

APOHA SEMANTICS—SOME SIMPLE-MINDED QUESTIONS AND DOUBTSⁱ

Bob Hale, University of Sheffield

1. INTRODUCTORY REMARKS

Some Buddhist Nominalists thought, or can at least plausibly be interpreted as claimingⁱⁱ, that a commitment to universals can be avoided by explaining the meaning of general predicates or kind terms by exclusion, appealing to their theory of apoha. In broad terms, their idea seems to have been that whereas there is a strong temptation to think that different particulars to which a general predicate such as ‘animal’ can be correctly applied must have something in common, there is no such temptation to suppose that different particulars to which such a general predicate does not apply must do so, so that if we can understand the positive general predicate as conveying no more than that the particulars to which it applies are not non-animals—as ‘excluding the other’ (anyapoha)—we can see that there is no commitment to shared universals after all.

I shall pursue two main questions: (1) How is the central idea—of avoiding ontological commitment to universals by recourse to doubly negative predications—to be understood, and how is it supposed to work? (2) What difficulties does it face, and is it possible to overcome them?

2. WHAT IS THE DOCTRINE?

I want to begin by identifying what I see as the main elements of the Buddhist Nominalist solution to the problem of universals. As I understand it, the doctrine of apoha involves the following leading ideas or claims:

1. The meaning of a kind-term K is to be explained in terms of negation as: not non-K

As Mark Siderits (“Buddhist Nominalism”, 95) puts it

... the meaning of a kind-term is the 'exclusion of the other' (anyapoha).

This builds on the idea that since a given predicate determines a bipartition of the world, mastery of that predicate may be expressed either as the ability to tell when the expression does apply or the ability to tell when it does not.

So...the meaning of 'crow' may be given as: not non-crow. What all the things that are called crows have in common, then, is that they are not in the class of non-crows.

2. By giving the meaning of kind-terms in this way, one can avoid commitment to universals.

The thought here (cf. Siderits, loc cit) is that while we might be tempted to think that there has to be a single shared character which all crows have in common, there is little or no temptation to think that there must be a single shared character which all non-crows—as Siderits observes, a very disparate assortment, including 'ostriches, teapots, the number 7, and the wizard Gandalf'—have in common.

3. The two negations in 'x is not non-K' are importantly different—'non-K' is the nominally-bound negation of the kind-term 'K', whereas 'not' is verbally-bound negation.

The terms 'nominally-bound' and 'verbally-bound' are taken from B.K.Matilal (cf. Siderits, loc cit). I think we can put the intended point by saying that the prefix 'non' applies to nouns to form a kind of term-negation, whereas 'not' is used as a sentential operator, since if we think of negation as attaching to the verb in a simple sentence, this is equivalent to taking it to be the sentence as a whole (or the proposition expressed by it) that is negated. Crucially, 'x is non-K' and 'x is not K' are to be distinguished. The latter is understood like negation in classical logic, in which the Principle of Bivalence is assumed, so that 'x is not K' is true if and only if 'x is K' is not true, and the classical logical Law of Excluded Middle holds (i.e. ' $p \vee \neg p$ ' is true for all choices of p). But while 'x is non-K' is incompatible with 'x is K', it is not its

contradictory. So 'x is K or x is non-K' is not a logical law—it may be that neither disjunct is true, just as Jones may be neither kind nor unkind.

The difference is important because while 'x is not not K' (i.e. '¬¬ x is K') collapses, classically, to 'x is K', 'x is not non-K' does not. This may seem to enable defenders of the doctrine to avoid the objection that our ability to determine what belongs to the class of non-crows, for example, is parasitic upon our ability to determine which are the crows. The objection would be good, if both negations were classical, because then—applying the principle of Double Negation Elimination—'x is not non-K' would simply reduce to 'x is K'. But 'the application of verbally bound negation to a predicate formed through the nominally bound negation of some predicate P does not yield a straightforward assertion of P ... To say of something that it is not non-P is just to refuse the characterization of the thing in question as non-P without commitment to any positive characterization' (Siderits, "Buddhist Nominalism", 95).

These, then, are the main ideas or claims I shall discuss.

3. SOME DOUBTS AND DIFFICULTIES

Most of the difficulties I want to discuss focus upon one general doubt about the central idea of apoha semantics—that the meaning of a kind term K is to be explained as not non-K—or arise in connection with responses to it. This general doubt turns on the widely accepted idea that some sort of principle of compositionality has to be respected by any satisfactory semantical theory for a particular language. But I'll begin by briefly airing another doubt, about the thought which Siderits presents as underlying this central idea.

POSITIVE AND NEGATIVE APPLICATION CONDITIONS

The thought, as Siderits puts it in the passage I've quoted, is that predicate mastery can be taken equally to consist either in the ability to determine when the

predicate applies or in the ability to determine when it does not apply. In general terms, this idea has a good deal of plausibility. But this initially plausible idea—or at least its application in the present context—seems to me to be undermined by the introduction of nominally-bound negation.

The point is this: when a predicate P effects an exhaustive partition of the universe of objects into two exclusive classes, one containing the objects to which P applies and the other those to which it does not apply, then if we have learned to determine when P does not apply, and we know that P applies to all and only those objects x such that it is not the case that P does not apply to x, we are in position to know just which objects are the ones to which P does apply. So we can arrive at an understanding of P indirectly, in this way, just as well as directly, by learning P's positive application conditions. To put the essential point another way: if I am told which objects do not belong to the extension of P, I can tell, for any object whatever, whether it belongs to the extension of P or not. For given any object x, I can ask myself: Is x one of the objects which I have been told does not belong to the extension of P? By hypothesis I know the answer. If, and only if, the answer is that x is not one of those objects, I can infer that it belongs to the extension of P.

In the case with which we are concerned, however, things are more complicated. We are, as in the case just described, to arrive at an understanding of the application conditions for P indirectly, by way of first grasping the application conditions for non-P and then taking P to apply just in case non-P does not. But while non-P effects an exhaustive partition of the universe of objects into two exclusive classes, just as P does, it does not apply to all and only those objects to which P does not apply, only to some proper subclass of them. So here is the picture:

P	Not P	Not P
Not non-P	Not non-P	non-P

It is a good question what ability can constitute my grasp of the meaning of 'non-P' in this case—whether it must consist in my being able to determine which objects lie in its extension, or whether it might instead consist in an ability to determine which objects are the ones to which 'non-P' does not apply. But the question that most concerns us is not this one, but what my grasp of the meaning of 'P' could consist in. In particular, could it consist in my ability to determine which objects are not non-P? And it seems to me the answer must be that it can't. For if all I know is that an object is not non-P, I am in no position to infer that it is P—as my picture indicates, it may just as well not be P, even though it is not non-P.

In short, the very feature of nominally-bound negation that distinguishes it from verbally-bound or sentential negation appears to render it unfit to serve as a starting point for an indirect route to a mastery of a positive predicate P by exclusion (i.e. x is P if and only if x is not non-P). If this is right, there is an awkward—and perhaps fatal—tension in the doctrine we are examining. For the interpretation of 'excluding the other' in terms of a combination of sentential and term-negation—rather than just double sentential negation or double term-negation—is required to give plausibility to the idea that by understanding general predicates in terms of exclusion can avoid commitment to universals. If both negations are of the same type, the doubly negative predication just collapses back into the positive predication, and it is hard to see how the latter's unwanted apparent commitments are supposed to be avoided. But if 'x is not non-P' isn't equivalent to 'x is P', and so to ' $\neg\neg x$ is P', how can one explain the meaning of the positive predications by means of the doubly negative one?

I shall not pursue this difficulty further just now, but will return to it briefly later. First I want to press what appears to be, if anything, an even more serious difficulty.

COMPOSITIONALITY

An initial—and at least initially serious—cause for concern arises from the theory’s apparent disregard for considerations of compositionality. It is very widely agreed that a competent speaker’s knowledge of the meanings of most of the expressions of her language isn’t, and couldn’t be, acquired through a separate process of training in the use of each expression piecemeal. Rather, her understanding of complex expressions in general results somehow from her knowledge of the meanings of their parts together with a grasp of the semantic significance of the way they are put together. In broad terms, the worry is that any attempt to understand or analyse a positive predication ‘x is P’ as some sort of doubly negative predication ‘x is not non-P’ offends against the extremely plausible idea that the meaning of a complex expression has to be understood in some such way—as a function of, or as determined by, or somehow composed out of, the meanings of its parts. If, for example, the claim is that ‘x is P’ means, or can be analysed as, ‘x is not non-P’, the objection will be that the contained sentence ‘x is non-P’ and the whole containing ‘It is not the case that x is non-P’ cannot be understood save on the basis of a prior and independent grasp of the meaning of the constituent predicate ‘P’—so there can be no question of explaining what ‘x is P’ means by recourse to the more complex paraphrase. This would just get the meaning-dependency relations back-to-frontⁱⁱⁱ.

Two quite radical responses to this worry would be:

- (i) To accept the principle of compositionality for senses or meanings, but reject compositionality for reference, and argue that the proposal does not require claiming that ‘x is P’ has its sense or meaning given by ‘x is not non-P’—and thus only requires rejecting compositionality of reference.
- (ii) To accept both forms of compositionality, but reject the idea that the proposal requires any claim of equivalence, either in terms of sense or in terms of reference/existential commitment. Instead, the idea is that the nominalist could agree that positive predications carry commitment to universals, but can argue that no

important loss of information is incurred by using ‘x is not non-P’ instead of ‘x is P’, thereby avoiding the acknowledged commitments of ‘x is P’

I call these radical responses because they both involve jettisoning what seems to me the distinctive idea of apoha semantics, which is a claim about what positive predications mean. It is mainly for this reason that I am going to set them aside, at least for now. Instead I want first to explore the question whether the central idea of apoha semantics can be presented and developed in such a way that—so far from conflicting with plausible principles of compositionality for sense or reference—it actually constitutes a kind of compositional theory of meaning which explains how general predicates have meaning in a nominalistically acceptable way, i.e. without implying or presupposing the existence of shared characteristics belonging to the objects to which they apply^{iv}.

4. NOMINALISTIC COMPOSITIONAL SEMANTICS

In any compositional semantics, the expressions of the language are divided into two classes—there will be a base class of simple expressions which are, in a certain sense, primitive, together with a class (the derived class) of complex expressions, whose meanings can be seen as built up out of those of simpler expressions, and ultimately from those of the simplest expressions. Normally, the base class will be finite, and normally the class comprising all other expressions belonging to the language will be, at least potentially, infinite. Expressions in the base class will typically be syntactically simple, but anyway are always semantically simple in the sense that their meanings are not further decomposed, and so cannot be worked out on the basis of the meanings of other expressions and therefore have to be learned directly in some way, such as ostensive training. Other expressions will typically be syntactically complex, but need not be, since syntactically simple expressions may be semantically complex, being introduced by definition on the basis of other expressions. Such definitions may be explicit—as illustrated by the

definition of 'vixen' as meaning 'female fox'—or implicit—as illustrated by the usual recursive definition of '+' by means of the recursion equations: $a + 0 = a$, $a + b' = (a + b)'$. Some—but not all—of the expressions in the base class may be associated with constructions by means of which complex expressions are formed from simpler ones—their meanings will be given by rules which in some way fix the meanings of the resulting complex expressions in terms of those of their constituents. A particularly important case of this kind is the explanation of the (classical) logical constants by stipulations or rules which give the truth-conditions of complex sentences formed by their means.

From our present point of view, the interesting question is what (kinds of) expressions may be taken as lying in the base class for a nominalistic semantic theory. I shall not try to answer this question fully, but will give what I hope is a plausible partial answer—this will be enough for my purposes. Two points—one negative, the other positive—are, I think, pretty obvious, but important enough to justify emphasising them.

The negative point is that the nominalist's base class cannot—with crucial exceptions to which I shall soon attend—include general predicates. His goal is to explain the meanings of these expressions in a way which avoids any assumption of a common characteristic belonging to all and only the entities to which such expressions apply. This is to be accomplished by showing how their meanings result, in a nominalistically acceptable way, from those of expressions in the base class.

The positive point is that there is no reason why the base class should not include names of particulars of one sort or another, and it surely will include semantically simple such names. Thus if one took Plato to be a particular, and took the use of the proper name 'Plato' to be learned by direct association with its bearer, 'Plato' might be taken to belong to the base class. But a complex name for the same particular—such as 'the author of Theaetetus and The Sophist'—would not do so, because it involves a general relational predicate.

Clearly, however, the nominalist's base class must comprise some expressions of some other kinds, since from simple names alone, no sentences can be formed, only lists of particulars. Without some further expressions, it will not be possible to say anything using only expressions belonging to the base class. But if sentential operators and other expressions are to be explained in a way that is at all plausible, there must be semantically basic sentences—atomic sentences—to anchor those explanations.

I claim that there must be at least one general predicate. But there is a strong constraint governing the admission of predicates into the base class—any admissible basic predicate must be such that it is evident, or can easily be seen, that its application does not require the existence of any shared characteristic belonging to the objects to which it applies.

I think that perhaps only one predicate—the general predicate of identity—can plausibly be taken to satisfy this constraint, and I shall assume that the nominalist can include it in his base class. In this case, one might argue as follows: Obviously each object is identical with itself and with no other object. Let a and b be distinct objects. Then each of them is identical to itself— $a = a$ and $b = b$. But clearly the truth of these two statements does not require that a and b should have anything in common—they can be as different as you like. And if we consider the whole universe of objects, then there is a true statement of self-identity for each of these objects, but this plainly does not require that there be some shared characteristic which all objects whatever have in common.

It might be thought that a similar argument could be given to justify including the opposite of the identity predicate—the distinctness predicate—in the base class. One would need to argue that where $\{a,b\}$ and $\{c,d\}$ are disjoint pairs of distinct objects, there need be nothing in common between the two pairs (i.e. in order for it to be true that $a \neq b$ and that $c \neq d$). Maybe one could argue for that—although I'm not sure how—but in any case the nominalist will be able to introduce the distinctness

predicate anyway, provided he can include (sentential) negation. For he can then just define $x \neq y$ to mean $\neg x = y$.

I can see no reason why the nominalist can't have sentential negation, along with the other standard sentential operators (conjunction, disjunction, etc) in his base class. But as we know, a proponent of apoha semantics requires another kind of negation operator, which can be applied to nouns rather than complete sentences (or, equivalently, in effect, to verbs). There is nothing to prevent the nominalist from introducing a form of term- as opposed to sentence-negation. So far, the only nouns in our nominalistic language are proper names of particulars. So it will be to these to which term-negation is applied, at least in the first instance, at the basic level, rather than common nouns like 'crow', 'man', etc. With this point in the open, it is also clear that more needs to be said about how term-negation is supposed to work. If term-negation ('non-') is applied to a general term, such as 'man', we naturally assume that the resulting compound 'non-man' is another general term—so in this kind of case, 'non-' forms general terms from simpler general terms. But this can't be how it works at the basic level, where it is proper names rather than common nouns which are negated. Here we take a proper name—'Socrates', say—and form a new expression—'non-Socrates'. Pretty clearly the output expression has to be of a different logico-syntactic category from the input—'non-Socrates' isn't a proper name. Whatever negating a proper name gives us, it won't be another proper name. Obviously the idea has to be that, at this basic level, applying 'non-' to a proper name results in a general term or predicate.

There is nothing objectionable in that, provided that the general predicates thereby produced are nominalistically acceptable—that is, don't require that the objects to which they apply have some shared character. But it seems clear that this condition is met—'non-Socrates', for example, is true of each and every one of those objects which is distinct from Socrates, and it is perfectly clear that there need be nothing which all of these objects have in common^v.

But there is a snag. Suppose 'n' is any name. Then 'x is non-n' will be true if and only if 'x is not n' (i.e. $\neg x = n$) is true. So 'x is not non-n' will be equivalent to 'x is not not n' (i.e. $\neg\neg x = n$), and so to 'x is n' (i.e. $x = n$) after all. That is, the complex predicate '... is not non-n' can, it seems, be true of one and only one object, namely n itself—so our attempt to exploit the combination of sentence- and term-negation to get a general predicate which clearly assumes no shared character but is potentially applicable to many objects seems to have fallen flat.

Siderits confronts this snag, but is undaunted, because he thinks there is a way around it. What is needed, evidently, is to explain 'non-n' in such a way that it is not equivalent to the general predicate 'distinct from n' (i.e. $\neq n$). Indeed, we need 'non-n' to be not merely different in meaning from ' $\neq n$ ', but to diverge from it extensionally—more precisely, we want 'non-n' to be true of no more than a proper subset of the universe of objects apart from n, i.e. $V - \{n\}$. How is that to be achieved? Siderits proposes that we associate the original name 'n' with a "paradigm image p_n which is formed so as to be manifestly incompatible with some but not all of the remaining particulars in the universe" ("Buddhist Nominalism", 96). So—to continue with his example—if n is a particular crow, we might associate with n an image which is incompatible with the perceptual images we get from ostriches, tumbleweeds, ... and so on for many other kinds of thing, but is not incompatible with the perceptual images we get from other crows. Then—or so, I take it, the thought runs—non-n will apply more selectively, in the way desired, not to every object distinct from n, as before, but to just those objects which are incompatible with the image p_n . And the class of objects which are not non-n will comprise, not just n itself, but with it any other objects which are not incompatible with p_n —all the crows, in the example.

5. FURTHER DOUBTS

Does the paradigm image manoeuvre work? In the remainder of this paper, I shall try to do two things. First, I'll try to argue that if it does work, it can turn the trick for the nominalist without reliance on any special doctrine about doubly negated predicates. Second, I'll raise some independent doubts about whether the paradigm image manoeuvre can work.

In fact, I can't see that much argument is needed for the first of these claims. The point is quite simple. If, by associating with a particular object n a certain paradigm image, we can ensure that the negative term 'non- n ' applies, not to everything in the universe other than n itself, but only to some of the objects distinct from n —all the non-crows, say—then what is to prevent us from directly introducing a non-negative general term ' n^+ ' with the stipulation that it is to be true of exactly those objects which are compatible with the paradigm image p_n associated with n ? The play with the not non- n construction is just an idle wheel—all the real work is done by the paradigm image manoeuvre. Another way to put this point is that 'non- n ' will differ in extension from ' $\neq n$ ' only if the term to which 'non-' is prefixed applies to more objects than just n itself, and so is already (functioning as) a general term. If so, we don't need the doctrine that meaning of a kind term is the exclusion of the other to get general terms without universals.

But does the paradigm image manoeuvre work? An obvious objection to it is that it simply assumes that images can possess, unproblematically, the kind of generality of application which the nominalist agrees he needs to argue words can have. But why suppose there is any less of a problem in seeing how an image can fit many things without their having some shared quality or qualities, than there is in seeing how words can apply to many things without shared qualities? Doesn't the manoeuvre simply pass the buck from words to images?

As I read him, Siderits confronts essentially this problem in the central pages of his paper, and proposes an ingenious solution to it. Here is his formulation of the problem:

How does one learn to form [the paradigm image]? ... What we want to know is how one can learn to see ostriches but not other crows as distinct from a particular crow if the distinctness in question is not qualitative distinctness ... How can one learn to conform one's linguistic behaviour to that of others if there are no objective features to form the basis of this discriminative capacity? ("Buddhist Nominalism", 96)

And here is his solution:

...the appeal to social convention does bring out the point that the ability in question depends on practices that are responsive to human wants and interests. ... For it turns out that there is something all the non- s_n may be said to share, when 'non- s_n ' is formed in accordance with the relevant convention, namely that they all fail to satisfy a certain desire, say the desire to eat crow. Obtaining an ostrich or tumbleweed will fail to satisfy that desire; the desire persists in their presence. Obtaining s_n satisfies the desire. But the same hold for s_{n+1} , s_{n+2} , and s_{n+3} , so these also belong to the class of things correctly asserted to be not-non- s_n . This is why all four things are said to be crows. (ibid)

So, the ingenious move, as I understand, is this: to the realist's insistence that there has to be a shared characteristic, the nominalist can, and should, respond: "There is, but what is shared in not an objective feature, but a common relation to us. There are no objective common features—so no universals—but that doesn't mean there can't be non-objective common features, grounded in the relations between classes of particulars and their different relations to our wants and interests. These non-objective common features are acceptable to the nominalist, and they suffice to account for our use of general terms."

Siderits goes on to offer an equally ingenious defence of this move against the objection that different particulars could not all equally serve to satisfy a certain desire if they themselves did not have a shared nature—in essence, the reply is that the capacity of different particulars to satisfy some single desire no more requires that they have some common nature than does the capacity of different drugs to relieve fever. In general, different things may play the same causal rôle without there being something in common between them all in virtue of which they do so.

There is room, I suspect, for further rounds in this debate. It might be claimed, for instance, that while a functional property may well supervene on different base properties, it is in virtue of particulars' possession of general properties that the functional property supervenes—it is at best unclear how different particulars could all serve in the same rôle if their capacity to do so were not underpinned by some general properties, albeit not necessarily the same ones in all cases. But I shan't pursue this further now. Instead I want to comment briefly on some other aspects of the proposed solution, and to air a different cause for concern about appealing to the capacity of different particulars to satisfy a single desire, such as the desire for crow pie (or, generalising, to stand in some single relation to some want or interest).

First, then, it seems to me that although the play with the capacity of different particulars—e.g. different non-crows—all to fail to satisfy a certain desire is introduced to answer a question about how we form suitable paradigm images, in fact the images just drop out of sight. They play no discernible rôle—and certainly no essential one—in explaining how we achieve conformity in linguistic behaviour, or how we come to apply a single word to many different particulars. Once again, if the explanation works, it seems that it can work perfectly well without any introduction of intermediary images—so they may as well be ditched.

Second, although Siderits puts his proposal in terms of a non-objective shared feature common to all the objects which are non- s_n (non-crows)—viz. they all fail to satisfy the desire for crow—thereby leaving work to be done by equating the positive

with the doubly negative, it again appears that, if the proposal can work at all, it can be implemented directly, appealing to a different non-objective similarity among all the particulars which are s_n (i.e. the ones to which we apply ‘crow’)—viz. that they all satisfy the desire for crow. So once again, the distinctive doctrine of apoha semantics—that a kind term K means not non-K—is doing no essential work.

But does the proposed solution work? My final worry about it concerns the legitimacy of appealing to desires such as the desire to eat crow at this fundamental level of semantic theorizing. Roughly, my fear is that appealing to any such general desire—that is, desire with a general content which has to be specified in general terms, if it is to be satisfiable by many different particulars—simply re-introduces our old worry about circularity. We can all agree that the nominalist must aspire, in his semantics, to explain how we can meaningfully ascribe to ourselves desires and other attitudes which are general in the sense that their objects are not particulars—they are desires, not for this or that particular, but for something of a certain general kind. For example, I want a pint of beer, but there is no particular pint of beer of which it can correctly be said that it is the pint of beer I want—I just want it to be the case that I have a pint of beer, and I don’t mind which (so long as it tastes good, etc.). My desire for a pint of beer is like Quine’s desire for a sloop—no particular sloop, just relief from slooplessness^{vi}. And the nominalist must hope to be able to explain this in a way that avoids commitment to universals. I myself see no reason why he shouldn’t do that in just the way Quine advocated—by construing my desire for a sloop, when it is the desire for relief from slooplessness, rather than some covetous desire for my neighbour’s sloop, as a desire to the effect that a certain existential proposition be true: namely, that $\exists x(x \text{ is a sloop} \wedge \text{I own } x)$. But this solution to that problem highlights precisely my cause for concern over appealing to general desires to explain how different particulars can all fall under some general term—the explanation of how desires can, in the relevant way, be general presupposes that the terms involved in specifying their objects are already general, in the sense of being potentially applicable to many distinct particulars. To put the

point another way: a nominalistic compositional semantic theory, for English say, is going to generate a partial ordering of the sentences of English—an ordering which at least roughly reflects the order in which understanding of the expressions has to be acquired. In that ordering, one would expect sentences like ‘This is a sloop and I own it’ and ‘This is crow pie and I am eating it’ to come well before sentences like ‘I want a sloop’ (with the relief from slooplessness meaning) and ‘I want some crow pie’.

It may be said that the proposed solution can agree that the understanding of sentences ascribing desires cannot precede that of sentences describing what is, in effect, the satisfaction of those desires—but claim that this does not make it illegitimate to appeal to facts about what will or won’t satisfy a given desire in explaining how words get to apply to many things. This might be right, if by what satisfies my desire for crow pie, say, one simply means what will bring it about that I no longer want crow pie. But what satisfies my desire for crow pie—in this sense of satisfy—may very well not have parts of something which once had wings and black feathers in it. It may be a glimpse of a far more appetising roast pheasant, or the news that I’ve won the lottery, or a bang on the head^{vii}. So appealing to what satisfies or fails to satisfy desires for X is liable to give X the wrong extension. But if by what will satisfy a desire, one means what it is a desire for, we are back with the preceding worry.

Perhaps it can be argued that the circularity I have argued afflicts the appeal to desires (and other attitudes) is not vicious—but I have to confess that at this stage, I don’t see how.

6. CONCLUDING REMARKS

I have raised two main kinds of doubt about the doctrine of apoha and the associated defence of nominalism. First, I have suggested that the doctrine of apoha by itself does not provide a means of avoiding commitment to universals, and that when it is supplemented with other ideas in the way Siderits describes, it is the other

ideas that do the work, and that nothing essential is contributed by the apoha doctrine of exclusion of the other^{viii}. But I have also expressed some scepticism about whether the other ideas can anyway really give the nominalist a way of explaining how we can talk in general terms without commitment to universals.

I should emphasise that the main difficulties I have raised for the apoha doctrine have been developed on the assumption that it will involve explaining the meaning of general terms by the way of 'exclusion of the other'. In particular, it is this which gives rise to the difficulty over compositionality which I have been mainly concerned to press. So far, it remains possible that a defence of nominalist which preserves something of the spirit of the apoha doctrine might be mounted adopting one or other of the two more radical responses to the problem of compositionality which I mentioned earlier in the paper. Roughly put, the idea would be that instead of claiming that one can give a semantics for general terms under which 'x is K' is explained as meaning 'x is not non-K', one might claim that negation can be viewed as a device which cancels existential commitments carried by expressions occurring in what is negated. Perhaps one could then argue that 'x is not non-K' is precisely not equivalent to 'x is K'. Rather, it is a weaker—ontologically less committed—statement, which the nominalist may offer as a replacement for the more problematic 'x is K'. This would involve granting that general positive predications of the form 'x is K' are fully intelligible without benefit of nominalistic re-interpretation, but claiming that this leaves the nominalist free to introduce 'x is not non-K' as a replacement which avoids the unwanted commitment to universals associated with what it replaces. Whether a stable version of nominalism can be developed along these lines is a good question, but it not one I can discuss here^x.

ⁱ This paper is based upon a talk given at the conference Apoḥa Semantics and Human Cognition held in Cret Berard, Switzerland, 24-28 May 2006. In sharp contrast with all the other speakers at the conference, I have no expertise in this area, and am certainly not competent in the exegesis of the Buddhist Nominalists whose ideas and doctrines were the subject of the conference. I have, therefore, relied heavily upon my fellow speakers' contributions, and am greatly indebted to them for helping me to acquire some understanding of the intriguing, but perplexing, ideas I am trying to talk about in my paper. I need hardly say—but I will—that none of them can reasonably be held responsible for any misunderstandings and errors of interpretation which may remain.

ⁱⁱ I have drawn especially upon Mark Siderits's paper "Buddhist Nominalism"

ⁱⁱⁱ I think what Pascale Hugon, in her conference paper, calls the charge of circularity/interdependence can be seen as an alternative formulation of the objection I am considering here. She considers, and I think means to endorse, the response she takes to have been Dharmakīrti's—that the realist is equally open to the same charge. But I do not myself think that this ad hominem argument against the realist works. The realist can agree that when he fixes the convention for the application of the term 'tree', for example, he also settles what is not to count as a tree. But this does not mean that his theory makes the understanding of 'tree' depend upon understanding 'not non-tree'. I think formulating the problem for the nominalist in my

way helps to make it clear that it is a problem that distinctively afflicts the apoha doctrine.

^{iv} In fact, as we shall see, a good deal of what Mark Siderits says in “Buddhist Nominalism” fits in very well with such an approach, even though he does not explicitly address questions about compositionality.

^v Obviously, they may be said to have it in common that they are all distinct from Socrates, but there is nothing to unsettle the nominalist in that. The essential point is that there is, or at least need be, no characteristic they share, in virtue of which they are each distinct from Socrates.

^{vi} Cf. Quine “Quantifiers and Propositional Attitudes”

^{vii} This is the main problem with the theory of desire advanced by Bertrand Russell in Lecture III of The Analysis of Mind. This problem with Russell’s theory—which identifies whatever it is that causes the feeling of discomfort which he takes to be involved in any desire to go away as what the desire is for—is discussed at some length by Anthony Kenny in Action Emotion and Will, 101-110.

^{viii} In this respect, what I have said is very much in line with a number of remarks in Tom Tillemans’ paper.

^{ix} I should like to record my gratitude to the other participants in the Cret Berard conference both for their helpful discussion and for their patience with my amateurish efforts in the area of their expertise. I am especially grateful to the organisers, and to

Professor Shoryu Skatsura for helping me towards a better grasp of the ideas I have tried to discuss in this paper.