

The construction of ontological categories

Jan Westerhoff

Australasian Journal of Philosophy, 2004, forthcoming

2nd December 2003

Abstract

I describe an account of ontological categories which does justice to the facts that not all categories are ontological categories and that ontological categories can stand in containment relations. The account sorts *objects* into different categories in the same way in which grammar sorts *expressions*. It then identifies the *ontological* categories with those which play a certain rôle in the systematization of collections of categories. The paper concludes by noting that on my account what ontological categories there are is partially interest-relative, and that furthermore no object can belong essentially to its ontological category.

The aim of the following remarks is to give an explication of the notion of an ontological category. Although this notion is central to ontology and metaphysics (it is, after all, what these disciplines are about) it is hardly ever carefully discussed. References to ontological categories are made frequently and often at central places in philosophical discussion, but if we start to look at attempts of actually *defining* what an ontological category is we are likely to be disappointed. Moreover, the few attempts which we do find, and which try to define ontological categories by appealing to such notions as generality,¹ substitution² or criteria of identity³ all face a particular fundamental problem. The properties they (plausibly) ascribe to ontological categories are not had by the items the definitions pick out. Whatever it is that is defined there, it is certainly not the notion of an ontological category.⁴ In this paper I do not want to repeat criticisms of these definitions but set out to do what they do not achieve: describing a theory of ontological categories defining entities which actually have two central features of such categories (which are both intuitively plausible and frequently referred to in the literature) and which will also explain *why* ontological categories have these features.

¹Norton (1977).

²Sommers (1963).

³Dummett (1981, 75–80).

⁴An extended discussion of various attempted definitions of the notion of an ontological category is given in Westerhoff, ch. 2.

1 Two features of ontological categories

What are these two plausible features of ontological categories? They are the following:

DISTINCTNESS Not all categories are ontological categories.

STRUCTURE Ontological categories form a hierarchy.

Culinary implements, items of furniture, and different sorts of buildings all constitute kinds or categories, but these are not the categories ontology talks about. Ontology talks about abstract and concrete objects, about individuals, properties and relations, about substance and accident, about events, collections, tropes, facts and similar things, but never about categories as specific as knives and forks, tables and chairs, or churches and palaces.

This point is also noted by Hoffman and Rosenkrantz in their study of substance. They remark that

not all kinds divide up the world in ontologically important ways. Examples of kinds which are not ontological categories are: being a green thing, being a triangular thing, being a widow and (the disjunctive kind of) being a substance or an edge.⁵

This distinct place of ontological categories is stressed by Jonathan Lowe, who furthermore gives it an epistemological gloss.

Categorial structure is an *a priori* matter. By contrast, taxonomic relations between natural kinds are an *a posteriori* matter of natural law. [...] [A] categorial scheme, being *a priori*, should not be open-ended or provisional [...]⁶

Ontological categories are thus seen as differing in an important and well-defined way from other categories. Ontological and non-ontological categories are fundamentally different sorts of things and the dividing-line between them is not just vague. Unfortunately this feature is not reflected in the definition of ontological categories Hoffman and Rosenkrantz actually give⁷ (Jonathan Lowe does not himself propose an explicit definition of what an ontological category is but confines himself to enumerating various of the properties such categories should have.). Their definition does not give us the resources for formulating a clear criterion which allows us to distinguish the ontological from the non-ontological categories.

⁵Hoffman and Rosenkrantz (1997, 46–47).

⁶Lowe (2001, 185).

⁷Hoffman and Rosenkrantz (1994, 16–21).

To have such a criterion is of fundamental importance for ontology. Its absence implies that we do not have a clear answer to the question why certain distinctions are ontologically important while others are not. From an ontological point of view it makes no difference whether something takes ten seconds or ten minutes, but it *does* make a difference that it has a duration at all. On the other hand it makes no difference whether something is a spoon or a fork, *nor* whether it is some piece of cutlery or other. If we admit that there is no clear answer to the question why ‘having a duration’ is an ontologically relevant feature while ‘being a piece of cutlery’ is not, we may as well say that we do not know what we are talking about when doing ontology. This would be just as problematic as it would be for logic if we had no story to tell about the fundamental difference between logical and non-logical notions.⁸ The current account sets out to rectify this deficiency by describing an account which incorporates a clear distinction between ontological and non-ontological categories.

Apart from their distinctness from other kinds, ontological categories are generally not assumed to be simply unrelated, disjoint sorts of things. Perhaps the most obvious fact about systems of ontological categories is that they are generally arranged in some sort of structure, often even depicted in a tree-like diagram.⁹ This structure is usually understood in terms of class inclusion. Some categories contain others: the category of individuals contains those of abstract and concrete objects, abstract objects in turn contain properties, relations, sets and propositions, concrete objects material objects and events and so on. The basis of this structure is the assumption that objects in different categories can share properties: material objects and events are both located in time; properties and relations have the one-over-many property, sets and propositions cannot be created or destroyed.

This hierarchical structure of ontological categories is another feature traditional accounts fail to account for.¹⁰ They do not have the resources for explaining that ontological categories can be related in a tree-like manner. The categories they define are all of the same level of generality and therefore do not form a hierarchy. At best this constitutes a subset of the set of ontological categories, but it cannot possibly be comprehensive. This further divergence between plausible assumptions about ontological categories and the actual properties had by items which the supposed definitions of such categories pick out is the second problem which the account presented in this paper avoids.

Whether we believe that the examples of ontological categories men-

⁸For example in terms of invariance properties as discussed by Tarski (1986), McGee (1996), van Benthem (1989).

⁹For some examples see Chisholm (1996, 3), Lowe (2001, 181), Grossmann (1992, 87) and Hoffman and Rosenkrantz (1997, 48).

¹⁰It is particularly problematic for accounts based on criteria of identity proposed by Dummett (1981, 75–80), Wiggins (1971, 27–40), and Stevenson (1975).

tioned above are indeed such categories, or whether we think that they are actually related by containment in the way described is not important for our argument. These positions, all of which are defended at different places in the recent ontological discussion are merely introduced for the sake of illustration. My argument only demands that *any* theory of ontological categories implies the two features of distinctness and structure, but not that it does so in any specific way. My exposition of the notion of an ontological category is designed to do justice to these two metaontological assumptions, but not to any assumptions which may be held on the ontological object-level.

2 Two assumptions about types

My exposition of the notion of an ontological category is based on two methodological assumptions about types of things.

1. Types of things are defined in the same way as types of expressions.
2. Types of things can be systematized.

The first assumption suggests that our knowledge of the different kinds of objects there are in the world is generated in a way which is structurally analogous to that in which knowledge of the different types of expressions of our language is produced. Types of objects are like linguistic types, with the difference that they are not classes of expressions, but classes of non-linguistic items. Expressions are frequently sorted into types by considering substitution-patterns in sentences. I want to argue that constituents of states of affairs (which are bits of the world, not bits of language) can be sorted in a similar way by considering substitution-patterns in states of affairs.

Once this is done we have a collection of different types of things. But not all of these can be ontological categories; some of them will be the ‘too specific’ categories of tables, chairs etc. mentioned above. The idea is to identify the set of ontological categories with a special subset of the types of things. This subset plays a central rôle in the systematization of types. The concept of a systematization of types is not something which is readily encountered when considering linguistic types. The underlying intuition is that sets of types can exhibit redundancy: some of the types can be dispensed with because they (or something very much like them) can be generated from the remaining types together with some suitable construction operation. Systematizing a set of types therefore means cutting it down to a proper subset and then introducing a construction operation which allows us to regain the lost types, or at least some plausible substitutes for them. I want to identify ontological categories with types of things which are in this

subset, i.e. with those types of things which can in some suitable sense be regarded as having the capacity of generating all the other kinds of things in the world. We will also call this subset a *basis* of the set of types.

The remainder of this paper will be concerned with spelling out the understanding of types of things in terms of substitution-types and the concept of a systematization of a set of types in greater detail. It will then be shown how the notion of an ontological category which emerges from this satisfies the two conditions of distinctness and structure.

3 Types of expressions and types of things

3.1 Substitution in grammar and ontology

The substitutional account of types of expressions or grammatical categories is certainly the most straightforward and the most standard.¹¹ The idea behind it is that two words belong to the same grammatical category if we can take a sentence in which the first word occurs and substitute the second for it (and vice versa) and the result is still a grammatical sentence. Thus we see that only a certain group of words will go into the blank to make ‘... can be very annoying’ a grammatical sentence. Another group will go into ‘Many people like to ...’ and yet another into ‘Joe buttered the toast ...’. Members of the first group will be words like ‘linguistics’ or ‘satsumas’, while those going into the second blank will be e.g. ‘swim’ or ‘live’, while those in the final one are expressions like ‘quickly’ or ‘carefully’. We will call the first group nouns, the second verbs and the third adverbs. Of course the idea is not that for example *only* words which go into the first blank constitute the grammatical category of nouns, but rather that we can construct enough similar contexts which will *all together* allow us to pick out the nouns by intersubstitutability.

The idea of defining types in terms of substitution-classes also made its way into ontology, most importantly in the work of Gilbert Ryle.¹² Ryle, however, does not consider grammaticality but meaningfulness. He argues that certain substitutions in a sentence just affect its truth-value (that is they turn it from a true sentence into a false sentence or vice versa) while others affect its meaningfulness: they turn it from a meaningful sentence into a meaningless one. Or, to put it differently, only certain ways of filling the gap in a sentence-frame will produce something which can be true or false; for other substitutions the result will not have a truth-value. Ryle then equates ontological categories with types of things picked out by expressions intersubstitutable *salva congruitate*.¹³

¹¹Bradford (1988, 60, 62), Oliver (1999, 250–251).

¹²Ryle (1938), Ryle (1949).

¹³See Sommers (1963) for a further development of this idea.

We will not be concerned with this way of utilizing the notion of substitution for developing an account of types of things, not only because Ryle's account faces serious problems.¹⁴ We consider substitutions *in the world* rather than in language because this ensures a certain degree of independence of our account of ontological categories from linguistic considerations. Rather than considering cases where expressions cease to be meaningful when intersubstituting subexpressions we are concerned with collections of constituents of states of affairs which cease to 'fit together' into a state of affairs once certain constituents have been substituted for others.

Substitution in states of affairs

But what exactly do we mean by talking about *substitution in states of affairs*? This may sound obscure at first, but the underlying idea is in fact very simple. Suppose someone asserts that 'Adam loves Becca' and someone else replies 'No, it's Charles rather than Adam'. What the second speaker means is that it is not the state of affairs in which Adam loves Becca which obtains, but rather a similar one, in which Charles takes the place of Adam (that is, the state of affairs in which Charles loves Becca). We might also want to say that the latter state of affairs is the result of *substituting* the individual Charles for the individual Adam in the state of affairs in which Adam loves Becca. Similarly we can regard the state of affairs in which Charles admires Becca as the result of substituting the relation of admiring for the relation of loving in the state of affairs in which Charles loves Becca.

It is now important to note that the ability of a collection of objects to go together in forming a state of affairs is not stable under substitutions. Consider Lego blocks. For some collections of Lego blocks you can build some single Lego structure which uses all the blocks. But if you exchange some blocks for others, this is not necessarily the case any longer. It may be the case that either you cannot build any structure at all, since the new collection of blocks does not fit, or you might be able to build something, but be unable to use all the blocks. The very same thing can happen with states of affairs. Take the three objects Adam, Becca and the loving relation. They can form a state of affairs together. (In fact they can form two: that Adam loves Becca and that Becca loves Adam). Now substitute the loving relation by the relation 'have 3 as their greatest common factor'. There is no state of affairs which consists just of Adam, Becca, and *this* relation, now matter which order we arrange the elements in. The relation is of the wrong type: it is a relation between numbers, whereas loving is one between persons. Similarly, consider substituting the loving relation by the property 'is male'. There is no state of affairs consisting just of Adam, Becca and 'is male' either, since 'is male' takes only one individual, whereas 'loves' takes

¹⁴Smart (1954).

two. Parts of the collection of Adam, Becca and ‘is male’ can go together to form a state of affairs, but the whole collection cannot.

We should stress at this point that the idiom of constituents of states of affairs fitting together is not just a *façon de parler* for words fitting together to form meaningful sentences. We could not just transpose the discussion of what states of affairs could obtain to what sentences make sense. States of affairs and meaningful sentences are closely related but differ in important ways. It is not the case that if some objects $a, \dots, n$ fit together to form a state of affairs there is a meaningful sentence containing expressions referring to the $a, \dots, n$: our language might not contain names for all the constituents. It is not even the case our language always *could* contain such names. Depending on the extent of one’s Platonist inclinations one might want to claim that all relations between the real numbers (or some other set of the same cardinality) constitute states of affairs. But names being finite strings of symbols, there are not enough of them to go around for every real number. Nor is it the case that if a sentence picks out a state of affairs with constituents $a, \dots, n$ there are expressions referring to $a, \dots, n$ in the sentence (none of the parts of an equation a Gödel number picks out correspond to the digits in the Gödel number). States of affairs and sentences stand in a close structural relation, but not in one which is so close as to make them identical for our purposes.

This does not mean, however, that we cannot use information about sentences in order to gain knowledge about whether certain objects referred to by subsentential components can be joined into states of affairs. *Grammaticality* (or rather the lack of it) can be some indication: that the relations of loving and sitting between belong to different types can be inferred from the fact that ‘sits between’ cannot be plugged in for ‘loves’ in ‘Adam loves Becca’ *salva congruitate*. *Meaninglessness* is another one: we can argue that being prime and being green belong to different types since ‘green’ cannot be intersubstituted with ‘prime’ in ‘17 is prime’ *salva significatione*. Note that the meaningless sentences we consider for these purposes are a very restricted class, namely just those sentences which are *grammatical* and where *the meanings of all the constituent expressions are clear*, but which still lack meaning as a whole.

Taking the difference between states of affairs and meaningful sentences into account we realize that some recourse to linguistic information does not imply that our exposition of the notion of an ontological category is fundamentally linguistic, in particular since we do not rely *only* on linguistic evidence.

Our ability to *imagine* possible but not actual states of affairs is a further source of information about the joining behaviour of their constituents. None of us have ever encountered states of affairs featuring pink elephants, talking donkeys or worlds with four spatial dimensions. But we find it quite possible (with a bit of practice in the last case) to imagine what such states

of affairs would be like. In other cases, however, we fail no matter how hard we try: even considerable practice does not supply us with a sufficiently clear impression of what a state of affairs containing a green prime number or the square root of a symphony would be like. In cases like these we do not get behind the mere words to ascertain the referents of the respective sentences.

There are thus at least three ways of finding out about the joining behaviour of constituents of states of affairs: by considering grammar, by considering meaninglessness and by considering our imaginative capacities. The reader will have noticed that I have not made any statement about what underlies the fact that some constituents of states of affairs can go together while others cannot. I have described ways in which we can *acquire knowledge* about it, but I have been silent about what is actually *behind* the whole thing. What makes it the case that only certain bits of states of affairs fit while others do not?

One possible answer could appeal to *essential properties* of objects. Numbers, it could be argued, are essentially non-spatial, and since everything which has a colour must be spatially extended the property ‘being green’ cannot apply to any number in a state of affairs. Which object goes together with which other objects in states of affairs then would be taken to be a direct consequence of the different objects’ inner nature. But note that this is not the only possibility of explaining joining behaviour.

Another explanation could appeal to *neurobiology*. The assumption would be that our brains are set up in such a way that certain concepts just cannot be combined. Our inability to imagine certain states of affairs would then not have anything to do with the essential properties of the objects involved, nor would it be a reflection of lack of imagination on our part. It would be the case that it is biologically impossible for us to make a connection between certain mental representations. This could be explained by the fact that particular representations necessitate certain codings in the brain, and that the neuronal structures brought about by these codings cannot be merged, due to the way our brain is constructed. The impossibility of imagining a particular state of affairs would be the same kind of impossibility as that of learning a language which violated certain language universals: an impossibility resulting from the hard-wiring of our brains.¹⁵

Finally we could try to explain the different joining behaviour on purely *linguistic* grounds. Once we have expressed the ‘deep structure’ or ‘logical form’ in a suitable way we see that the sentence ‘The number seven is green’ is just as ungrammatical on the ‘deep’ level as ‘Caesar is and’ is ungram-

¹⁵I do not claim that there is actually any neurobiological evidence for such a foundation of inconceivability in brain structure. The sole purpose of this hypothesis is to show that appeal to essences is not the only possible way of explaining information about joining behaviour acquired via considerations of imaginability.

mathematical on the surface level.¹⁶ That only certain bits of states of affairs fit together would then have nothing to do with either the essences of objects or with the way our brains work, but would be a result of facts about the ‘deep grammar’ of the language employed for speaking about the world.¹⁷

Now which of the three explanations, if any, is the right one? This is a deep and difficult question to which I do not pretend to know the answer. For this reason I set up my theory in such a way that answering it is not necessary for arriving at an account of ontological categories. I treat the ‘fitting relation’ between constituents of states of affairs as primitive and develop a theory on this basis. Given the fundamental epistemic and semantic rôle of states of affairs¹⁸ this relation seems to be sufficiently well-entrenched in our cognitive lives to make it an interesting object of study. Since I was able to describe some possible ways in which knowledge about this relation could be acquired its epistemological background should be sufficiently clear for further investigations. But its ontological background is far from clear. It could be a manifestation of facts about essence, or brains, or deep structure, or something else completely. It is therefore so important for my account that it relies on information about this primitive ‘fitting relation’ only.

It might be useful to compare our approach to some other ‘structuralist’ theories in philosophy, which are similarly restricted to a single primitive relation. Carnap bases the ‘logical structure of the world’ on a primitive relation of ‘recollection of similarity’ between elementary experiences and sets out to construct all other phenomena, among them individuals and properties, the space-time world, physical and psychological objects from this.¹⁹ In a similar vein theories of ‘particularized individuals’ or tropes aim to construct all sorts of metaphysical entities based solely on a resemblance relation between tropes.²⁰ Now none of these theories sets out to deny that there is something else *in the objects* (the elementary experiences or the tropes) which is responsible for the similarity relation holding between them. The point is rather that all our epistemological access has to proceed via the similarity relation — what is behind it is either epistemologically

¹⁶Such a view of formal languages was famously held by Carnap (1959, 68).

¹⁷Note that whichever option we go for, we have a way of explaining what our judgement that the number seven and the colour green cannot go together in as state of affairs is about that does not appeal to the ontologically rather problematic notion of an impossible state of affairs. In the first case the objects of our judgements are the respective essences of the number seven and the colour green, in the second case we talk about the structure of neuronal representations and how they relate to the structure of our brain and in the final case we refer to consequences of the ‘deep’ grammatical rules of the language we speak.

¹⁸Westerhoff, 98-122.

¹⁹Carnap (1928).

²⁰For an account of trope theory which treats tropes in much the same way in which I treat states of affairs see Bacon (1995).

inaccessible or not sufficiently clear to build a theory around it. Carnap can agree that something makes the elementary experiences resemble one another, but what this is is not clear to us. The tropist does not have to deny that there might be something which makes all the red tropes resemble one another, but he claims that the only way of developing an account of this has to be based solely on the similarity relation between tropes. Restricting ourselves to the primitive relation is therefore as good a starting point for our theorizing as we can get, and the aim is to find out how far we can make it from there.

The same motivation is behind our restriction to a ‘fitting relation’ between constituents of states of affairs in order to develop an account of ontological categories. There might be all kinds of things responsible for such fitting (facts about essences, about brains or about deep structure), but it turns out that not of these are sufficiently clear for developing a satisfactory theory of ontological categories.²¹ We will therefore bracket further other considerations as far as possible and proceed by developing a theory of ontological categories based on a single primitive relation.

4 Generating types from substitution-patterns

We have now seen that certain collections of objects fit together to form states of affairs, while others do not. How does this help in sorting objects into types?

Let us say that two objects belong to the same type if there is a state of affairs containing the first as a constituent and the result of substituting the second for it is also a state of affairs. So Adam and Becca will belong to the same type, since substituting Becca for Adam in the state of affairs that Adam loves Becca gives us the one in which Becca loves Becca, which is a state of affairs as well. The number 17, however, will not belong to this type since it does not ‘fit into’ the place of Adam: that 17 loves Becca is not a state of affairs.

We will call this the *weak* conception of types. It is weak because it only demands intersubstitutability in *at least one* state of affairs to ensure that two constituents belong to the same type. The strong conception, on the other hand, demands intersubstitutability in *all* states of affairs.

The distinction between these two kinds of types is familiar from grammatical accounts of defining categories of expressions by substitution. A simple way of doing this assigns two expressions to the same grammatical category if all sentences in which the one occurs remain sentences if we exchange it for the other.

²¹For some criticism of attempts of building such a theory around the notion of essence or deep structure see Westerhoff, chapter 2.

Unfortunately this criterion does not always give us what we want. Words which intuitively belong to the same grammatical category, such as ‘Boolean algebra’ and ‘Wednesday’, which are both common nouns, can be shown to belong to different categories by suitable sentence-frames. ‘I visited my aunt last Wednesday’ is grammatical, while we would be hard pressed to say the same of ‘I visited my aunt last Boolean algebra’. Of course this is just a problem for strong types. So what about picking weak types instead? Obviously a single context cannot do, else ‘the violin’, ‘while it is raining outside’ and ‘because he is bored’ would all turn out to belong to the same grammatical category since they can all be plugged in for the blank in the sentence-frame ‘Peter plays ...’ to produce grammatical sentences. Thus we need more than one context. Unfortunately it is very difficult to give a precise answer as to how many one needs if intersubstitutability is to work as a criterion for picking out grammatical categories. Grammarians generally avoid this problem by adding further criteria (such as morphological or semantic considerations) to supplement the account of grammatical categories in terms of intersubstitutability.²²

Weak and strong types have different structural properties. Strong types constitute a partition on the set of constituents of states of affairs: every constituent belongs exactly to one type. Weak types will generally exemplify a containment structure, so that constituents can belong to more than one type. Indeed on the assumption that there are parts of states of affairs corresponding to what are sometimes called ‘high predicates’, such as ‘is self-identical’, ‘can be thought about’, ‘may be referred to by a meaningful sentence’, *everything* will be intersubstitutable in states of affairs containing these. Substituting any other object (including the number 17) for Becca in the state of affairs that Becca is self-identical is also a state of affairs. Thus there will be types everything belongs to. On the weak conception of types each constituent of a state of affairs is typically a member of several very inclusive types (such as the type of all self-identical objects), as well as of other, more specific types included in these (such as the type of persons). On the strong conception of types, on the other hand, there will only be the more specialized types, since types cannot contain one another. Because we want to develop a conception of ontological categories which satisfies the structure condition it is natural to concentrate on weak types. They are already equipped with a ready-made containment structure whereas in the case of strong types such a structure would have to be subsequently introduced.

But we cannot without further ado equate ontological categories with weak types without violating the distinctness condition. This is because some of the more specific types will be too specific to qualify as ontological categories. Remember that for some objects to constitute a weak type it is

²²See e.g. Bradford (1988, 63).

sufficient that there is *one* state of affairs such that they are all intersubstitutable in it. Therefore a single state of affairs with a particular property as a constituent which only takes objects from a very restricted set will make this set into a weak type.

Take an example. Consider the state of affairs denoted by the statement ‘Peter was gated’.²³ We cannot simply intersubstitute any other person for Peter in this state of affairs since the property ‘to gate’ only takes Oxbridge undergraduates as arguments. Just a non-person could not love anything, a Non-Oxbridge undergraduate could not be gated. Being gateable presupposes being an Oxbridge undergraduate, so it is inconceivable that someone lacks the latter property while still having the former. But if we identify ontological categories with weak types, it will turn out that since any Oxbridge undergraduate (and nothing else) could take the place of Peter in the state of affairs that Peter was gated, there is a weak type consisting of all and only the Oxbridge undergraduates. But we will hardly want to say that these form an ontological category!

A closely related point is already present in Smart’s criticism of Ryle who argues that on the latter’s proposal even tables and chairs would constitute ontological categories, since only names of chairs could be meaningfully plugged in for the blank in such sentences as ‘... has a hard seat’.²⁴ In fact this is quite a widespread phenomenon. The property of having a green back door is only applicable to buildings, the property of having more than 200 pages only to books, and that of being abelian only to mathematical groups or categories. So all of these would have to be considered as ontological categories as well. But this is clearly not in harmony with the condition of distinctness described above. What we need therefore is a procedure for identifying all the too specific weak types of constituents of states of affairs in order to retain only those which are sufficiently unspecific to qualify as ontological categories.

Thinking about what distinguishes the more specific from the less specific types suggests an answer. It becomes apparent that all the more specific types are only there because the less specific ones are. There is a sense in which types containing Oxbridge undergraduates, or buildings, or groups are parasitic upon other, less specific types, namely those containing persons, or material objects, or sets. They rely on these other types for their existence. These very specific sets are also in a certain sense redundant: if all the Oxbridge undergraduates, all buildings and all groups suddenly vanished overnight, but persons, material objects and sets were still around, the loss would not be considered fundamental. It would not have affected the ontological richness of the world. We could still regain the lost categories

²³The verb ‘to gate’ means, according to the OED ‘to confine an undergraduate at the Universities of Oxford or Cambridge to the precincts of the College, either entirely or after a certain hour’. See also Stubbings (1991, 26).

²⁴Smart (1954).

from the ones which are still there.

The fundamental idea I want to use in this exposition of the notion of an ontological category is that systems of weak types can be *systematized* by constructing some of the types from other types. I will then argue that only the types used in these constructions, types which constitute the constructional basis of the set of types, should be considered as fundamental and therefore as ontological categories.

It is this consideration of constructing types from other types which distinguishes my account of ontological categories from those of Ryle and Sommers. First of all constituents of states of affairs are sorted into types — this is still very close to the Ryle-Sommers account in terms of substitution criteria. However, as we saw above, this is only the basis for a theory of ontological categories, but cannot constitute the whole account. Types like those discussed by Ryle and Sommers include some which are far too specific to constitute ontological categories. Therefore we have to find some way of ‘filtering out’ these too specific ones in order to arrive at a satisfactory definition of the set of ontological categories. This is done by appealing to the notion of construction of types, something which is not part of the account of types given by Ryle and Sommers.

But what exactly is the construction of types from types? Let us now turn to the description of this notion in greater detail.

5 Constructing types

5.1 Two kinds of construction

First of all it is important to note that there are two fundamentally different kinds of constructions which I will call *replacement* construction and *surrogate* construction. The two can easily be distinguished by considering an example from chemistry. Suppose we have three chemical substances, Natrium, Chlorine and Natriumchloride and get rid of the Natriumchloride. Afterwards we produce some of the Natriumchloride from the Natrium and the Chlorine. The substance we got rid of initially has been *replaced* and the collection we end up with contains the same chemical substances (although in different amounts) as the original collection.

Now consider the different case of having a collection of three substances A, B and C. Again we get rid of C. This time, however, we use A and B to produce some substance D which has a different chemical constitution from C but resembles it in several important respects. D thus acts as a *surrogate* for C.

Replacement constructions will be familiar from logic. For example if we want to systematize some set of logical truths we get rid of all but some of them (‘the axioms’) and use some constructional operations (‘the rules of inference’) to reconstruct them from the axioms. Similarly in giving a

recursive set of syntactic rules for constructing well-formed formulae of our logical language we show how the endless variety of logically complex formulae can be constructed in a step-by-step manner from the atomic formulae together with the specific rules of formation.

It might be argued (at least if we have any structuralist leanings) that surrogate constructions are frequent in the foundations of mathematics, where real numbers are constructed as sets of rationals, natural numbers as pure sets, points as sets of volumes or volumes as sets of sets of sets of points. Regarding these we would want to say that neither numbers nor volumes nor points *are* sets, but that the constructs in question behave in all important (i.e. mathematical) respects just like the originals they are intended to replace.

The relation between construct and original is not completely straightforward. It might for example be initially convincing to demand an extensional identity of the expressions referring to the construct and the original. Something would be a good construct of an original if what is true of the original is also true of the construct, and vice versa. This was advocated by Carnap, who argued that the fact that some *a* can be constructed from *b* and *c* meant that each expression referring to *a* can be replaced by one referring to some construct from *b* and *c* *salva veritate*.²⁵ For example, Carnap argues that whenever we refer to prime numbers we can also refer to natural numbers having the property of being divisible only by 1 and by themselves. But in fact demanding substitutability *salva veritate* is far too restrictive. As Goodman has shown, this is obvious in the case of identity-statements involving constructs and originals.²⁶ To know whether replacing ‘point’ by ‘set of volumes’ in ‘a point is the same as a set of volumes’ changes the truth-value of the statement we first have to know its truth-value. Once we do know it, however, we no longer need to appeal to the notion of extensional identity to see whether they really are the same. Thus the ‘test’ presupposes what it is supposed to show. It is furthermore important to note that this problem will not just arise in the case of identity-statements but also in the case of all statements depending on such identity-statements. To determine whether replacing ‘point’ by ‘set of volumes’ in ‘a point has no members’ changes the truth-value of the statement, for example, we first have to know whether points really are the same as sets of volumes.

The failure of extensional identity of construct and original is in harmony with the fact that we often give different extensionally non-identical constructions of the same original, as in constructing a point as a set of volumes or as a pair of intersecting lines. We can thus agree that far from being identical with the originals the constructs are in many respects quite differ-

²⁵Carnap (1928, 47).

²⁶Goodman (1951, 8). See also Quine (1976).

ent.²⁷ The point of devising a construction ‘consists not in showing that a given entity is identical with a complex of other entities but in showing that no commitment to the contrary is necessary’.²⁸ The important point in determining whether something is a construction of an original is therefore, as Goodman notes, that the truth-value of the sentences we ‘care about’ is preserved.²⁹ The existence of extensionally non-identical constructions of the same objects is evidence enough that there are always some sentences we do not care about. In the above case, for example, we do not care about whether points have members. Similarly in the case of set-theoretical constructions of the natural-number sequence we do not care about whether the number 2 is a member of the number 4 (which is the case with von Neumann’s but not with Zermelo’s construction).

5.2 Construction in ontology

Let us now consider a simple example of a surrogate construction of a weak type from other weak types. Take the case of events. Events are identified as a weak type by noting that certain parts of states of affairs such as ‘the battle’, ‘the funeral’ and so on can form states of affairs with other parts of states of affairs such as ‘was interrupted twice’ or ‘lasted for two days’ and that this fact distinguishes them from other parts of states of affairs. These parts constitute a type, and intuitively we will want to say that this is the set of events. The aim is now to construct this type from other types.

Jaegwon Kim suggested a way in which this could be done.³⁰ He equates events with ordered triples of individuals, properties and time-instances. The individual is the bearer of the property which is exemplified by that individual when the event occurs. Thus the event of me lighting a candle is spelt out as the triple consisting of the candle, the property of its being lit and the particular time when this happened. Kim uses this account amongst other things as a basis for a theory of criteria of identity for events and for analyzing their rôle in causal contexts. We are here interested not so much in whether this account is completely satisfactory, but rather in the underlying thesis that if it is and if we have types containing individuals, properties, time-instances and the set-theoretic membership relation, the type of events can be dispensed with. The idea is therefore that the existence of events does not contribute anything genuinely new to the kinds of things there are: events are just a special configuration of other kinds. Therefore the type of events will can be constructed out of other types and will therefore not be regarded as an ontological category.

Similarly we might want to say that rational numbers behave quite dif-

²⁷See Gottlieb (1976, 59).

²⁸Goodman (1951, 29).

²⁹Goodman (1951, 23).

³⁰Kim (1976, 161).

ferently from integers (after all, this is why they have been introduced) but there is a sense in which this difference of behaviour which makes them into a different kind of number is not fundamental. It does not mark a diversity which could not be achieved in a world in which there are only integers and the pairing function. The different behaviour shown by the rationals is the very same thing as the different behaviour shown by equivalence classes of pairs of integers. On the basis of this reduction we will therefore want to say that the rational numbers do not constitute a genuinely new kind of entity in a world with integers and the pairing function.

Furthermore, note that e.g. in claiming that events do not qualify as ontological categories we do not imply that some particular events (such as the battle of Leuthen or Victor Hugo's funeral) do not belong to any ontological category, since the category of events is not granted such a privileged title. Everything belongs to some ontological category. If we follow Kim's line, these events are all particular kinds of sets (i.e. ordered triples) and the category of sets is the ontological category they belong to. This holds for members of constructed types in general: the ontological category they belong to is one of the types it has been constructed from.

My theory is not just committed to the constructability of relatively general types, such as events, but also to that of more particular ones, such as tables, chairs or buildings. Given their greater specificity the constructional processes for these are much more intricate than those required in the case of events. Surrogate constructions of tables, chairs and buildings cannot be just regarded as complexes of physical objects, but will presumably also have to include some psychological items, such as the human *intentions* that the lump of physical stuff is to serve as something to sit on or to live in. This, however, is not a qualitative, but only a quantitative difference. In the same way in which the construction of imaginary numbers out of pure sets is more complicated than that of the integers, that of more specific categories is more complicated than that of less specific ones. There is, however, no difference in the *nature* of the different constructional processes involved.

5.3 When is a construct adequate?

Considering the *original* event and the *construct* ordered triple which is supposed to be a surrogate for it we realize that their features properly overlap. There are some features which the original has but which the construct lacks (events start and stop, sets are timeless), some features of the construct are not had by the original (ordered sets have members, events do not), and some are shared by both (both the set and the event involve properties and individuals). The features they share will determine whether a construction of an original is adequate or not. To give a precise specification of the amount or nature of the features the two must share to make something an

adequate construct is surprisingly hard in the general case.³¹ The extent to which one can give precise criteria for the adequacy of a construct depends crucially on how precise our understanding of the original is. In the case of arithmetic, for example, nothing is an adequate construction of the natural numbers which does not satisfy the Peano axioms. A construct containing only finitely many items, or one where two items could have the same ‘successor’, can under no circumstance count as an adequate substitute for the natural numbers. The Peano axioms thus provide us with a necessary condition for the adequacy of constructs in the case of natural numbers. Unfortunately, such necessary conditions cannot always be formulated with this degree of precision. However, we might get something less precise but still useful. It seems evident that any construct which is supposed to act as a substitute for events must somehow incorporate the fact that they exist in time.³² Similarly anything which can go proxy for buildings must fulfil certain minimal conditions. The notion of an *essential feature*, of something an object cannot lose without ceasing to be that very object, suggests itself here. The natural number structure cannot lose the features described in the Peano axioms, events cannot stop being in time without ceasing to be events, and so on. We might therefore say that something is an adequate construct of some original if it has all the essential features of the original. This criterion is of course only as clear as our conception of the essential features of the original under consideration. If these are relatively well-defined we seem to get round problems such as the ‘unicorn-construction’ discussed by Gottlieb.³³ The idea there is that we have a set-theoretic construction of the natural numbers which is completely standard apart from the fact that the successor-function $S'xy$ is defined as $(x = y \wedge (\exists x)(Ux)) \vee Sxy$, where U stands for the property ‘being a unicorn’. This construct is extensionally equivalent to the standard one: all sentences true in it are also true in the unicorn-construct. The problematic thing about the unicorn-construct is that it *entails* things the standard construct did not entail. For example the truth that no number is its own successor $((\forall x)\neg(x = x \wedge (\exists x)(Ux)) \vee Sxx)$ entails that there are no unicorns $(\neg(\exists x)(Ux))$. It is clear that we must find some way of ruling out this kind of construction. If the natural number structure has any essential features at all, not entailing anything about unicorns is one of them. The unicorn-construct therefore cannot be an adequate construct of the original natural number structure.

³¹See Gottlieb (1976, 67–69) for some attempts.

³²It might be argued that this shows Kim’s construction does not work: sets, being abstract, cannot be in time and thus cannot act as surrogates for events. Two ways of repairing this immediately suggest themselves. We might use a different composition operation instead of set formation which does not generate abstract objects (such as mereological fusion), or we could opt for an interpretation of set theory which supplies sets with spatio-temporal locations (for example along the lines of Maddy (1980)).

³³(1976, 64).

6 The resulting picture

Now the idea of construction has been clarified we realize how the problem of the too specific weak types is to be avoided. We select a proper subset of the set of weak types and use this to produce surrogate constructions of the remaining objects in the set. These objects will possess all the essential features of the weak types they are supposed to replace. We will then consider only those weak types which can function in this way as a constructional basis to be ontological categories.

How do we know that the weak types we end up with are really less specific than the ones we have discarded? This is evident from what it means to construct one object from another one. For an operation to be a construction operation it must be complexity *increasing*, rather than decreasing. Roughly speaking, an operation is complexity increasing if it makes more complex objects out of less complex objects, rather than the other way round. Obvious examples of complexity increasing operations are encountered when constructing sets from their members, constructing the concept ‘bachelor’ from ‘male’ and ‘unmarried’, or constructing molecules from atoms. Decreasing operations work the other way round: they produce members from the sets containing them, ‘male’ and ‘unmarried’ from ‘bachelor’ and atoms from molecules.

The difference between complexity increasing and complexity decreasing operations can be made more precise by considering the complexity of algorithms or instructions for generating objects.³⁴ It is straightforward to calculate the complexity of these algorithms in terms of the number of computational steps and the memory capacities needed. For each object call the least complex algorithm generating it its *recipe*. Now an operation is complexity increasing if the recipe of its input is always less complex than the recipe of its output, and complexity decreasing if it is the other way round. For an example consider the following two simple operations O_1 and O_2 which take pure sets as inputs. O_1 returns for any set ϕ the output $\{\phi\}$, O_2 makes ϕ out of inputs of the form $\{\phi\}$. Now clearly if the complexity of the recipe of some set is n , that of the recipe of the result of applying O_1 to it will be greater than n and that of the recipe of the result of applying O_2 to it will be smaller than n . Operations which use O_1 are therefore complexity increasing, those which confine themselves to O_2 complexity decreasing.

So my account demands that operations constructing weak types out of weak types are complexity increasing. What is the argument for this? My main concern here is a desire for qualitative economy regarding composition operations. We can achieve the same compositional results by using complexity increasing operations we get by using increasing and decreasing ones, but not by using decreasing operations alone. Consider an exam-

³⁴See e.g. Chaitin (1987a), (1987b).

ple from chemistry. Suppose we have some collection of atoms and some molecules consisting of atoms of this kind. If we want to select some subset of the collection and construct all other members of the collection from the subset employing some compositional operation there are two ways of doing this. We can choose a complexity increasing operation which synthesizes all the molecules from the set of atoms. Or we can take a set of molecules (which contains a sufficiently large variety of atoms as constituents), break them up into their constituent atoms and then assemble all the remaining molecules from these. In this second case we employ first a decreasing and then an increasing operation. Just using a decreasing operation would not be sufficient in the general case. Although it is sometimes possible to construct one molecule from another one by breaking up the latter (for example by separating the hydroxyl group from an alcohol) this usually does not give us enough variety for constructing all the molecules in the remaining subset. In order to this we have to go via the atoms and then reassemble from these.

Now the situation with categories is exactly analogous. Though it may sometimes be possible to construct a simpler category from a more complex one, complexity decreasing composition operations are not sufficient for allowing us to come up with a ‘constructional systematization’ (i.e. a set of ontological categories) for a set of categories in general. Given that we want to restrict the types of composition operations employed we select the complexity increasing ones, which *can* achieve such a systematization. But this implies that the types used in the construction are general and unspecific, while the constructed ones are less general and more specific.

I hope to have convinced the reader by now that our account of ontological categories which fuses an understanding of them in terms of substitution classes with the idea that they can be systematized in terms of construction is a satisfactory explication in that it obeys the conditions of distinctness and structure mentioned at the beginning. The substitutional conception ensures that the ontological categories can stand in containment relations while introducing the idea of construction allows us to draw a line between those types which are ontological categories and those which are not. I would like to conclude by describing two philosophical implications of this conception.

6.1 Relativism

If we consider a set of types, there are several ways in which its members may be constructionally related. It may be the case that no subset is able to construct any other one or that precisely one subset may be able to construct all the other types. The situation is a bit more complicated if there are several sets which can construct all other sets. There may be some types which are thus not constructible from anything. We will call the set of these the *core of the basis*. Equally there might be types, which

are always constructed, never used to construct other types. We call these *redundant types*. By means of illustration we might want to say that types like the type of sets or mereological sums, which turn up in a number of ontological constructions belong to the core of the basis for most sets of types. Redundant types, such as culinary implements or kinds of furniture, are not usually employed to construct any other types, but can in turn very well be constructed.

If we want to decide which types are used as a basis of construction, and which are constructed the cores of the basis and the redundant types are obviously to be included in the two.³⁵ Regarding any other type there is room for negotiation. Depending on which basis is taken as most plausible, some type which is neither in the core of the basis nor in the set of redundant types may come to lie on either side of the divide between the constructing and the constructed types. Relative to a given world (a given set of types) there may be more than one set of types which can be chosen as a basis, and thus more than one plausible candidate for the set of ontological categories in that world.

But how do we settle which of the different bases of the set of types under consideration is to be regarded as the set of ontological categories? There are certainly cases where this is a matter of argument. Considering the types of individuals and properties it is evident that there have been attempts to construct the first from the second and the second from the first. Constructing properties from individuals is usually associated with different forms of nominalism, while constructing individuals from properties (sometimes called ‘universalism’) is what bundle theory sets out to perform.³⁶ For the nominalist, individuals would be the constructing and properties the constructed type. For the universalist, it would be exactly the other way round. Similarly, the firm believer in events may well argue that in a world in which there are sets, individuals, properties, time-instances and events, it is not the events which should be constructed from other types. Rather should sets and events be taken as basic, while individuals, properties and time-instances could be conceived of as set-theoretic constructions from events.

If both accounts worked, and worked equally well, this would indeed leave us in a tie regarding which of the types in question are supposed to be basic. But in ontology the situation is usually not like this. Apart from the fact that it is relatively hard to provide a satisfactory construction even in *one* direction, the two constructions would presumably differ enough in their internal details to provide some criterion for accepting the one but not the other.

³⁵If they exist, that is. There may be radically different bases which can construct an entire set of types so that there is no core of the basis, or it may be the case that every type can be employed to construct some others, so that there are no redundant types.

³⁶See Armstrong (1978, I, §§ 2–9) for an overview.

It is apparent that there is no determinate procedure for selecting the ‘best’ of the bases. We rather employ pragmatic criteria, such as the size of the basis, the nature and simplicity of constructions and so forth in order to rate some bases as more plausible than others and then pick one of the most plausible ones. It should be noted, however, that the fact that it may in this way be relative whether some particular type is used to construct others and is thus an ontological category is no deficiency of the metaontological account described here. What we were looking for was a satisfactory account of ontological categories; whether a particular class of objects fulfils that definition is then a further question, and it is ontology rather than metaontology which has to answer it.

It is important to notice that this relativism does not extend to the types themselves. What types there are is settled by the world, rather than by our decision to organize the world in a certain way, simply due to the fact that objects go together to form states of affairs in a certain way. But a list of types is not yet an ontological theory. A list of types contains all sorts of kinds of things, even those which are far too special to fall within the precinct of ontology, which is after all supposed to be a theory of the most general kinds of things there are in the world. Relativism comes into play when we begin to *systematize* the list of types under some constructional construction. This kind of relativism is benign; while there is a problem with relativism in the case of ways of being, for they must be one way or another, there is none in the case of ways of systematizing: there is no necessity that there should be a unique or even a best way amongst the most fundamental ways of systematizing information about the world.

The comparison of a system of ontological categories with the axiomatization of a theory might be helpful at this point. While the answer to the question whether proposition A or $\neg A$ is a *truth* of Euclidean geometry does not depend on us but is settled by something else (although it is surprisingly difficult to specify exactly what this ‘something else’ is), the answer to the question whether proposition A rather than B is an *axiom* of Euclidean geometry is not settled by this something else, whatever it may be. Assuming that A and B are both true, it is settled by us and by our desire to systematize the truths of Euclidean geometry in a certain way.

It therefore turns out that we have to say goodbye to the conception of ontological categories as a unique and objective fundamental set of objects which encompasses the most general kind of things there are. In the same way in which contemporary mathematics does not regard axioms as self-evident truths any more, but as truths which play a certain rôle in a particular systematization of some body of knowledge, contemporary ontology should regard ontological categories as kinds of things which play a certain rôle in our systematization of the phenomenological plenitude of the world.

6.2 Holism

Our second conclusion is a holist conception of ontological categories. Whether a set of objects constitutes an ontological category, and whether some object belongs to a particular type does not depend on the nature of the object or objects considered, but on what *other* objects there are.

Let us consider the first point first, that whether a set of objects constitutes an ontological category depends on what other objects there are. The relativist view of ontological categories entails that certain types could come out as ontological categories or not, depending on the constructional systematization. But remember that above we mentioned the possibility that some types might belong to the intersection of all bases, and could thus not be constructed from any other type (we called the set of these types the core of the basis). We might now be tempted to assume that these types, which come out as fundamental in *any* systematization of the set of types, could be regarded as being essentially ontological categories. But in fact this is not the case. Whether a type belongs to the core of the basis depends on what *other* types there are. Relative to one world (one set of types) a particular type might be in the core, but relative to another world it might not be. But this of course means that it cannot be due to the nature or essence of the objects in the type that it belongs to the core. We therefore see a holist picture emerging. The place of a set of objects in the set of types is determined by the objects themselves together with the other types there are, since these are responsible for the joining behaviour the objects show, which in turn determines what type they belong to. Whether some set of objects constitutes an ontological category is thus fixed by the whole world, rather than by individual objects in isolation.

The above holism also implies that whether some object belongs to a particular type depends on what other objects are around. Here it is helpful to make a comparison with *semantic* holism. This claims that we cannot look into a word to see what it means, but that we must look at its relations with other words. We have to adopt a behaviourist approach towards meaning: meaning is not something to be found deep down in the nature of a word, but something arising from the interrelationship between many different words.

Our *categorial* holism incorporates an exactly parallel view. We cannot tell the type something belongs to by merely looking at *it* — we have to see how it behaves relative to other objects in the formation of states of affairs. But of course how something relates to other objects depends on what other objects there are. Therefore the type an object belongs to can change if the collection of other objects present in the world changes. That two object have the same joining-behaviour (and thus belong to the same type) might just be a product of their present environment: relative to some other environment they might not be in the same type.

At this point we might be tempted to assume that the difference in

joining behaviour an object shows in different worlds (that is in the presence of different collections of objects it can form states of affairs with) is somehow determined by the object's inner nature or essence. As we saw in section 3.1, this would be one way of accounting for it, but not the only one. Since I do not know how to settle the matter I made the structuralist assumption of restricting myself to information about this joining behaviour only in order to set up our theory. We now realize that the conclusions arrived at in this way actually contradict an explanation of joining behaviour in terms of essences. Since the joining behaviour and thus a form set an object belongs to varies from world to world (depending on the other objects there are in the world), and since essences are supposed to be invariant across world-shifts, essences cannot be what is behind the joining behaviour of constituents of states of affairs. Note that a 'brain-based' explanation of joining behaviour would do better here. Its assumption is that the limits of conceivability are to some extent hard-wired into our brain, so that certain representations could not be combined. But since the existence of alien objects in a different world would presumably entail that the neuronal representation of these objects is different from any representations we have at present, facts about which representation can 'go together' will be different in the other world too. And this is just as it should be, given that the joining behaviour of objects is not assumed to be stable across worlds.

The anti-essentialist conclusion just arrived at of course contradicts the usual ontological position that it is a necessary part of the nature of an object to belong to a particular ontological category. Membership in the ontological categories an object belongs to, it is argued, is one of the properties an object cannot lose without ceasing to be *that* object. Individuals are essentially individuals, properties essentially properties, abstract objects essentially abstract objects, and without their category memberships they would cease being what they are. But on our account of ontological categories ontology systematizes information about how objects can go together to form states of affairs. Thus whether some set of objects is an ontological category and whether some object belongs to a particular type is fixed by the whole world. Both can change if new objects with new joining behaviour are introduced into the world: a type which is not constructible may become constructible, an object might now belong to a different type since its joining behaviour is changed due to the presence of the new objects. But none of this will mean that the nature or essence of the old objects is changed thereby.

We will therefore have to give up the idea that information about ontological categories supplies us with information about the essences of objects.³⁷ It provides us with a unified account of how objects in this world fit

³⁷Note that I do not deny that objects have essential properties. These even have a place in my account (see section 5.3). But I do deny that membership in an ontological

together into states of affairs. But since what things there are in the world is a contingent matter, claims about ontological categories cannot have the modal force attributed to them when it is claimed that they provide us with information about the essential properties of things.

Now I hear the friend of essences cry for the application of *modus tollens*: since basing our theory on a primitive fitting relation between constituents of states of affairs will not let category-membership come out as an essential property, he argues, we should dump a theory based on such a primitive. There is exactly one condition under which I would agree to go down that route: if the defender of essences is able to come up with a satisfactory theory of ontological categories based on the notion of essence as a primitive. As I have argued elsewhere,³⁸ all the attempts at doing this which can be found in the literature are fundamentally flawed and I think the feat cannot be done. Until I see an account which actually achieves this, I claim that the essentialist's justification for appealing to *modus tollens* at this point is not sufficient.

It therefore turns out that our theory of ontological categories which set out from the attempt to do justice to two of their central features, namely the fact that they are distinct from other categories and the fact that they form a hierarchy entails two surprising philosophical conclusions. The first is the relativistic view that what ontological categories there are is to a certain degree interest-dependent and a result of the 'conceptual scheme' we choose to pick, i.e. that it depends on the features we want to bring out in the constructive systematization. The second is the holist position that belonging to a particular ontological category is not a property which somehow flows from the essential nature of the object concerned, but is rather a place in the structure constituted by the joining behaviour of the objects there are and is therefore dependent on these other objects. We can now conclude that on a viable understanding of ontological categories which does justice to their central properties these categories turn out to be much more epistemological than ontologists would like to think.*

References

David M. Armstrong. *Nominalism & Realism*. Cambridge University Press, Cambridge, 1978.

John Bacon. *Universals and Property Instances. The Alphabet of Being*. Blackwell, Oxford, Cambridge MA, 1995.

category is such a property.

³⁸Westerhoff, chapter 2.

*I would like to thank two anonymous referees for the *Australasian Journal of Philosophy* for numerous apt critical remarks which helped to express the central points of the paper more clearly.

- Andrew Bradford. *Transformational Grammar. A First Course*. Cambridge University Press, Cambridge, 1988.
- Rudolf Carnap. *Der logische Aufbau der Welt*. Weltkreis-Verlag, Berlin, 1928.
- Rudolf Carnap. The elimination of metaphysics through the logical analysis of language. In A.J. Ayer, editor, *Logical Positivism*, pages 60–81. Free Press, Glencoe, 1959.
- Gregory Chaitin. *Algorithmic Information Theory*. Cambridge University Press, Cambridge, 1987a.
- Gregory Chaitin. Information-theoretic computational complexity. In *Information, Randomness & Incompleteness. Papers on Algorithmic Information Theory*, pages 23–32. World Scientific, Singapore, New Jersey, Hong Kong, 1987b.
- Roderick Chisholm. *A Realistic Theory of Categories*. Cambridge University Press, Cambridge, 1996.
- Michael Dummett. *Frege. Philosophy of Language*. Duckworth, London, 1981. 2nd edition.
- Nelson Goodman. *The Structure of Appearance*. Harvard University Press, Cambridge, MA, 2nd edition, 1951.
- Dale Gottlieb. Ontological reduction. *Journal of Philosophy*, 73(3):57–76, 1976.
- Reinhardt Grossmann. *The Existence of the World. An Introduction to Ontology*. Routledge, London, New York, 1992.
- Joshua Hoffman and Gary S. Rosenkrantz. *Substance among other categories*. Cambridge University Press, Cambridge, 1994.
- Joshua Hoffman and Gary S. Rosenkrantz. *Substance. Its nature and existence*. Routledge, London, New York, 1997.
- Jaegwon Kim. Events as property exemplifications. In Myles Brand and Douglas Walton, editors, *Action Theory*, pages 159–177. Reidel, Dordrecht, 1976.
- E. J. Lowe. *The Possibility of Metaphysics. Substance, Identity and Time*. Clarendon Press, Oxford, 2001.
- Penelope Maddy. Perception and mathematical intuition. *Philosophical Review*, 89(2):163–196, 1980.

- Van McGee. Logical operations. *Journal of Philosophical Logic*, 25:567–580, 1996.
- Bryan G. Norton. *Linguistic Frameworks and Ontology. A Re-examination of Carnap's Metaphilosophy*. Mouton, The Hague, New York, Paris, 1977.
- Alex Oliver. A few more remarks on logical form. *Proceedings of the Aristotelian Society*, 159:247–272, 1999.
- Willard Van Orman Quine. Ontological reduction and the world of numbers. In *The Ways of Paradox*, pages 212–220. Harvard University Press, Cambridge, MA, London, 1976.
- Gilbert Ryle. Categories. *Proceedings of the Aristotelian Society*, 38:189–206, 1938.
- Gilbert Ryle. *The Concept of Mind*. Hutchinson's University Library, London, 1949.
- J.J.C. Smart. A note on categories. *British Journal for the Philosophy of Science*, 4:227–228, 1954.
- Fred Sommers. Types and ontology. *Philosophical Review*, 72:327–363, 1963.
- Leslie Stevenson. A formal theory of sortal quantification. *Notre Dame Journal of Formal Logic*, 16(2):185–207, 1975.
- Frank Stubbings. *Bedders, Bulldogs and Bedells. A Cambridge ABC*. Cambridge University Press, Cambridge, 1991.
- Alfred Tarski. What are logical notions? *History and Philosophy of Logic*, 7:143–154, 1986. Edited by John Corcoran.
- Johan van Benthem. Logical constants across varying types. *Notre Dame Journal of Formal Logic*, 30(3):315–342, 1989.
- Jan Westerhoff. *Orders of things. An inquiry into the notion of an ontological category*. Forthcoming.
- David Wiggins. *Identity and Spatio-Temporal Continuity*. Blackwell, Oxford, 1971.