of enduring structured wholes, he believes that the requisite boundaries are drawn for him, so to speak, by the persistence conditions of the objects referred to by singular terms in ordinary discourse and the kinds to which these objects belong.

¹⁵ Expressions are stimulus-synonymous when they have the same stimulus-meaning; that is, when they coincide with respect to 'the native's total battery of present dispositions to be prompted to assent to or to dissent from' the expressions, when presented as one-word sentences (Quine 1960, p. 39).

meaning would register no difference. When from the sameness of stimulus meanings of 'Gavagai' and 'Rabbit' the linguist leaps to the conclusion that a gavagai is a whole enduring rabbit, he is just taking for granted that the native is enough like us to have a brief general term for rabbits and no brief general term for rabbit stages or parts. (Quine 1960, pp. 51–52)

But suppose now that, in addition to single word sentences in which the expressions of the F-type and those of the G-type occur separately, the language under discussion also allows for complex expressions of the type, $F_i G_i$, i.e., expressions stimulus-synonymous with English sentences like 'A white rabbit!'. Evans argues that a difference does now surface between the various interpretational hypotheses entertained by Quine. For suppose we interpret a particular expression of this foreign language, F_i , as dividing its reference over all and only the world's white things; and we interpret a particular expression, G_i , as dividing its reference over all and only the world's rabbits (undetached rabbit-parts, rabbit-stages, and so on). Then how are we to interpret the compound expression, $F_i G_i$, stimulus-synonymous with the English 'A white rabbit!', 'White undetached rabbit-parts!', 'White rabbit-stages!', and so on? In particular, how are we to read the semantic contribution of F_i as combining with that of G_i , in order to predict correctly the assent- and dissent-behavior prompted by the compound expression, $F_i G_i$, in speakers of this language?

In some cases, the semantics of compound expressions might be correctly captured by requiring nothing more than that there be (reasonably sized) spatial overlap between the denotations of the expressions of the F-type and those of the G-type. Take for example a complex expression of this language, $F_2 G_2$, which, let's say, is stimulus-synonymous with the English 'Red water!' ('Red undetached water-parts!', 'Red water-stages!', etc.); suppose moreover that the expression, $F_2 G_2$, would be assented to by speakers of the language in a scenario in which a few drops of red ink are immersed into a swimming pool (Evans 1975, p. 32). A case of this sort, Evans argues, does not yet require the introduction of the machinery that is, in his view, involved in full-fledged predication, namely that of objects and their boundaries as well as the idea that properties attributed to such objects must be distributed within their boundaries in a characteristic manner. Together, the apparatus of objects and their boundaries yields what is required for identity and individuation (or what Quine would call the 'apparatus of divided reference'): what is needed

for saying when one object of a certain kind ends and another begins and when we encounter the same object again on a later occasion.

In the scenario just imagined, on the other hand, no single object seems to be clearly indicated whose instantiation of redness within its boundaries would account for the assent behavior prompted among the speakers of this language by an utterance of $F_2 G_2$:

‘Red Water?’ Certainly. But which object is it, even roughly, whose satisfaction of the predicate ‘Red’ makes the remark true? Is it the whole pool, or just the water immediately diluting the ink, or one of the indefinitely many intermediate alternatives? The language provides us with no way of answering these question, to which we must somehow find an answer if the construction is predication. (Evans 1975, p. 32)

I take it that what Evans means by saying that the construction at issue need not be interpreted as involving full-fledged predication is that we could for example get by (i.e., capture the assent- and dissent-behavior generated by it in speakers of the language correctly) even if we interpreted the occurrences of the expressions, F_2 and G_2 , in the compound expression, $F_2 G_2$, as both playing the role of singular terms and the ‘tie’ between them as requiring merely that their denotations overlap spatially. There are various options as to what, under these circumstances, we might take these expressions to denote, if they are in fact construed as playing the role of singular terms in the compound expression, $F_2 G_2$. We could for example go the Quinean route (Quine 1960, Ch. 3, §19ff.) and interpret F_2 and G_2 as denoting a particular scattered object or mereological sum, viz., the totality of the world’s red things and the totality of the world’s water, respectively.¹⁶ In that case, we would predict that an utterance of the expression, $F_2 G_2$, will prompt assent-behavior among speakers of the language just in case the totality of the world’s red things and the totality of the world’s water share one or more parts in the region of spacetime with which the utterance in question is concerned.

If Evans is correct, however, this interpretational route will not succeed across the board. For in order for an utterance of $F_1 G_1$ (stimulus-synonymous with the English ‘A white rabbit’, ‘White undetached rabbit-parts!’, ‘White

¹⁶ Evans (1975, p. 32, n. 8) indicates that he favors an analysis in terms of Hartry Field’s apparatus of ‘partial denotation’ and ‘partial satisfaction’ (see, e.g., Field 1974), but does not expand on the details of this idea.

rabbit-stages!', and so on) for example to prompt assent-behavior from speakers of this language, more may be required in certain circumstances than that the denotation of F_1 merely spatially overlap that of G_1 . To bring out this point, Evans quite ingeniously imagines a scenario in which several brown rabbits are so arranged that their white tails are contiguous (Evans 1975, p. 34). Certainly, the condition that there be spatial overlap between the world's white things and the world's undetached rabbit-parts is satisfied in this scenario in the particular region of spacetime with which the utterance in question is concerned. In fact, the overlap in question between the totality of the world's white things and the totality of the world's undetached rabbit-parts may even be large enough in size and appropriately shaped to compare with that of one or more single white rabbits. And yet speakers of this language may nevertheless on such occasions dissent from an utterance of $F_1 G_1$, since the imagined scenario does not contain one or more single rabbits which have whiteness distributed within their boundaries in a way characteristic of white rabbits: a brown rabbit with a white tail does not make a white rabbit; nor do several such brown rabbits with their white tails arranged contiguously.

Consider also the interaction between several expressions of the F-type and an expression of the G-type, as in $F_3 F_1 G_1$, which, let's say, is stimulus-synonymous with the English 'A bloodstained white rabbit!' ('Bloodstained white undetached rabbit-parts!', 'Bloodstained white rabbit-stages!', and so on). As Evans remarks, it is not sufficient for speakers to assent to $F_3 F_1 G_1$, that they assent to both $F_1 G_1$ and $F_3 G_1$ separately, since $F_3 F_1 G_1$ may only hold in situations in which it is the very same rabbit within whose boundaries whiteness and bloodstainedness are distributed in a characteristic fashion. Moreover, what counts as a characteristic distribution of the property within the object's boundaries is different in the case of whiteness and bloodstainedness: in order for a rabbit to count as white, it must be (roughly) predominantly white; but in order for it to count as bloodstained, only very few of its parts need to be bloodstained. Finally, in order for a single rabbit to count as both white and bloodstained, it is not necessary that any of its white parts be bloodstained; thus, in order for $F_3 F_1 G_1$ to prompt assent among the speakers of this language, it is not necessary that there be spatial overlap among the denotations of all three expressions, F_3 , F_1 and G_1 , since the compound expression does not require that there be any single undetached rabbit-part that is both white and bloodstained.

The notion of a boundary drawn around single objects, as well as the characteristic distribution of a property within such a boundary, according to Evans, is necessary not only for a proper understanding of predication, but is also required to capture the semantics of negation in a language like ours (Evans 1975, pp. 35–36). For we may for example adopt a rule governing negation which would predict that an object satisfies ‘not-white’ just in case the object in question does not satisfy ‘white’. However, in order for this rule to capture the force of negation in the language, it must be the case that when ‘white’ and ‘not-white’ are simultaneously applied to one and the same object, a contradiction results. And in order for this outcome to obtain, according to Evans, it is again necessary that we can tell the difference between where a single object of a certain kind in the domain stops, where another begins, and how we are supposed to find the feature in question distributed with respect to the boundaries of this object. For a contradiction will only arise when it is said of one and the same object that a given property both is and is not distributed within its boundaries in a certain characteristic manner. In this way, Evans’s point concerning the characteristic distribution of a certain property within the boundaries of a single object applies not only to the proper analysis of constructions involving predication, but also to those involving negation.

6. Objects and their Boundaries

Although Evans’s discussion is mostly focused on the role of general terms within the language, the lesson we learn from his arguments straightforwardly extends to singular terms as well. For reference and predication are but two different sides of a single semantic coin: whatever apparatus is required for a proper understanding of constructions involving full-fledged predication will also be implicated in a genuine contrast between singular terms and general terms. Singular terms, after all, are just those expressions that are used to refer to the objects over which the general terms contained in the language divide their reference. And general terms are ways of expressing the similarities and differences present among the referents of singular terms. Thus, if the presence of objects with genuinely demarcated boundaries is necessary for a proper understanding of predicative constructions in a language like ours, then this same apparatus will carry over also

to what is required for an understanding of the referential properties of full-fledged singular terms contained within such a language.

As was noted in section 3, our ordinary practice of using singular terms, in particular proper names, if taken at face value, leads us to believe that we typically employ such terms as names for commonly recognized kinds of things: names for human beings, fictional or real; names for places; names for significant historical events, periods, or agents within history; names for artistic creations; brand names or names for products; and so on. Moreover, as Strawson remarks, proper names in particular appear to provide a community of language-users with a means of referring to an object that is both stable (in that it preserves sameness of reference in the face of various sorts of changes) and shared (in that it bridges differences in the informational background available to members of the group).

In one respect the conclusion reached by Evans in his discussion of Quinean indeterminacy, viz., that a proper semantic analysis of a language like ours requires a domain of objects with well-demarcated boundaries, accords well with the conclusion reached by David Lewis, as observed in section 4. For Evans and Lewis both converge on the idea that our singular terms, somehow, successfully manage to refer to objects with well-demarcated boundaries. Thus, whatever one's particular approach to reference, it should come out at the end of the day, in both Evans's and Lewis's view, that a particular cat, say, is a more highly eligible referent for the name, 'Bruce', when employed in its intended role as a name for a cat, than a cat-shaped chunk of miscellaneous and ever-changing matter that occupies a region of spacetime which only partially overlaps but does not coincide with that occupied by the cat in question.

The means, however, by which Evans and Lewis envision these boundaries to be drawn are quite different. For Lewis, highly eligible referents are those which are well-demarcated by highly natural properties. And while Lewis's natural/non-natural distinction is intended to apply to objects and singular terms as well, it does so only derivatively: the distinction is formulated in the first place with properties and predicate-meanings in mind. For Evans, on the other hand, the domain itself, over which the singular and general terms of the language range, already comes pre-divided, so to speak, into highly eligible referents. That is, the boundaries of these objects are not drawn in a derivative manner by means of a distinction whose primary realm of application is something other than objects and singular terms;

rather, the idea of a single object with its boundaries, capable of enduring through space and time, is taken as fundamental.¹⁷

But the question of how our singular terms successfully manage to pick out the highly eligible referents from among all the possible referents is not even worth asking unless the domain over which our singular terms and general terms range is construed in such a way that it contains ineligible or less highly eligible referents alongside the highly eligible ones. If such gerrymandered objects never make it into the domain over which our singular and general terms range to begin with (e.g., because a particular ontology does not count such ‘non-entities’ as lying within the realm of things that exist), then the need for a distinction between eligible and ineligible referents vanishes and, with it, the need to account for the observation that speakers tend to ignore the ineligible or less highly eligible referents in their ordinary discourse. We should thus count it as among the advantages of the type of approach to reference endorsed by Evans over that endorsed by Lewis that such an approach does not need to settle the question of how to distinguish the highly eligible from the less highly eligible referents.

Lewis is for independent reasons already committed to a domain of objects for singular terms and general terms which includes, in addition to cats and other such commonly recognized kinds of things, also all manner of more or less gerrymandered chunks of miscellaneous and ever-changing matter in the vicinity of these commonly recognized kinds of things. This commitment follows from Lewis’s endorsement of the principle of unrestricted mereological composition, according to which for any plurality of objects, there exists a mereological sum of these objects (Lewis 1986b, pp. 211ff.). Recall from our earlier discussion (at the beginning of section 4) that Lewis identifies two plausible sources towards which we should look in our search for constraints on reference: (i) use and (ii) intrinsic eligibility. The first, which locates the source of at least some of the content-determining powers of our terms in the speaker or the speaker’s linguistic community, is not directly relevant to our current topic of discussion and hence has not played into the foregoing remarks. The second, intrinsic eligibility, is supposed to trace some of the content-determining powers of our language to the

¹⁷ An approach of this kind is also taken by Strawson; in *Individuals*, for example, he remarks as follows: “Material bodies, in a broad sense of the word, secure to us one single common and continuously extendable framework of reference, any constituent of which can be identifyingly referred to without reference to any particular of any other type” (Strawson 1959, p. 54).

intrinsic nature of the referent itself. But now, given Lewis's acceptance of standard mereology, we know that the referent itself, on Lewis's conception, does not have the power to do this work on its own: rather, the work of providing highly eligible referents with well-demarked boundaries must be accomplished by a different piece of apparatus, the natural/non-natural distinction, whose primary target is properties and predicate-meanings and which, if successful, yields distinctions among objects and singular terms only derivatively.

One kind of handicap the Lewisian imposes upon himself, then, which others manage to avoid, is that he is on independent grounds committed to an ontology which creates a need for a distinction between more or less highly eligible referents. Those, on the other hand, who embrace a less promiscuous ontology, into which the gerrymandered chunks never find their way to begin with, have an easier time in at least this respect: they do not to the same extent face the need to distinguish the highly eligible referents from the less highly eligible ones. And, for this reason, foregoing a commitment to standard mereology has the advantage, among other things, that the source for one kind of constraint on reference really can be located in the referent itself, genuinely and not merely in a derivative manner.

7. Quinean Indeterminacy and the Natural/ Non-Natural Distinction

I began by noting that there seems to be a *prima facie* mismatch between the persistence conditions apparently exhibited by the referents of singular terms, as they are standardly used in ordinary discourse, and the persistence conditions we would expect these objects to have if our ordinary singular terms in fact referred to mereological sums. Lewis proposes to remedy this *prima facie* mismatch by means of his natural/non-natural distinction, which is formulated so as to apply directly to the extensions of predicates and only indirectly to the referents of singular terms. By appeal to the natural/non-natural distinction, Lewis hopes to be able to make out the difference for example between a highly eligible referent for the name 'Bruce', when employed in its intended use as a name for a particular cat, and the many not-so-highly eligible chunks of miscellaneous and ever-changing matter in the vicinity of Bruce, the cat. Of course, given Lewis's independent

commitment to standard mereology, Bruce, the cat, in Lewis's view, is itself just a chunk of feline tissue.¹⁸ Nevertheless, so the proposal goes, the boundary between Bruce, the cat, and its surroundings is better demarcated by differences in the distribution of highly natural properties (such as density of matter, relative abundance of chemical elements, loci of causal chains, and so on) than the boundary between many of the other chunks of matter that are in the vicinity of Bruce, the cat, and their surroundings.

When I implied just now that there is just a single highly eligible referent for the name, 'Bruce', as it is employed on this particular occasion as a name for a specific cat, I was really oversimplifying. For, in addition to the components of Lewis's approach already mentioned, we must also add to the mix his 'Many, But Almost One' doctrine (Lewis 1993a). According to this doctrine, in any ordinary situation in which we may find ourselves employing a name like 'Bruce', in its intended role as a name for a particular cat, we will really be faced with a plurality of 'many-but-almost-one' cat-candidates that are present in the vicinity of what we would normally think of as a single cat (e.g., aggregate C_1 of feline tissue minus one particular molecule; aggregate C_2 of feline tissue minus a different molecule; and so on). Since these 'many-but-almost-one' cat-candidates, though strictly speaking numerically distinct, are mostly overlapping and overwhelmingly similar in all other respects relevant to our cat-related discourse, Lewis believes that our best strategy is simply to acknowledge a certain (moderate) degree of indeterminacy in our singular terms. Just as no semantic decision has been made or is called for concerning for example a term like 'heap', as to exactly how many grains constitute a heap, similarly no semantic decision has occurred, is likely to occur, or needs to occur, which would single out some particular one from among the many perfectly good cat-candidates in the vicinity of what we normally think of as a single cat. If we combine Lewis's 'Many, But Almost One' doctrine with the other components of his approach already cited, we arrive at the conclusion that, according to Lewis's view, while the natural/non-natural distinction may distinguish for us some of the roughly cat-shaped chunks of feline tissue in the vicinity of what we would normally consider to be a single cat, we will nevertheless

¹⁸ I am using the term 'feline tissue' in a way that presupposes that feline tissue is present only where (living) cats are present. Thus, given this use of the phrase, the presence of detached cat-parts, cat-corpses, and the like would not for present purposes be sufficient to ensure the presence of feline tissue. This assumption is of course beneficial to the Lewisian.

be left with a plurality of equally eligible potential referents for the name, 'Bruce', when employed on a particular occasion in its intended use as a name for a specific cat.

Earlier in this chapter, we discussed an argument mounted by Gareth Evans whose intended purpose was to show that in fact not all of the different interpretational hypotheses entertained by Quine do an equally good job in accounting for the assent- and dissent-behavior we would actually expect to observe on particular occasions in ordinary speakers of natural languages. In particular, Evans's argument against Quine purports to show that we can in fact tell the difference between the following two competing interpretational hypotheses: the hypothesis according to which a term like 'rabbit' divides its reference over undetached rabbit-parts and the hypothesis according to which the term in question divides its reference over enduring rabbits. For, among other things, if Evans is right, then the former hypothesis cannot, while the latter hypothesis can, account for the fact that speakers would not assent to a compound expression of the form, 'White rabbit?', when faced with several brown rabbits whose white tails are arranged contiguously.

Notice that Lewis's metaphysical and semantic commitments precisely put him into the camp of those who are subject to Evans's argument against Quine, since in Lewis's view cats just are (certain) mereological sums of undetached cat-parts. These mereological sums therefore provide the available choices of objects over which the general term, 'cat', can divide its reference and to which the singular term, 'Bruce', can refer, when it is employed in its intended use as a name for a particular cat. The main relevant difference between the interpretational hypothesis considered by Quine and that endorsed by Lewis is that Lewis invokes one additional piece of apparatus not present in Quine's account, viz., the natural/non-natural distinction. Our last remaining task at this point in the dialectic, then, is to see whether Lewis's appeal to the natural/non-natural distinction can help him overcome Evans's objection against Quine. Not surprisingly, my project in the remainder of this chapter will be to argue that Lewis's invocation of the natural/non-natural distinction by itself is not sufficient to ward him off against Evans's objection to Quine.

For the purposes of establishing this conclusion, I will make several simplifying assumptions, which I hope will not detract from the substantive issues at hand. Let's suppose, first, for the sake of simplicity, that our world is much

less complicated at the level of the perfectly natural properties than it in fact is. For example, we might imagine that the world at the most fundamental level is how the Presocratic philosopher, Empedocles, pictured it, with just four perfectly natural properties, viz., those corresponding to the general terms, ‘earth’, ‘air’, ‘fire’, and ‘water’. And let’s assume furthermore that all the other predicates in our language can be defined somehow (in more or less lengthy ways) in terms of these four predicates which denote the perfectly natural properties, together with numerical ratios which express the proportion of earth, air, fire, and water that is present in a given object at its most fundamental level of composition. Given this picture, it might turn out to be the case for example that the extension of the predicate, ‘cat’, consists of mereological sums which exhibit the following ratio of elements: two-eighths earth, two-eighths fire, three-eighths air, and one-eighth water. (We might think of this numerical ratio of elements as analogous for example to the chemical formula for water, viz., H_2O , which tells us not only what sorts of elements constitute water but also in what proportion these elements are present in this substance.) Finally, let’s make a further simplifying assumption, which Empedocles did not make, namely that we can think of the four elements as occurring in the form of small indivisible particles or atoms, so that the world according to this fictional account consists at the most fundamental level of four different kinds of atoms, viz., earth-, air-, fire-, and water-atoms.¹⁹

According to this fictional account, then, we have an item in the extension of the general term, ‘cat’, whenever we have earth-, air-, fire-, and water-atoms present in the proportion of two parts earth, two parts air, three parts fire, and one part water.^{20,21} Suppose now that a given speaker is employing the name, ‘Bruce’, on a particular occasion as a name for a

¹⁹ Even according to this highly simplified picture, there are many more possible dimensions along which entities at the most fundamental level could vary. For example, atoms could differ from one another not only with respect to the kind to which they belong (earth, air, fire, or water), but also along such dimensions as size, shape, weight, and so on. I will abstract away from these possible additional parameters of variation and assume, for the purposes of this argument, that all atoms belonging to a given kind are uniform with respect to these further parameters.

²⁰ Perhaps, our atomistic Empedoclean version of Lewis would want to add to this formula determining the extension of the predicate, ‘cat’, that the atoms in question must also be in close spatial proximity to one another. Without explicitly mentioning this additional condition, I will assume that it is met by the mereological sums in question which qualify as possible cat-candidates.

²¹ Since our real-world knowledge tells us that cats can in practice differ from one another, among other things, in terms of their size, it would not be plausible to assume (even within the confines of this highly simplified fictional story) that we could give a definition of what belongs into the extension of the predicate, ‘cat’, not just in terms of numerical ratios of elements, but rather in terms of precise numbers

specific cat. Assuming that the speaker manages to carry out his referential intentions successfully on this particular occasion, then (given our current assumptions) we know that at least one mereological sum must be present in the vicinity of the speaker which consists of two parts earth, two parts air, three parts fire, and one part water. In fact, in practice, whenever one such mereological sum is present in the vicinity of the speaker, there are many more such mereological sums which are present alongside the first one and which also exhibit the right ratio of elements to qualify as cat-candidates.²² For example, suppose C_1 is the aggregate of feline tissue which is composed of the following earth-, air-, fire-, and water-atoms: $[(E_1 + E_2) + (A_1 + A_2) + (F_1 + F_2 + F_3) + W_1]$. A different aggregate of feline tissue, C_2 , also currently in the vicinity of the speaker, might consist of the following parts: $[(E_2 + E_3) + (A_1 + A_2) + (F_1 + F_2 + F_3) + W_1]$. C_1 and C_2 , though strictly speaking numerically distinct, differ from one another only with respect to the presence or absence of a single atomic part: C_1 has E_1 among its proper atomic parts and lacks E_3 , while C_2 has E_3 among its proper atomic parts but lacks E_1 . Otherwise, they overlap with respect to all of their other proper atomic parts. It is therefore consistent with Lewis's philosophical commitments to assume that, due to their largely overlapping nature, C_1 and C_2 would be equally highly eligible referents for the use of the name, 'Bruce', in its use on that particular occasion. Moreover, if this scenario is anything like how matters would stand in a comparable situation in the actual world, then, in addition to C_1 and C_2 , the speaker would find himself on this particular occasion in the vicinity of many more such largely overlapping mereological sums, C_3, C_4, \dots, C_n , many of which would qualify as equally highly eligible referents for the name, 'Bruce', in its use on that particular occasion.

of atoms (e.g., along the lines of: a mereological sum, x , belongs into the extension of the predicate, 'cat', just in case x has among its atomic proper parts k_1 earth-atoms, k_2 air-atoms, and so on, where ' k_1 ', ' k_2 ', ..., denote natural numbers).

²² What if the scenario we are imagining is different from the actual world not only in that it is Empedoclean and atomistic at the most fundamental level, but also in the following respect: on most normal occasions, when speakers take themselves to be referring to a single cat, there is in fact only a single mereological sum present in their vicinity which exhibits the appropriate ratio of elements to be placed in the extension of the general term, 'cat'? Given my assumptions, such a world would have to be one in which cats for example can never sit on mats; for if they did, then additional mereological sums composed of the four kinds of elements would be in the vicinity of what ordinary speakers in that world think of as a single cat. What matters for our present purposes, however, is that it is in no way part of Lewis's account that the world in fact must be so highly austere as we just imagined it to be, in order for ordinary speakers to be frequently successful in their referential intentions.

Given the identity conditions governing mereological sums, each of C_1, C_2, \dots, C_n has its proper parts essentially. It is essential to the aggregate of feline tissue, C_1 , for example that its proper atomic parts are $E_1, E_2, A_1, A_2, F_1, F_2, F_3, W_1$. In the same way, the other mereological sums exhibiting the cat-ratio that are in the vicinity of the speaker on this particular occasion are also related essentially to their respective proper atomic parts. At the same time, given Lewis's commitments, it cannot be essential to Bruce, the cat, that it have any of these atoms among its proper atomic parts. For any of C_1, C_2, \dots, C_n , so we have assumed, are equally highly eligible referents for the name, 'Bruce', in its intended use on this particular occasion as a name for the specific cat in question; and C_1, C_2, \dots, C_n , while largely overlapping, each differ from one another at least with respect to the presence or absence of a single atomic proper part. Hence, it cannot be a requirement, in order for a given mereological sum to qualify as a potential referent for the name, 'Bruce', in its intended use on this particular occasion, that any specific earth-, air-, fire-, or water-atom present in the vicinity of the speaker, $E_1, \dots, E_n, A_1, \dots, A_n, F_1, \dots, F_n, W_1, \dots, W_n$, number among its proper atomic parts. All that is required of such sums is that they exhibit the ratio of elements characteristic of items in the extension of the general term, 'cat', viz., two parts earth, two parts air, three parts fire, and one part water; but, for any specific atom present in the vicinity of the speaker on this occasion, there is no requirement that *it* number among the proper atomic parts of the cat in question.^{23,24}

²³ Perhaps there are further requirements which a given mereological sum must satisfy in order for it to be selected as a potential highly eligible referent for the name, 'Bruce', in its intended use as a name for a particular cat. For example, such a candidate might have to bear certain cross-temporal relations to other cat-candidates which existed at earlier times and are connected to the speaker in specific ways. For example, events might have occurred prior to the present occasion which the speaker would describe in the following way, "A year ago, I adopted Bruce from the local animal shelter." In that case, in order for a mereological sum which currently exhibits the cat-ratio and is present in the vicinity of the speaker to qualify as a highly eligible referent for the name, 'Bruce', this sum would have to bear certain cross-temporal relations to earlier mereological sums exhibiting the cat-ratio which were present in the vicinity of the speaker when the events he now describes as "A year ago, I adopted Bruce from the local animal shelter" occurred. But we can ignore these possible additional constraints on potential candidates for the reference of the name, 'Bruce', for the time being. For the purposes of my current argument, we can assume for example that the speaker has laid eyes on Bruce, the cat, for the very first time just moments ago.

²⁴ One might think that the situation in which the Lewisian currently finds himself can be improved if we permit him to require of the mereological sums in question that they satisfy additional *structural* (i.e., *configurational*) constraints in order for them to qualify as proper cat-candidates. For example, we might

At this point, we are in a position to see why Lewis's account, despite its appeal to the natural/non-natural distinction, is still susceptible to Evans's argument against Quine. Given the semantic and metaphysical machinery at Lewis's disposal, nothing prevents us from interpreting both an apparent singular term like 'Bruce', in its intended role as a name for a particular cat, and a general term like 'cat' as merely a marker indicating the presence of cathood in the environment of the speaker. Our discussion above has brought out that a speaker's assent to an utterance of 'Cat?' or 'Bruce?', on the Lewisian picture, is not only compatible with, but also in practice a reliable indicator of, cathood being instantiated many times over, by the many numerically distinct but largely overlapping mereological sums exhibiting the cat-ratio that are present in the vicinity of the speaker on a particular occasion. What the Lewisian account has *not* told us (and, as far as I can see, lacks the resources to tell us) is how a term like 'cat' or 'Bruce' would manage to mark the boundary between a certain particular object satisfying the cat-ratio and some other such object or objects.

In this way, the Lewisian account wrongly assimilates the semantic behavior of a general term such as 'cat' to that of a paradigmatic mass term like 'water': for a mass term like 'water' precisely serves to indicate the presence of wateriness in the vicinity of the speaker, but without marking the boundary between a single particular object that exhibits wateriness and another. Similarly, while an apparent singular term like 'Bruce', in its intended role as a name for a particular cat, on Lewis's account, may serve to mark the presence of some specific mereological sums exhibiting the cat-ratio in the vicinity of the speaker, no one object in particular need be singled out by a felicitous use of the term in question.²⁵

consider adding a condition of the following sort to the formula by means of which the extension of the predicate 'cat' is specified: a mereological sum, x , belongs in the extension of the predicate, 'cat', iff (i) the proper parts of x are atoms in the ratio two parts earth, two parts fire, three parts air and one part water; and (ii) these atoms are arranged in *circular* fashion. Still the situation will be unchanged: for whenever a speaker appears to be using a term like 'Bruce' successfully, as a name for (what the speaker thinks of as) a single cat, he will in fact be in the presence of a plurality of mereological sums whose proper parts satisfy the cat-ratio and exhibit a circular arrangement. (Or, at least, a successful use of a term like 'Bruce', in its intended role, is compatible with the presence of a multiplicity of such circularly arranged cat-candidates in the vicinity of the speaker.)

²⁵ What also comes to mind in this context is Strawson's notion of a feature-placing expression, of which a paradigmatic example is 'It is raining' (see for example Strawson 1953). (Not surprisingly, the model of feature-placing expressions was used by Strawson and others as a potentially fruitful way of thinking of the semantic behavior of mass terms). While, as Strawson notes, a feature-placing construction, such as

On this picture, then, we would wrongly expect a compound expression like ‘White cat?’ to behave semantically along the lines of Evans’s ‘Red water?’, i.e., as producing assent on the part of a speaker when whiteness is instantiated, cathood is instantiated, and their respective instantiations overlap spatially. But what is needed, so Evans’s argument against Quine indicated, in order to capture the assent- and dissent-behavior we would in fact expect to observe among speakers, is that we are given the resources to tell when whiteness is distributed in certain characteristic ways within the boundaries of one or more particular cats in the vicinity of a speaker on a particular occasion.

Lewis’s natural/non-natural distinction by itself, as I have tried to establish in this section, is not the right kind of tool to close the gap between the mere presence of mereological sums instantiating cathood and single enduring cats. After all, the discovery that the chemical formula corresponding to water is ‘H₂O’ did not suddenly transform the term, ‘water’, which up to that point had been used standardly as a mass term and continues to be used so today, into a singular count noun. In the same vein, the ability to state what the cat-ratio is in terms of the perfectly natural properties also would not by itself yield well-demarcated boundaries around individual cats. But such well-demarcated boundaries are precisely what is needed to account for the semantic behavior of singular count nouns like ‘cat’ and genuine singular terms, such as ‘Bruce’, which can be used to pick out such individual cats.²⁶

‘It is raining’, marks the presence of a feature (viz., in this case, the occurrence of rain) in the environment of the speaker, its function is not to direct the discussion towards a particular object or objects in the vicinity of the speaker which are supposed to instantiate the feature in question. For example, it would sound rather strange to follow up someone else’s utterance of ‘It is raining’ with ‘What’s raining?’ or ‘Where is the thing that is raining?’ Rather, the use of ‘it’ in ‘It is raining’ is what is referred to by linguists as an “expletive ‘it’”: this kind of use of ‘it’ in a sentence does not play the role of a genuine subject term, indicating a particular object to which the predicate contained in the sentence ascribes a certain attribute. Rather, it is the whole sentence which marks the presence of a certain feature in the environment of the speaker, without containing a constituent whose purpose is to single out a particular object instantiating the feature in question.

²⁶ It is also unlikely that an appeal to Ramsey sentences would be of any assistance to Lewis’s particular package of views, in helping him to overcome Evans’s objection against Quine (see for example Lewis 1970, 1997). For Ramsey sentences take the form of existentially quantified statements and hence would also not be sufficient to enforce singular reference in the face of multiple, equally eligible, candidates for the denotation of any particular expression.

8. Conclusion

In sum, while I am sympathetic to Lewis's position that a credible account of reference and predication will need to appeal to the presence of well-demarcated boundaries among the objects referred to, I have tried to bring out why I am skeptical that his already heavily burdened natural/non-natural distinction among properties (i.e., predicate-meanings) can bear the additional weight of drawing boundaries in the right places around highly eligible candidates for the reference of our singular terms and the denotations of our general terms. If my arguments in this chapter are on the right track, then Lewis's natural/non-natural distinction among properties does not do the work required to avoid Quinean indeterminacy. But Evans's argument against Quine shows us why this result is undesirable: for, unless we manage to rule out some of the Quinean interpretational hypotheses, we in effect fail to predict correctly how actual speakers would behave on particular occasions of utterance.

Those who are in the business of giving an analysis of constructions involving full-fledged reference and predication, as contrasting with constructions which require merely spatial overlap among the denotations of expressions or the presence of a certain feature in the environment of the speaker, should thus feel tempted to go in for an ontology that places more stringent structural constraints than those that are supplied by standard mereology on the objects that are available as candidates for the denotations of our singular and general terms. When evaluating competing approaches to the metaphysics of material objects, we should also keep in mind that the choices we make on the metaphysical side may have unwanted semantic consequences.

