The Last Dogma of Type Confusions

Ofra Magidor

§1. Introduction

For a large part of the twentieth century there was a widely accepted view about a certain kind of sentences involving "type confusions". The kind of sentences in question were those which are standardly labelled as 'category mistakes', i.e. sentences such as 'The theory of relativity is eating breakfast', 'Colourless green ideas sleep furiously', or 'Wednesday is asleep', and the claim was that such sentences are meaningless.[1] The view certainly still has its defenders these days,[2] but it is no longer very popular and nowadays most philosophers seem to think that view is wrong: category mistakes are perfectly meaningful sentences, and the claim that they are meaningless is no more than an old-fashioned philosophical dogma.

Consider in contrast a different kind of "type confusion". Let us start with a perfectly grammatical sentence and replace an expression of one grammatical category (e.g. a name) with an expression of another category (e.g. a predicate). The result is an ungrammatical string which I shall label by 'grammatical type confusion' (GTC). For example, if we replace the name 'John' in 'John eats' with the predicate 'runs' the result is the GTC 'Runs eats'. And if we replace the predicate 'eats' in 'John eats' with the name 'Jack', the result is the GTC 'John Jack'.[3] Interestingly, in contrast to the case of

[1] For some defendces of the view see Russell (1967), Ryle (1949), Strawson (1952), Chomsky (1957), Routley (1966), and Drange (1966). Evidence for how prominent this view once was can be found in Prior's remark that an objector of the view "must nowadays count himself among the heretics". (Prior (1954), p. 160).

[2] See for example Diamond (2001), Sorensen (2001) p. 89, Beall & van Fraassen (2003), p. 125, and Fine (2003) pp. 207-208.

[3] I leave it somewhat underspecified how precisely the concept of a GTC is defined. I will merely say that a necessary condition for being a GTC is being a finite string of meaningful words that do not form a grammatical sentence. This leaves open the question of what are the sufficient conditions for being a GTC. (Are extremely ungrammatical strings such as 'has has the the the' GTCs? How about mild ungrammaticalities such as 'John drink beer'?). I will not address these questions here, but I will take as paradigmatic examples of GTCs strings such as 'Runs eats' which somewhat resemble grammatical sentences in their structure and where the source of their ungrammaticality has something to do with using expressions of the wrong broad syntactic type.

category mistakes, the canonical view concerning GTCs is that they are meaningless, and this view is rarely if ever contested.[4]

There is thus a notable contrast between the now prevailing attitude towards traditional category mistakes and between that taken towards GTCs.[5] This contrast is particularly striking since there are certainly similarities between what is seen to be, at least on a naïve take, the source of the infelicity involved in the two kinds of type confusions. Take for example the following definition of 'category mistake' given in *The Oxford Companion to Philosophy*: "The error of ascribing to something of one category a feature attributable only to another".[6] The thought is that what is wrong with a sentence such as 'Two is green' is that it involves an attempt to ascribe to something (the number two) of one category (numbers or abstract objects) a feature (the property of greenness) attributable only to entities of another category (concrete objects). But at least at a first pass it seems that this is exactly the problem that arises with GTCs: what is wrong with the string 'runs eats' is that it involves an attempt to ascribe to something (the property of running) of one category (first-order properties) a feature (the property of eating) attributable only to entities of another category (individuals).

We are thus faced with two kinds of type confusions which are at least superficially similar: category mistakes and GTCs. Assuming we accept that category mistakes are indeed meaningful (and I shall assume so for the remainder of the paper)[7], the question naturally arises: are there any good reasons for holding on to the view that GTCs are meaningless, or is this merely another incorrect dogma – the last prevailing dogma of type confusions?

---

[4] Chomsky (1965), p. 151 is a rare example of a defence of the view that at least some ungrammatical sentences are meaningful, but it is not clear that Chomsky allowed for extreme ungrammaticalities such as 'runs eats' to count as meaningful.

[5] For a classic example of the view that traditional category mistakes are meaningful and truth-valued but grammatical type confusions are meaningless, see Frege (1997b), p. 189.

[6] Honderich (1995), p. 126. I do not think that this is an ultimately satisfactory explanation of the kind of 'error' involved in category mistakes, but it does represent a standard naïve view on the matter.

[7] I defend this assumption in detail in Magidor (2007), ch. 3.

In this paper, I explore precisely this question. I discuss and evaluate a variety of putative reasons for accepting the view that GTCs are meaningless (I shall call this view - somewhat provocatively - 'The last dogma'). In **§2** I discuss metaphysical and semantic reasons, in **§3** I discuss logical reasons, and in **§4** I discuss syntactic reasons. I conclude that none of the reasons discussed is decisive, and thus that we should at least be willing to question this last dogma of type confusions.

Before I proceed, I would like to make one note that will hopefully help clarify the project I am engaging with. A simple-minded response to the question of whether GTCs such as 'runs eats' are meaningful might be the following: strings of words are only meaningful or meaningless relative to a given language. The string 'runs eats', the response goes, is trivially not a meaningful sentence of English simply because it is not a sentence of English at all. On the other hand, the thought continues, it is trivially true that 'runs eats' is meaningful relative to *some* possible language. For example, there is a possible language in which 'runs eats' means that Jane is reading. Moreover, even if we restrict our attention to languages in which 'runs' and 'eats' have the same meanings as in English, there is surely some such language where 'runs eats' is a meaningful grammatical sentence (for example, it might mean that someone is both running and eating). It thus seems that there is no interesting sense to be made of the question of whether 'runs eats' is meaningful.

I think this line of thought is too simple-minded. Consider by analogy the following case. Let English* be the language that consists of all and only English sentences which are up to thirty words long. Let *T* be the best semantic theory for English*. Plausibly, considerations of simplicity, elegance and so forth will dictate that the semantic theory *T* will not simply enumerate the finite but long list of sentences of English* and their respective meanings, but would rather be compositional in nature. Moreover, it seems plausible that other than the artificial restriction that *T* is only supposed to apply to sentences of English*, *T* would look very much like the best semantic theory for English. Now consider a sentence *s* of English that is thirty two words long. Does the semantic theory *T* assign it a meaning? In one strict sense the answer is perhaps 'no': *T* is only officially a theory about sentence of English*, and *s* isn't a sentence of English*. However, if we just choose to ignore this extra restriction there is a perfectly natural way to extend the meaning specifications recommended by *T* so as to apply to *s* and assign it a meaning. It is in this extended sense that it seems fair to say that *s* is after all meaningful according to semantics of English*.[8]

Analogous considerations apply to the case of GTCs. Even though a string such as 'runs eats' is not strictly speaking a sentence of English, the question is whether there is a natural way to apply the semantics of English so as to assign a meaning to such GTCs or is there some principled reason for why this cannot be done. I take the orthodox view expressed by the last dogma to claim that GTCs cannot be assigned meanings in this way, and it is this orthodox view that I will go on to question.

§2 Metaphysical reasons

The main reason for thinking that GTCs are meaningless has to do with an interaction between semantics and metaphysics. The thought is the following: terms of different syntactic types receive semantic values that belong to different metaphysical categories. Moreover, it is part of the metaphysical nature of entities of one such category that they can only be 'cohesively conjoined' with entities of a certain other such category. What is wrong with GTCs is that they involve an attempt to conjoin together two items of non-matching categories, i.e. two items that, as part of their metaphysical nature, cannot be conjoined.

A paradigmatic version of this idea is the Fregean metaphysical picture according to which properties, as opposed to objects, are 'unsaturated' or 'incomplete'.[9] The rough picture is that first-order one-place properties have a 'gap' or a 'hole' that can only be filled by an individual object but not by other properties; first-order two-place relations have two gaps; second-order one-place properties have one gap, but one that can only be filled by first-order properties; individuals have no gaps of any kind; and so forth. On this

[8] In fact, it is not implausible to think that the *best* semantic theory for English* will not include the redundant restriction to English* sentences only, in which case *s* will be meaningful in English* even in a strict sense.

[9] See Frege (1997a) and Frege (1997b).

picture, a crucial ingredient in the explanation of how sentences such as 'John runs' succeed in expressing a proposition is the fact that the individual John can fill the gap in the property of running, fusing them into one unified proposition.[10] The flip side of this picture is that if we attempt to construct the proposition which 'Jack John' supposedly expresses, we would be trying to place the individual Jack into an object that had no gap at all, and thus the attempt is doomed to fail. Similarly, if we try to construct the proposition that 'Runs eats' supposedly expresses, we would be trying to fit the property of running into the property of eating. While the latter contains a gap, it is not the right kind of gap: we cannot fit a first-order property into a gap that is intended for an individual, and thus this attempt will fail as well. The upshot is that in both cases it is a matter of metaphysical reality, combined with the semantic values of the different lexical items, that GTCs cannot express propositions or receive meanings.[11]

Unfortunately, the version of the metaphysical picture I have just sketched is hopelessly metaphorical. No semanticist seriously thinks that the semantic values of predicates are literally entities which contain gaps into which we try to fit other entities.[12] It would therefore be helpful to consider the kind of entities that semanticists do seriously think that expressions of different grammatical types denote, and consider whether it in fact follows from the metaphysical nature of, for example, the entity denoted by 'eats' that it cannot be applied to the entity denoted by 'runs'. Since semanticists disagree about which kinds of entities are denoted by different kinds of expressions, I will consider in turn three major views on this issue.

§2.1 The set-theoretic approach

One standard semantic framework, at least among philosophers, is the set-theoretic approach. According to the set-theoretic approach proper names denote individuals (let us for the moment put aside the question of what it is to be an individual and assume that we

[10] For the importance of the problem of the unity of proposition for motivating Frege's picture, see Bell (1979), particularly pp. 7-10.

[11] There is a subtle issue here about whether the metaphysical picture is supposed to show that GTCs are truth-valueless, that they fail to express propositions, or that they are meaningless. I am here primarily interested in the picture's success in motivating the latter view.

[12] Frege himself explicitly admits this. See for example his claim that "'Complete' and 'unsaturated' are of course only figures of speech; but all that I wish or am able to do here is to give hints." (Frege (1997b), p. 193).

have a given collection of entities that we call 'individuals'); first-order one-place predicates denote sets of individuals; second-order one-place predicates denote sets of sets of individuals, and so forth. On a simple version of this view the predicate 'runs' denotes the set of running things and the name 'John' denotes the individual John. The truth-conditions of the sentence 'John runs' are determined so that it is true if and only if the entity denoted by 'John' (namely John) is a member of the entity denoted by 'runs' (namely the set of running things).

It is hard to see how the last dogma can be justified within the scope of the set-theoretic approach. After all, it seems perfectly obvious how one should extend the set theoretic truth-condition specifications to GTCs such 'runs eats': the GTC would be true just in case the entity denoted by 'runs' (namely the set of running things) is a member of the entity denotes by 'eats' (namely the set of eating things). Since the set of running things is not a member of the set of eating things the string 'runs eats' is simply false and there is no reason (at least no metaphysical reason) to think it must be meaningless. Only slightly more problematic is the case of 'John Jack'. Here the natural extension of the truth conditions would entail that this string is true if and only if the entity denoted by 'John' (namely John) is a member of the entity denoted by 'Jack' (namely Jack). Now one might worry that since Jack is not a set one cannot make sense of the claim that something is a member of it. But if Jack is not a set then the sentence 'John is a member of Jack' is at worst an ordinary category mistake, which (we are assuming) entails that we *can* make sense of it. Again, the result is that the GTC is simply false rather than meaningless.

It is also worth mentioning a further complication for the defender of the last dogma in the context of the set-theoretic approach. On the set-theoretic approach expressions of different syntactic types can receive exactly the same semantic value. To see this note first that there are good reasons to think that at least some sets should count as individuals: we can use singular terms to denote sets (for example: '{John}' or 'the set containing all and only running things'), we can apply first-order predicates to these terms (for example: 'This set is large'), and if we think that there is a model the domain of individuals of which includes absolutely all existing things, then one would have to include sets as individuals in this model.[13] Second, even putting aside sets as individuals we will get expressions of different syntactic types receiving the same semantic value. Take for example a first-order empty predicate ('being blue all over and red all over') and a second-order empty predicate ('being both instantiated and not instantiated'). On the set-theoretic approach these predicates denote exactly the same entity: the empty set. And this is not the only example: compare a second-order predicate true only of the empty first-order property, and the third-order predicate true only of the second-order empty property. Both will denote exactly the same entity: the singleton of the empty set. So even if sets are never individuals we will get sets such as $\emptyset$, $\{\emptyset\}$, $\{\{\emptyset\}\}$, and $\{\emptyset,\{\emptyset\}\}$ that can be the values of expressions of indefinitely many different types.

The upshot is that on the set-theoretic approach one cannot maintain the initial metaphysical picture according to which expressions of each syntactic type correspond to semantic values of a certain metaphysical category (at least not if the metaphysical categories in question are supposed to be disjoint). Moreover, once expressions of different syntactic types are allowed to receive identical semantic values, at least some GTCs are bound to be meaningful. For example, a GTC such as 'is identical to John is a set' involves an application of the predicate 'is a set' to exactly the same entity (the set {John}) as is involved in the perfectly acceptable sentence 'The singleton set containing John is a set'. I conclude that the set theoretic approach thus renders little support to the last dogma.

§2.2 The functional approach

The second semantic approach that I would like to consider is the functional approach. This approach is probably the most popular amongst semanticists, particularly for those working in the broad framework of type-theoretic semantics such as Montague Grammar. According to a simple version of the functional approach, singular terms denote individuals; first-order predicates denote functions from individuals to truth-values; second-order predicates denote functions from functions from individuals to truth-values,

[13] Though note that considerations concerning Russell's paradox entail that on the set-theoretic framework one cannot include all sets as individuals, at least not without substantial revisions to the framework. (It will not do to merely postulate that the domain of individuals is a class rather than a set. We would also have to revise our account of predicates. For example, since every individual falls under the predicate 'is self identical', then if every set is an individual the predicate 'is self identical' cannot denote a set).

to truth-values; and so forth.[14] The functional approach is in many respects very similar to the set-theoretic approach: whereas the functional approach claims that ‘green’, for example, denotes a function *f* from individuals to truth-values, the set-theoretic approach claims that it denotes a set of individuals which can be defined in terms of *f* as {x: *f*(x)=1}. And given a first-order property on the set-theoretic construal, i.e. a set S of individuals, one can construct an analogous first-order property on the functional construal as the function *f* from individuals to truth-values, such that *f*(x)=1 if and only if x∈S.[15]

This apparent inter-definability between the functional and set-theoretic approaches has led many to view the two approaches merely as notational variants of each other,[16] at least for the fragment of language which includes expressions that both approaches can adequately represent. (The functional approach is standardly taken to be richer, because it can assign to some expressions functions whose co-domain differs from the set of truth-values. For example, adverbs can be taken to denote functions from properties to properties). However, it seems to me that at least in the context of the current debate there are substantial differences between the two approaches.

According to the functional approach, what is wrong with GTCs such as ‘runs eats’ is that ‘eats’ denotes a function from individuals to truth-values and ‘runs’ is not an individual. It follows that in predicating ‘eats’ to ‘runs’ we are attempting to apply a function to an argument which is not in its domain, and the attempt is thus doomed to fail.

---

[14] In other versions of the approach the details differ a bit. In Montague Grammar, for example, singular terms denote second-order properties, and all functions are amended to account for intensionality. But these details do not matter for the purposes of the current discussion.

[15] This correspondence between functions and sets is not to be confused with the standard correspondence between functions and sets of ordered pairs. The question of whether functions can be reduced to sets of ordered pairs or are primitive entities that cannot be so reduced is basically orthogonal to the discussion in this sub-section. In particular note that even if functions can be reduced to sets of ordered pairs, this does not settle the issue of how we are to interpret statements of the form ‘*f*(c) is P’, where c is not in the domain of *f*. This could be interpreted as saying that there exists an x, such that <c,x>∈*f* and P(x), in which case the statement in question would be straightforwardly false. Alternatively, one might treat ‘*f*(c)’ as a non-denoting singular term (‘the x such that <c,x>∈*f* ’, on one reading of the description) which would yield an analogous treatment one gets if functions are taken to be primitive.

[16] For example, Gamut (1991) say: “Sets and their characteristic functions really amount to the same thing…The concepts *characteristic function of a set X* and *set X* are interchangeable” (pp. 83-84).

This story has several apparent advantages for the defender of the dogma over the one given by the set-theoretic approach. On this story it is easier to understand what we mean when we talk of a property not being attributable to an entity. As supporters of the view might put it, the ingenuity of using functions is that they give a precise interpretation to the metaphorical picture of properties having 'gaps' into which one can only 'fit' entities of a certain kind: the 'gaps' are interpreted as the argument places of the function; the operation of 'fitting' is interpreted as functional application; and the kind of entities that can be fitted into a certain 'gap' is interpreted as the domain of the function (for the relevant argument place).[17] A further advantage of the functional approach is that it allows us to eliminate some of the overlaps between different categories: on this approach an empty first-order predicate and an empty second-order predicate do *not* denote the same entity: the former denotes a function f from individuals to truth-values such that for every individual x, $f(x)=0$ and the latter denotes a function g from functions from individuals to truth-values, to truth-values such that for every function f in g's domain, $g(f)=0$. So while the two functions resemble each other in that they both yield the value 0 for any argument, they differ in having different domains. Thus the strict distinction between properties of different types is maintained.[18]

However, the functional approach is not problem-free for the defender of the dogma either. First, the functional approach does not completely escape the problem of overlap of categories. Just as with the set-theoretic approach, we have good reasons to think that at least some functions are also individuals (functions can be denoted using singular terms, one can apply first-order predicates to them, and so forth).[19] But if some predicates denote functions that are also individuals then some GTCs are bound to be meaningful: if, for example, 'runs' denotes a function which is also an individual then the semantic value of 'runs' *is* after all in the domain of the function denoted by 'eats', so 'runs eats'

[17] Of course, it is a good question whether talk of 'argument places' of functions or of 'applying a function' aren't equally metaphorical.
[18] This is my main reason for rejecting the standard claim that properties on the set-theoretic construal and on the functional construal are interdefinable. One can read the set-theoretic property off the functional one but not vice versa: given a set, it is not determined what the domain of the relevant function is supposed to be.
[19] Though as before considerations related to Russell's paradox suggest that not every function can count as an individual.

involves no failure of functional application, and we have no reason to think it is meaningless (at least no semantic reason).

Second, even if we ignore the possibility of functions as individuals, there is a further problem for the defender of the dogma. To see the problem, consider a standard sentence such as 'John runs'. Suppose that the semantic value of 'John' is John, and of 'runs' is a function *f* from individuals to truth-values. It is clear that the meaning of 'John runs' cannot simply be the item denoted by *f*(John), for that item is merely a truth-value. If meanings were given in this fashion then there would be only two meanings for sentences to have (*the true* and *the false*), which is absurd. Nor will it help to amend the framework so that *f*(John) denotes a function from possible worlds to truth-values, rather than a mere truth-value. Even if functions from possible worlds to truth-values are deemed fine-grained enough to act as *propositions*, they are certainly not fine-grained enough to act as *meanings*: the sentences 'Fermat's last theorem is correct' and '2+2=4' clearly have different meanings.

So what is the meaning of 'John runs' according to the functional approach? There are broadly two lines of thought that the proponents of the approach can take here. According to the first, the meaning of 'John runs' is not quite the entity denoted by *f*(John) but some sort of intentional specification of this entity. For example, one might argue that 'John runs' means (roughly) that the result of applying the function *f* to John is the value 'true'. Even if 'Jack eats' receives the same truth-value as 'John runs' the relevant intentional specification of this truth-value will differ in this case, yielding a different meaning.[20]

The second line of thought of thought is this: the functional-theoretic semantics computes the meaning of sentence in a systematic compositional manner, typically via a semantic analysis tree. The semantic analysis tree for 'John runs' will indeed contain *f*(John) (i.e. a truth-value) at its root node. But according the second line of thought, the meaning of 'John runs' is given by the complete analysis tree for the sentence rather than being merely read-off the root node of the tree. Suppose for example that 'John runs' and 'John

[20] Defenders of this line of thought usually attribute it to Frege (see for example, Dummett (1981), p. 227, Heck & May (MS), and Heim & Kratzer (19998), pp. 20-22).

eats' have the same truth-value. Let *f* be the semantic value of 'runs' and *g* be the semantic value of 'eats'. The items of at the roots of the semantic analysis tree for the above sentences would be, respectively, *f*(John) and *g*(John) which will in this case be identical. But the trees themselves would not be identical: one would have a node containing the function *f* while the other will have a node containing the function *g*, thus allowing the two sentences to receive different meanings.[21]

Returning to the issue of GTCs, it strikes me that neither line of thought supports the view that GTCs are meaningless. Consider the GTC 'runs eats'. Following the first line of thought, it seems that the GTC can receive as its meaning some intentional specification such as 'The result of applying the function *g* to the function *f* is the value *true*'. Now plausibly this specification contains an empty definite description, so depending on one's view concerning such descriptions, the specification might be either truth-valueless or false. But either way, it is still a legitimate specification which endows the GTC with a meaning. And taking the second line of thought, it seems that we could assign to 'runs eats' a semantic analysis tree: the tree will look very much like the tree assigned to 'John eats', except for two difference: in place of the node containing John, we will now have a node containing the function *f*. And in place of the root note containing *g*(John), we will now have a root node which corresponds to the specification '*g*(*f*)'. Since (we are assuming) *f* is not in the domain of *g* it is plausible to interpret '*g*(*f*)' as failing to refer, and hence to interpret the root node as empty. But it is not clear why we cannot allow for analysis trees with empty root nodes. (In fact, quite independently of GTCs many semanticists are happy to allow for analysis trees with empty root nodes. For example, Heim and Kratzer take the definite article 'the' to denote a partial function of type <<e,t>,e>, which entails that sentences such as 'The king of France is bald' receives a semantic analysis tree with an empty root node).[22]

[21] As we have construed things the two sentences would still receive the same meanings if 'runs' and 'eats' happen to be co-extensional. But here the modification via possible worlds will help. (It will require us to accept that necessarily co-extensive predicates have the same meaning, but this is a lot less controversial than claiming that any two necessarily equivalent sentences have the same meaning).
[22] See Heim & Kratzer (1998), §4.4.

I conclude that according to either line of thought GTCs can receive a legitimate meaning.[23] Thus although it initially seemed more favourable to the last dogma than the set-theoretic approach, the functional approach does not validate the last dogma either.

§2.3 The nominalistic approach

The final semantic approach that I would like to consider is the nominalistic approach. According to the nominalistic approach the only linguistic expressions that denote anything are those that according to non-nominalistic accounts denote individuals. Other linguistic expressions such as first-order predicates simply do not denote at all. From a philosophical point of view the nominalistic approach has several strong motivations: It is favoured by philosophers who have metaphysical qualms about accepting abstract entities such as properties or sets. The approach also spares us the need to explain what entities we count as 'individuals' and how to motivate the distinction between individuals and other kinds of entities (as we have seen above, this issue gets quite slippery for the other accounts). Finally, the nominalistic approach is motivated by the wish to express notions such as absolutely unrestricted quantification – an apparatus which seems to run into a version of Russell's paradox on any fully denotational semantics.

The challenge, however, is to provide adequate semantics under such a limiting framework. At this point, we have nothing close to a fully developed nominalistic semantics.[24] Still, some preliminary steps have been made, and it seems to me that the most promising steps appear in the works of George Boolos and in recent expansions of his work by proponents of absolutely unrestricted quantification.[25]

---

[23] An interesting additional question is whether on the functional approach GTCs must at least be truth-valueless. The answer is that not necessarily: as we have seen, if one adopts the first line of thought coupled with a Russellian view of descriptions, then the GTC in question might be false. And if one adopts the second line of thought, it is still open to one to adopt an analogous move to the one Heim and Kratzer make with respect to sentences containing empty definite descriptions. Heim and Kratzer's semantics assign to such sentence a semantic analysis tree with an empty root node, but are reluctant to ultimately classify such sentences as truth-valueless. They thus allow for the possibility that the semantics ultimately interprets these empty nodes as standing for the value *false* (see Heim and Katzer (1998), pp. 77-78). By analogy, one could choose to interpret GTCs as being ultimately false (though for both cases there are complications concerning embeddings of such sentences under negation, which Heim and Kratzer do not discuss).

[24] There is obviously a wealth of literature defending nominalism in principle. But few nominalists take up the project of actually developing nominalistically acceptable formal semantics.

[25] See Boolos (1985), Rayo & Uzquiano (1999), Rayo & Williamson (2003), and Williamson (2003).

Boolos combines two ideas: defining the semantics so that no term of order greater than zero receives a semantic value (or at least not a *unique* de jure semantic value); and providing semantics for a second-order object-language in a meta-language that is itself second-order. Both of these ideas are in contrast with standard denotational semantics such as the set-theoretic or functional semantics. In standard denotational semantics, terms of any order receive a unique semantic value (a set or a function), and moreover, this unique semantic value is specified in the meta-language using a zero-order term (e.g. 'the set such that…' or 'the function such that…'). This essentially forces us to take a term of any order as denoting an object. On the other hand, on a Boolos-style semantics we represent in the meta-language what a predicate of the object-language expresses by talk of the different *values* that fall under the predicate (rather than the unique value that the predicate denotes), using higher-order terms in the meta-language to represent the values that fall under predicates of order greater than one. This allows us to provide a systematic and extensional description of the semantics of predicates, without committing ourselves to the claim that they denote anything. Of course, one can ask how the uses of higher-order terms in the meta-language are to be explained - don't they ultimately denote something? The Boolosean answer is that they do not: the semantics for the meta-language can be given in a meta-meta-language which is itself higher-order and so forth, all the way down. On this view, the semantics for higher-order terms simply does not need an ultimate denotational foundation.

Now at the outset one might expect that nominalistic semantics is the best avenue for the defender of the last dogma. It is no surprise, the thought goes, that denotational semantics such as functional or set-theoretic semantics ultimately render GTCs meaningful: by assigning terms of any syntactic type a semantic value which is essentially an individual, the semantics does not ultimately take type distinctions seriously enough. In contrast, nominalistic semantics take type distinctions seriously. There is a clear metaphysical difference between zero-order expressions and those of higher orders: the former denote *something* and latter do not denote at all. So from a semantic-metaphysical point of view, there is at least a crucial difference between singular terms such as 'John' and predicates such as 'runs'.

As appealing as this thought is it falls short of establishing that the metaphysical differences in question are relevant, in the sense that they suffice to establish that GTCs are meaningless. In order to consider carefully whether or not they are meaningless, it will help to have a particular semantics in hand. As a first toy model, let me define a semantics for a language that contains only zero-order constants (denoted by small case letters such as 'a') and first-order one-place predicates (denoted by superscripted uppercase letters such as 'P^1'). Our toy model will contain no quantifiers or variables. Assume also that first-order variables of the meta-language are denoted by lower case letters such as 'x', and second-order variables of the meta-language are denoted by upper case letter such as 'X'. Our toy semantics will work by defining a first-order two place predicate M, which can be thought of as expressing a kind of 'modelling relation'.[26] Thus instead of the conception of models as objects that appears in the denotational approaches, we have the idea of certain linguistic items being modelled by certain entities. We will assume that M is a legitimate modelling relation if the following holds:

(constraint): $\forall$x(x is a zero-order term of the object-language$\rightarrow\exists$!yM(x,y))

$\wedge\forall$x$\forall$y(M(x,y)$\rightarrow$x is a term of the object-language).

This means that every zero-order term is modelled by one and only one individual, and that the items being modelled are always terms of the object-language. As we shall see below, the semantics will be defined so that first-order predicates are modelled by each of the objects that fall under them.

Putting GTCs aside for the moment, we can define semantics for the 'legitimate' sentences of our language. In our toy model there is only one kind of legitimate sentence, namely sentences of the form P^1(a). Given a modelling relation M, we can define the truth-conditions for such sentences as follows:[27]

---

[26] Here and below my use of words such as 'property' or 'relation' are merely intended to facilitate the explanation. Such talk will not, of course, be literally acceptable for the proponents of the nominalistic approach.

[27] This is more or less the standard Boolosean semantics for such sentences (see Boolos (1985), p. 336).

(rule-1) A sentence of the form P^1(a) is true if and only if $\exists$y(M('a',y)$\wedge$M('P^1',y)).

What will happen with GTCs in our toy model? There are exactly three kinds of GTCs: a(b), a(P^1), and P^1(Q^1). I take it that the most natural way to extend the truth-conditions specified in rule-1 so as to apply to GTCs is to substitute 'a' in rule-1 with whatever appears in 'subject' position in the GTC (i.e. the expression in parenthesis), and substitute 'P^1' in rule-1 with whatever appears in 'predicate' position in the GTC (i.e. the expression to the left of the parenthesis).

Take the first kind of GTC: a(b). The truth-conditions of this will now be specified as follows: $\exists$y(M('b',y)$\wedge$M('a',y)). Note that not only are these well specified truth-conditions, but in fact, this specification might turn out to be true (it is true whenever 'a' and 'b' are modelled by the same individual). Next take GTCs of the form a(P^1). The truth-conditions will now be specified as follows: $\exists$y(M('P^1',y)$\wedge$M('a',y)). As before, note that there is absolutely no problem in specifying these truth-conditions, and as before, this specification might even turn out to be true (it is true whenever the item that is modelled by 'a' falls under the 'property' expressed by 'P^1', that is whenever P^1(a) will be true). Finally consider GTCs of the form P^1(Q^1). The specification of truth-conditions here will be $\exists$y(M('Q^1',y)$\wedge$M('P^1',y)). As above, this is a perfectly acceptable specification of truth conditions, and again this specification might turn out to be true (it is true whenever P^1 and Q^1 are co-instantiated). So far, the results are quite surprising: GTCs not only have perfectly well-specified truth-conditions, but a specification that sometimes determines that they are true. Moreover, note that as opposed to the case of the functional approach and the set-theoretic approach, these specifications are perfectly smooth and do not involve any 'problematic' specifications in the meta-language such as category mistakes or trees with empty nodes.

Things get somewhat more complicated when we look at slightly richer languages. Let us add to our toy model second-order predicates (denoted by expressions such as 'P^2'). This adds a new kind of 'legitimate' sentences to our language, namely sentences of the form P^2(P^1). Now the best way to extend our semantics so as to handle such sentences is to use a modelling relation that can take first-order predicates (of the meta-language) in its second argument position. This way, a second-order predicate (in the object-language) can be modelled by any first order-predicate (of the meta-language) for which it is true. However, using a modelling relation that takes first-order terms in its second argument place is problematic, because our constraint on the modelling relation and our rule-1 already assume that M takes zero-order terms in its second argument place. There are two ways to handle this problem. According to the first, we allow M to be a kind of multi-order relation that can take both zero-order and first-order arguments in its second argument place. On this proposal, we can keep the original rule-1 and add to it a second rule, which together yield the following semantics:

(rule-1) A sentence of the form $P^1(a)$ is true if and only if $\exists y(M(\text{‘a’},y) \wedge M(\text{‘}P^1\text{’},y))$.

(rule-2) A sentence of the form $P^2(P^1)$ is true if and only if $\exists X(\forall y(M(\text{‘}P^1\text{’},y) \leftrightarrow Xy)$

$\wedge M(\text{‘}P^2\text{’},X))$.

One might worry, though, that particularly in the context of a discussion on the status of GTCs, the use of multi-order relations is unacceptable. This brings us to the second way to define our semantics. According to the second way, M takes only zero-order terms in its first argument place and only first-order terms in its second argument place. But now we must modify both the constraint on the modelling relation and rule-1 so as to conform to this requirement. The best way to do this is to lift the type of the zero-order terms: For each individual there is a property that is true of it and only of it - call this ‘the characteristic property’ of that individual. The idea is that if before we had M holding between a term ‘a’ (whether zero or first-order) and an individual y, we would now have M holding between ‘a’ and the characteristic property of y. Formally, we get the following semantics:

(constraint*): $\forall x$(x is a zero-order term of the object-language

$\rightarrow \exists! y \forall X\, (M(x,X) \leftrightarrow \forall z(Xz \leftrightarrow z=y)))$

$\wedge \forall x \forall X(M(x,X) \rightarrow$ *x* is a term of the object-language).

(rule-1*) A sentence of the form $P^1(a)$ is true if and only if $\exists X(M(\text{‘a’},X) \wedge M(\text{‘}P^1\text{’},X))$.

(rule-2*) A sentence of the form $P^2(P^1)$ is true if and only if $\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow$

$(M(\text{‘}P^1\text{’},Z) \leftrightarrow Xy)) \wedge M(\text{‘}P^2\text{’},X))$.

In more informal terms: constraint* says that each zero-order term is modelled by one and only one characteristic property and that the only things modelled are terms of the object-language. Rule-1* says that $P^1(a)$ is true if and only if P^1 is modelled by the characteristic property corresponding to ‘a’. Rule-2* says that $P^2(P^1)$ is true if and only if there is a property X, such that the objects which fall under X are exactly those objects whose characteristic properties are modelled by P^1, and P^2 is modelled by X.

Personally, I prefer the first way of handling the semantics both because it is simpler and because the second way requires further revision in order to handle third-order terms. (Moreover, no such revision involving type lifting is possible if we want to handle a language with terms of any finite order). Nonetheless, for completeness I shall discuss both ways to handle the semantics. As we shall see, for the sake of the points I want to make below it makes little difference which of the two semantics we use.

In our new object-language there are seven different kinds of GTCs – the three we discussed previously, and in addition: $a(P^2)$, $P^2(a)$, $P^2(Q^2)$, and $P^1(P^2)$. Is there some natural way to extend the specified semantics so as to assign truth-conditions to these GTCs? Here we encounter a problem which has not occurred either in the denotational semantics (functional or set-theoretic) or for our simpler toy semantics above. For now we not only need to assign some kind of structure to the GTC (i.e. decide which term counts as standing in subject position and which in predicate position), but we also need to decide which of the two available rules to apply. To handle this problem, I consider all the options. In appendix-1, I summarise the results of all the possible applications when using the semantic specified by rule-1 and rule-2, and in appendix-2, I do the same for the semantics specified by rule-1* and rule-2*. I call a rule a ‘preferred rule’ if either the ‘subject’ of the GTC agrees in type with the subject specified in the rule or the ‘predicate’ of the GTC agrees in type with the predicate specified in the rule. This leaves open the option that neither rule is preferred, or that both are preferred. I have noted for each GTC which are its preferred rules (note that this does not depend on which of the two semantics we use).

The results of this exercise are similar to those we had for the simpler object-language: all possible GTCs receive a perfectly legitimate truth-conditional specification in the meta-language. This suggests that (contrary to our initial expectation) the nominalistic semantics does not render more support to the dogma than the previous approaches.

The defender of the dogma might, though, point out two seemingly unfavourable consequences of the attempt to specify truth-conditions to GTCs within this framework. The first is that there is more than one rule (sometime even more than one preferred rule) that we can apply in each case, and the results of applying different rules might yield different truth-values. For example, consider the GTC $P^2(a)$, and assume that 'P^2' expresses the property of being identical to 'a' and that P^2 is modelled by no individual. Both rule-1 and rule-2 are preferred in this case. But an application of rule-1 renders the GTC false, while an application of rule-2 renders it true. This fact, one could argue, is an indication that the whole attempt to assign truth conditions to GTCs on the nominalistic approach is highly artificial and should not be taken seriously. I don't this argument is decisive, though. One could just as well interpret those situations where applications of two rules yield different truth-conditions as showing that the GTC in question is ambiguous (it receives more than one truth-conditional specification) rather than showing that it is meaningless (i.e. it receives no truth-conditions at all).

The second seemingly problematic aspect of our exercise is that (even using only preferred rules) many GTCs come out not only as having truth-conditions, but also as being true. For example, if 'a' and 'b' denote the same individual then 'a(b)' comes out as true. This is problematic because one might expect that if GTCs are meaningful and truth-valued at all, then 'atomic' GTCs should at least come out as false. And here too, one might use this fact to argue that on the current approach the attempt to assign truth-conditions to GTCs is highly artificial and should thus be abandoned. However, I find this argument unconvincing as well: assuming the suggested nominalistic semantics is adequate in the first place, then perhaps we should simply accept that some 'atomic' GTCs are not only meaningful but also true. Moreover, it is worth pointing out that a similar result holds also in the set-theoretic and functional approaches. For example, on the set-theoretic approach, a GTC in which a third-order predicate which receives the set

{∅} as its semantic value is applied to an empty first-order predicate will be true. And assuming that the function which 'runs' denotes is an individual, then on the functional approach 'Runs is discussed in this paper' is a true atomic GTC. In so far as in the case of the denotational approaches we were willing to simply accept these results as showing that some GTCs are true, I see no reason not to accept a similar conclusion for the case of the nominalistic approach.[28]

One ought to be cautious not to make too much of the results of the exercise discussed in this subsection, for two reasons: first, I have only provided semantics for a very limited object-language. And second, these results turn out to be very sensitive to the particular semantics chosen: a different yet plausible semantics for the same language might yield different results. For example, given our constraint it will make no difference to the truth-conditions of grammatical sentences if we replace the right-hand side of rule-1 with the condition '$\forall y(M(\text{'a'},y)\rightarrow M(P^1,y))$'. However, this replacement would make a difference to the application of this rule to the GTC $P^1(Q^1)$: on the original rule it is true whenever P^1 and Q^1 are co-instantiated, whereas on the new rule it is true whenever every thing that falls under P^1 also falls under Q^1.[29]

Nevertheless, I hope the above discussion shows that it is far from obvious that on the nominalistic approach GTCs must be meaningless. And as we have seen, the same goes for the other approaches. Semantic and metaphysical considerations are certainly relevant to the question of the meaningfulness of GTC, but there does not seem to be any conclusive semantic-metaphysical reason to support the view that GTCs must be meaningless.

---

[28] The case of the truth-conditions assigned to the GTC 'a(b)' is particularly interesting: the specified truth-conditions entail that this GTCs has exactly the same truth-conditions as the claim 'a is b' - where the 'is' is read as identity. But there are several defences in the literature of the claim that 'is' is not ambiguous between the 'is' of identity and that of predication (see Lockwood (1975) and Gamut (1991), §6.3.9) and if this non-ambiguity claim is correct, perhaps it is should be quite natural to read a(b) as being true when a is identical to b.

[29] In fact, one might want to use this sensitivity in order to choose from various apparently adequate semantic formulations. For example, if on the current semantics some GTCs turn out to be counter-intuitively true perhaps we should take this as an indication that some details of the semantics ought to be amended so as to avoid this result.

§3 Logical differences?

Another reason to think that GTCs must be meaningless (or at least truth-valueless) stems from the thought that to suppose otherwise would lead us to a version of Russell's paradox, and hence to a logical contradiction.

Consider the following version of the paradox. Assume we have a language in which there is no division of predicates into orders and an application of any predicate to any other predicate counts as legitimate. Prima facie, we can define a predicate 'SA' (self applies) so that for any P, SA(P) is true just in case P(P). We can also define a predicate 'NSA' (does not self apply), so that for any P, NSA(P) is true just in case ¬P(P). We now ask whether NSA self applies or not. If it does, i.e. if SA(NSA), then by definition of SA (substituting 'NSA' for 'P'), NSA(NSA) so by definition of NSA, ¬NSA(NSA) which is a contradiction. And if it does not, i.e. if ¬SA(NSA), then by definition of SA, ¬NSA(NSA). So by definition of NSA, NSA(NSA) which is again a contradiction. We are thus presented with a paradox.

One of Russell's main solutions to the paradox involved the idea of introducing order-type distinctions. According to this solution, what went wrong in the construction of the paradox above is that our language was too promiscuous. However, the Russellean line goes, if we replace our language with one in which each predicate is given an order-type and one is only allowed to apply a predicate of order n to a predicate of order n-1 the paradox will not arise: in that case no formula of the form P(P) will be well-formed, and hence the predicates 'SA' and 'NSA' will not be well defined

Even if we are broadly sympathetic to Russell's type theoretic solution the paradox, we still face the following question: must we accept that GTCs are meaningless in order to avail ourselves of this solution? I would like to argue that the answer is 'no'. [30]

[30] Interestingly, Russell thought that it is crucial to his solution to the paradox that such GTCs be treated as meaningless and not merely false. His motivation for this was the thought that we reach a contradiction both from the assumption that SA(NSA) is true and from the assumption that it is false, so taking this sentence to be false will not help (see e.g. Russell (1919), p. 185).

Let us suppose that we did mark each predicate with an order-type distinction, but allowed formulas which were GTCs in virtue of order confusions ((e.g. $P^3(P^3)$ or $P^1(P^2)$)) to count as meaningful and truth-valued. Let us further suppose that all such atomic GTCs were deemed by default to be false. In such a setting, I argue, the paradox will not resurface (at least not in anything like its original form).

To see why, let us try to reformulate the paradox in the new language. For a start, we need to choose an order-type for the predicate 'SA'. Thus instead of one general predicate, we need to define a predicate 'SA^n' for each order n>1,. We would like SA^n to express the property of 'self applying' but also to be such that as applied to any predicate of an order other than n-1, it will be false by default. This naturally leads us to define the truth-conditions for SA^n as follows:

(1) If $m \neq n-1$ then $SA^n(P^m)$ is false.

(2) If $m=n-1$ then $SA^n(P^m)$ is true if and only if $P^m(P^m)$.

Of course, it is easy to see that since $P^m(P^m)$ is itself an atomic GTC and hence always false by default, SA^n will be an empty predicate.

Now, how should we define the negative correlate of SA^n – the property of not self applying? Here we are faced, I think, with several options:

Option-1: Define NSA^n as a predicate of order n, which works as follows:

(1) If $m \neq n-1$ then $NSA^n(P^m)$ is false.

(2) If $m=n-1$ then $NSA^n(P^m)$ is true if and only if $\neg P^m(P^m)$.

Option-2: One might think that there is no particular reason to define the negative correlate of SA^n as a predicate of order n. Instead for any n, k, one can define a separate predicate $(N\text{-}SA^n)^k$ as follows:

(1) If $m \neq k-1$ then $(N\text{-}SA^n)^k(P^m)$ is false.

(2) If $m=k-1$ then $(N\text{-}SA^n)^k(P^m)$ is true if and only if $\neg SA^n(P^m)$ is true.

Option-3: One might choose to think of '$N\text{-}SA^n(P)$' as a shorthand for the formula $\neg SA^n(P)$. On this construal '$N\text{-}SA^n$' is not taken to be a self-standing predicate and does thus not receive any order-type

Recalling that SA^n is always an empty predicate it is easy to see that according to option-1, $N\text{-}SA^n(P)$ is true if and only if P is a predicate of order n-1; according to option-2, $(N\text{-}SA^n)^k(P)$ is true if and only if P is of order k-1; and according to option-3, $N\text{-}SA^n(P)$ is true for any predicate P.

Now let us consider the paradox again. On the original formulation, the problematic issue was whether SA(NSA) was true or not. On the new version we have infinitely many versions of this question – but none of them yield a paradox. First, consider option-1. Since SA^n is an empty predicate, for any n, m, $SA^n(NSA^m)$ will be false and $\neg SA^n(NSA^m)$ will be true. And as in the original version $\neg NSA^m(NSA^m)$ will also be true, though note that this has nothing in particular to do with the definition of 'SA' or with the fact that $\neg SA^n(NSA^m)$ holds. Rather, it is true simply because $NSA^m(NSA^m)$ is an atomic GTC and thus false. Moreover, note that now we cannot infer from the fact that $\neg NSA^m(NSA^m)$ that $NSA^m(NSA^m)$ because the latter does not fall within the ambits of clause (2) in the definition of option-1, so no contradiction arises.

Next, consider option-2. Again, since SA^n is always an empty predicate, we will get that for any n, m, and k, $SA^n((N\text{-}SA^m)^k)$ is false, and $\neg SA^n((N\text{-}SA^m)^k)$ is true. And as above, we will get that $\neg(N\text{-}SA^m)^k((N\text{-}SA^m)^k)$ is also true, but that $(N\text{-}SA^m)^k\ ((N\text{-}SA^m)^k)$ is nevertheless false – simply because it is an atomic GTC. Again, no contradiction arises.

Finally, consider option-3. Here it is not even clear that the expression $SA^n(N\text{-}SA^m)$ is a legitimate one. The problem is not the expression is a GTC. Rather, the problem is that as I have formulated option-3, the standalone '$N\text{-}SA^m$' is not really an expression of our language – we have only introduced full formulas of the form '$N\text{-}SA^m(P)$'. At any rate, even if we somehow allow '$N\text{-}SA^m$' to count as a legitimate self-standing predicate (one which is true for any predicate P of any order if and only if $\neg SA^m(P)$), it will not receive any type-order, and hence in particular it will not count as a predicate of type n-1. But given this, it is plausible to interpret $SA^n(N\text{-}SA^m)$ as a kind of atomic GTC which is again false by default. This will imply that $SA^n(N\text{-}SA^m)$ is false, and for the case of m=n we would be able to infer from this $N\text{-}SA^n(N\text{-}SA^n)$ is true. But this does not imply that $SA^n(N\text{-}SA^n)$ is true because $N\text{-}SA^n$ is not a property of order n-1, and hence the second clause in the definition of SA^n does not apply to it.

Russell may have been right to think that order-type distinctions are logically important, in particular in the context of avoiding the paradox. But it does not follow from this that we have to take sentences involving confusions of these distinctions to be meaningless or truth-valueless: taking them to be merely false by default is sufficient in order to block the paradox. Logical reason do not, therefore, supply an adequate justification for the last dogma.

§4 Syntactic reasons

The final avenue for the defender of the dogma is to appeal to the most obvious feature of GTCs, namely that they are syntactically ill-formed. The argument might run as follows: the meaning of a (standard) sentence is determined by the meanings of the lexical items in the sentence together with some mode for combining these meanings. The mode of combination is given to us by the syntactic structure of the sentence. Now by definition, GTCs contain only meaningful words. However also by definition, GTCs are syntactically ill-formed and hence they receive no syntactic structure which in turn entails that there is no there is no mode for combining the meanings of the words in the GTC. So, the argument goes, GTCs must be meaningless.

The suggested argument consists of two claims. Claim one: since GTCs are syntactically ill-formed they do not have any syntactic structure. Claim two: if GTCs lack syntactic structure they cannot be meaningful. But both claims can, I think, be questioned.

Consider very mildly ungrammatical strings such as 'John drink beer' (rather than 'John drink*s* beer'). It seems plausible to think that despite not being strictly speaking generated by the syntax, such strings can be naturally assigned a syntactic structure. One way to explain how we might assign a structure to such syntactically ill-formed strings is to appeal to an idea of Chomsky's, in his discussion of the notion of 'degrees of grammaticality'.[31] For simplicity, let us fix on the syntactic model proposed by Chomsky in his *Aspects of the Theory of Syntax.* According to this model, one set of syntactic rules generates 'bare' phrase markers (syntactic analysis trees) which terminate with complex symbols: symbols which contain general syntactic labels such as 'Noun' as well as slightly more specific ones indicating features such as gender or number. In addition, the

[31] Chomsky (1965), pp. 148-153.

syntax includes lexical insertion rules that allow us to replace these abstract complex symbols with particular lexical items of the language so as to generate the final phrase markers for sentences. The lexical insertion rules prevent sentences with violations of number agreement such as in 'John drink beer' to be generated by the syntax. However, we could simply choose to ignore some of the restrictions imposed by the lexical insertion rules and generate phrase markers for ungrammatical sentences in this manner. For example, we can take the bare phrase marker for 'John drinks beer' and insert the plural verb 'drink' in place of the complex symbol which would otherwise receive the singular verb form 'drinks', thus assigning a syntactic structure to 'John drink beer'. Chomsky also defines a hierarchy of the restrictions on the lexical insertion rules, and maintains that violations of restrictions lower in the hierarchy make for a higher degree of grammaticality. The upshot is *not* that sentences such as 'John drink beer' are ultimately syntactically well-formed: to be well-formed a sentence needs to be grammatical to the highest degree – namely to violate no restrictions whatsoever. The upshot is rather that despite being syntactically ill-formed there is a natural way to assign a syntactic structure to the sentence. But if an idea along these lines is accepted, it clear that we can extend it so as to apply to GTCs such as 'runs eats': we start with the bare phrase marker for 'John eats', and then, ignoring most of the restrictions imposed by the lexical insertion rule for the complex symbol corresponding to 'John' (including the restriction that indicates that one expects a noun rather than a verb to be inserted in place of the symbol), one can generate a syntactic structure for the GTC 'runs eats'. It is of course open to debate whether a structure assigned in this way should indeed count as *the* structure of 'runs eats', but without further argument it is not clear that it should not.

This brings me to the second claim. Suppose a GTC does not have a syntactic structure. Does this entail that there is no mode for combining the meanings of the words in the GTC? This, too, is far from clear. One needs to separate two distinct reasons for thinking that we have no way of combining the meanings in such cases. The first is that in the case of a GTC such as 'runs eats' the lexical items involved in the sentence (namely 'runs' and 'eats') are such that, given the standard semantic modes of combinations available in our language, there simply does not exist any legitimate mode of combination that enables us to 'fit together' their meanings. This thought may be initially compelling, but I take it that the discussion in §2 of this paper shows that this thought is incorrect: whatever modes of combination are at play in the case of grammatical sentences (for example, set membership or functional application) these can be extended to the case of GTCs.

The second worry is that rather than there being *no* way to combine the meanings of the words in question the problem is that there is *more than one* way to combine these meanings. For example, in the case of 'runs eats' one can apply 'eats' to 'runs' or 'runs' to 'eats'. The problem, according to this worry, is that in the absence of a syntactic structure to dictate which mode of combination to use, any mode of combination we choose is arbitrary, and thus illegitimate. But this complaint too is not decisive. First, even if there is more than one available mode of combination, this does not mean that any choice is arbitrary: some modes of combination might be more natural than others on either semantic grounds (for example, we might prefer modes that do not render atomic GTCs to be true) or syntactic grounds (for example, we appeal to the similarities between the word order in the GTC and the word order in grammatical sentences). More importantly, even if there is no unique natural choice for a mode of combination, it seems to me that the situation ought to be interpreted as one in which the GTC in question is meaningful but structurally ambiguous rather than meaningless.

Syntactic consideration do perhaps indicate some degree of artificiality in the attempt to assign GTCs meanings, but the facts that GTCs are syntactically ill-formed does not provide a conclusive reason to think that they are meaningless.

§5 Conclusion

In this paper I have discussed what I take to be the main reasons for thinking that GTCs are meaningless. I have argued that none of these reasons are conclusive. To be clear, I do not take myself to have established that GTCs are meaningful. But I do hope to have shown that rather than taking it from granted, we need to consider carefully whether we ought to keep holding on to the last dogma of type confusions.[32]

[32] I am very grateful to ∅ystein Linnebo, Menachem Magidor, Nicholas Shea, Bruno Whittle, Timothy Williamson, and audience at the Mathematical Methods in Philosophy workshop in Bristol and in the Aristotelian Society meeting in London for helpful discussion and comments on this paper.

References

Beall & van Fraassen (2003) Beall, J. C., and van Fraassen, B. C., *Possibilities and Paradox*, Oxford University Press.

Bell (1979) Bell, D., *Frege's Theory of Judgment*, Clarendon Press.

Boolos (1985) Boolos, G., 'Nominalist Platonism', *Philosophical Review* 94, pp. 327-344.

Chomsky (1957) Chomsky, N., *Syntactic Structures*, Mouton.

Chomsky (1965) Chomsky, N., Aspects of the Theory of Syntax, MIT press.

Diamond (2001) Diamond, C., 'What nonsense might be', in Diamond, C., *The Realistic Spirit*, MIT Press.

Drange (1966) Drange, T., *Type Crossings*, Mouton.

Dummett (1981) Dummett, M., *Frege. Philosophy of Language*, 2ed ed., Harvard University Press.

Fine (2003) Fine, K., 'The Non-Identity of a Material Thing and Its Matter', *Mind* 112, pp.195-234.

Frege (1997a) Frege, G., 'Function and concept', trans. by Geach, P., in Beaney, M. (ed.), *The Frege Reader*, Blackwell.

Frege (1997b) Frege, G., 'On concept and object', trans. by Geach, P., in Beaney, M. (ed.), *The Frege Reader*, Blackwell.

Gamut (1991) Gamut, L.T.F., *Logic, Language, and Meaning*, vol. 2, Chicago University Press.

Heck & May (MS) Heck, R., and May, R., 'The composition of thoughts', (unpublished manuscript of the 2006 Gareth Evans Lecture).

Heim & Kratzer (1998) Heim, I. and Kratzer, A., *Semantics in Generative Grammar*, Blackwell.

Honderich (1995) Honderich, T., (ed.), *The Oxford Companion to Philosophy*, Oxford University Press.

Lockwood (1975) Lockwood, M., 'On predicating proper names', *Philosophical Review* 84, pp. 471-498.

Magidor (2007) Magidor, O., *Category Mistakes*, DPhil thesis, University of

Oxford.

Prior (1954) Prior, A. N., ‘Entities’, *Australasian Journal of Philosophy* 32, pp. 159-168.

Rayo & Uzquiano (1999) Rayo, A., & Uzquiano G., ‘Toward a theory of second-order consequence’, *Notre Dame Journal of Formal Logic* 40, pp. 315-325.

Rayo & Williamson (2003) Rayo, A., and Williamson, T., ‘A completeness theorem for unrestricted first-order languages’, in Beall, J.C., (ed.), *Liars and Heaps*, Oxford University Press.

Routley (1966) Routley, R., ‘On significance theory’, *Australasian Journal of Philosophy* 44, pp. 172-209.

Russell (1903) Russell, B., *The Principles of Mathematics*, Cambridge University Press, 1903.

Russell (1919) Russell, B., *Introduction to Mathematical Philosophy*, Allen and Unwin.

Russell (1967) Russell, B., ‘Mathematical logic as based on the theory of types’ in van Heijenoort, J. (ed.), *From Frege to Gödel*, Harvard University Press.

Ryle (1949) Ryle, G., ‘Categories’, in Ryle, G., *Collected Papers* (vol. 2), Hutchinson.

Sorensen (2001) Sorensen, R., *Vagueness and Contradiction*, Oxford University Press.

Strawson (1952) Strawson, P. F., *Introduction to Logical Theory*, Methuen.

Williamson (2003) Williamson, T., ‘Everything’, *Philosophical Perspectives* 17, pp. 415-465.

GTC Type	Preferred rule	Application of rule-1	Remarks on truth-value (rule-1)	Application of rule-2	Remarks on truth-value (rule-2)
a(b)	rule-1	$\exists y(M(\text{'b'},y) \wedge M(\text{'a'},y))$	True when ‘a’ and ‘b’ denote the same individual.	$\exists X(\forall y(M(\text{'b'},y) \leftrightarrow Xy) \wedge M(\text{'a'},X))$	Might be true, unless we restrict M so that zero-order terms are never modelled by first-order ‘entities’.
$b(P^1)$	rule-2	$\exists y(M(\text{'}P^1\text{'},y) \wedge M(\text{'a'},y))$	True if $P^1(a)$ is true	$\exists X(\forall y(M(\text{'}P^1\text{'},y) \leftrightarrow Xy) \wedge M(\text{'a'},X))$	Might be true, unless we restrict M so that zero-order terms are never modelled by first-order ‘entities’.
$a(P^2)$	neither	$\exists y(M(\text{'}P^2\text{'},y) \wedge M(\text{'a'},y))$	Might be true, unless we restrict M so that second-order terms are never modelled by an individual	$\exists X(\forall y(M(\text{'}P^2\text{'},y) \leftrightarrow Xy) \wedge M(\text{'a'},X))$	Might be true, unless we restrict M so that zero-order terms are never modelled by first-order ‘entities’.
$P^1(Q^1)$	Both	$\exists y(M(\text{'}P^1\text{'},y) \wedge M(\text{'}Q^1\text{'},y))$	True whenever P^1 and Q^1 are co-instantiated.	$\exists X(\forall y(M(\text{'}Q^1\text{'},y) \leftrightarrow Xy) \wedge M(\text{'}P^1\text{'},X))$	Might be true, unless we restrict M so that first-order terms are never modelled by first-order ‘entities’.
$P^1(P^2)$	rule-1	$\exists y(M(\text{'}P^2\text{'},y) \wedge M(\text{'}P^1\text{'},y))$	Might be true, unless we restrict M so that second-order terms are never modelled by an individual	$\exists X(\forall y(M(\text{'}P^2\text{'},y) \leftrightarrow Xy) \wedge M(\text{'}P^1\text{'},X))$	Might be true, unless we restrict M so that first-order terms are never modelled by first-order ‘entities’.
$P^2(a)$	Both	$\exists y(M(\text{'a'},y) \wedge M(\text{'}P^2\text{'},y))$	Might be true, unless we restrict M so that second-order terms are modelled by an individual	$\exists X(\forall y(M(\text{'a'},y) \leftrightarrow Xy) \wedge M(\text{'}P^2\text{'},X))$	Might well be true, e.g. if P^2 holds of the property of being identical to ‘a’.
$P^2(Q^2)$	rule-2	$\exists y(M(\text{'}Q^2\text{'},y) \wedge M(\text{'}P^2\text{'},y))$	Might be true, unless we restrict M so that second-order terms are never modelled by an individual	$\exists X(\forall y(M(\text{'}Q^2\text{'},y) \leftrightarrow Xy) \wedge M(\text{'}P^2\text{'},X))$	Might be true (e.g., if Q^2 is not modelled by any individual, and P^2 is true of the empty first-order ‘property’).

Appendix 1: *Analysis of GTCs according to nominalistic semantics with a multi-order modelling relation*

GTC Type	Preferred rule	Application of rule-1*	Remarks on truth-value (rule-1*)	Application of rule-2*	Remarks on truth-value (rule-2*)
a(b)	rule-1*	$\exists X(M(\text{'b'},X) \wedge M(\text{'a'},X))$	True whenever 'a' and 'b' denote the same individual	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'b'},Z) \leftrightarrow Xy))) \wedge M(\text{'a'},X))$	Will be true whenever 'a' and 'b' denote the same individual.
a(P^1)	rule-2*	$\exists X(M(\text{'}P^1\text{'},X) \wedge M(\text{'a'},X))$	True if P^1(a) is true	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'}P^1\text{'},Z) \leftrightarrow Xy))) \wedge M(\text{'a'},X))$	Will be true if P^1 expresses the characteristic property of 'a'.
a(P^2)	neither	$\exists X(M(\text{'}P^2\text{'},X) \wedge M(\text{'a'},X))$	Will be true if P^2 is true of the characteristic property of 'a'.	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'}P^2\text{'},Z) \leftrightarrow Xy))) \wedge M(\text{'a'},X))$	Will be true if the only characteristic property P^2 is true of is that of 'a'.
$P^1(Q^1)$	both	$\exists X(M(\text{'}Q^1\text{'},X) \wedge M(\text{'}P^1\text{'},X))$	True whenever P^1 and Q^1 are co-instantiated.	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'}Q^1\text{'},Z) \leftrightarrow Xy))) \wedge M(\text{'}P^1\text{'},X))$	Might be true (e.g. if Q^1 is true of exactly one individual, for which P^1 is true)
$P^1(P^2)$	rule-1*	$\exists X(M(\text{'}P^2\text{'},X) \wedge M(\text{'}P^1\text{'},X))$	Might be true (e.g. if P^2 is true of the characteristic property of some individual for which P^1 is true).	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'}P^2\text{'},Z) \leftrightarrow Xy))) \wedge M(\text{'}P^1\text{'},X))$	Might be true (e.g. if P^2 is true of the characteristic property of exactly one individual, and P^1 is true of that individual).
P^2(a)	both	$\exists X(M(\text{'a'},X) \wedge M(\text{'}P^2\text{'},X))$	True if P^2 is true of the characteristic property of 'a'.	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'a'},Z) \leftrightarrow Xy))) \wedge M(\text{'}P^2\text{'},X))$	True if P^2 is true of the characteristic property of 'a'.
$P^2(Q^2)$	rule-2*	$\exists X(M(\text{'}Q^2\text{'},X) \wedge M(\text{'}P^2\text{'},X))$	Will be true if Q^2 and P^2 are co-instantiated.	$\exists X(\forall y \forall Z(\forall x(Zx \leftrightarrow x=y) \rightarrow (M(\text{'}Q^2\text{'},Z) \leftrightarrow Xy))) \wedge M(\text{'}P^2\text{'},X))$	Will be true, if Q^2 is true of the property of being an individual for which P^2 holds of its characteristic property.

Appendix 2: *Analysis of GTCs according to nominalistic semantics with type lifting*