Should we be Pluralists about Truth?

Max Kölbel, final draft 4/3/2010

Pluralism about truth can take several forms: first, it might be the claim that the truth predicate expresses several truth concepts, and secondly, it might be the claim that the truth predicate, even though it expresses a single concept, corresponds to several truth properties, i.e. that truth is realized by several distinct properties. Recently, following suggestions made by Wright in his 1992, a number of theorists have defended pluralist theses of the second form. They claim that while there is a single generic concept of truth with application in all areas of discourse, truth is nevertheless realized by different properties in different areas of discourse. In this paper, I shall make some general observations about the way in which a pluralism of the second form can be motivated. Then I shall put forward some considerations in favour of the first sort of pluralism. Finally, I shall argue that pluralists of the second type cannot, without further ado, make use of standard frameworks of structured propositions, but rather need to re-think what type of entity, in their view, serves as truth-bearer.

1. Methodological Considerations: Expressions, Concepts, Properties

Examining Concepts

There are many concepts. Concepts are abstract objects, individuated in terms of the rules that govern them. Some of the many concepts are employed by us in thought, some of them are expressed in language, some of them are not employed in thought or not expressed in language. For each concept exists independently of there being any concrete episodes of thought and speech which could profitably described as an employment of it. Concepts are in principle susceptible to a priori investigation. We can fix on a particular concept by laying down, in stipulative manner, by which conceptual rules the concept is to be governed. Alternatively, we can fix on a concept by describing it as the concept that is expressed by this or that expression in a particular public language, or in a particular idiolect. Or we can fix on a concept by describing it as the concept employed by this or that person or persons in certain episodes of thought. Once we have fixed on a concept in

one of these ways, we can start examining it's properties. Some of these properties will be accessible to a priori methods, some may not be.[1]

Suppose we have fixed upon a concept by specifying a complete set of rules that govern it. Then we can directly move on to examine the a priori consequences of these rules. Such an examination may, for example, result in the finding that the concept is subsumed by another concept or that it subsumes another concept, if we have appropriate a priori information about these other concepts. It may result in the finding that the concept is definable in terms of certain other concepts (if the identifying set of rules didn't already come in the form of that definition). Alternatively, it may result in the finding that the concept is contradictory, or that it fails uniquely to determine an extension. Even if we have fixed upon a concept by laying down the rules that govern it, the concept may have some properties that are susceptible only to empirical methods. For example, the question whether a concept applies to anything will in many cases be an empirical question. If we have a sufficiently clear conception of a natural property or kind, we can also examine whether a concept's extension coincides with that of some natural property or kind, or maybe whether it does so necessarily.

Suppose we fix upon a concept by describing it as the concept expressed by a given expression in some public language. Then the question by which rules that concept is governed may be answerable only by recourse to empirical or quasi-empirical investigation. Thus, if we examine a concept identified as the concept expressed by some expression in a language with which we are not competent, we will need to employ empirical methods to discover any conceptual rules governing that concept. This may involve observing the use competent speakers make of the expression in question. If we examine a concept identified as the concept expressed by some expression in a public

---

[1] The approach to concepts described in this paragraph and assumed below is, of course, controversial. There are those, like Fodor (1975), Dretske (1981) or Laurence & Margolis (1999), who construe concepts as concrete psychological entities such as mental processes or mental symbols. This approach is often taken for granted in cognitive science. There are also those, who, like Dummett 1993, construe concepts as abilities. The current conception of concepts as abstract objects goes back to Frege's anti-psychologistic theory of "senses" (e.g. 1892), and is defended, for example, by Peacocke 1992, 2008 and Zalta 2001. I do not regard it as a foregone conclusion that the approach pursued in cognitive science is inconsistent, or even in genuine rivalry, with Fregean theories. For a defence of the role of abstract objects in theorising about thought and speech, with which I sympathize, see Matthews 1994. Part of what I am assuming here is that concepts are governed and individuated by norms. This assumption is supposed to be neutral between a number of different construals of such norms.

language with which we are ourselves competent then we can employ quasi-empirical methods to find out about the rules governing the concept. This may involve relying on our own intuitions as competent users of the language in question. But even in this case we can make use of straightforwardly empirical methods by observing the use we and others make of the expression. Once we have an idea of the conceptual rules governing the concept empirically identified, the situation is as before with concepts stipulatively identified: we can again investigate those of the properties of the concept that are consequences of these conceptual rules using a priori methods, and we may investigate other properties using empirical methods.

Suppose, finally, that we are examining concepts that are identified as the concepts that this or that thinker, or group of thinkers, is employing in this or that situation. Again, it is in principle an empirical or a quasi-empirical matter what the conceptual rules are that govern the concept. If the concept is identified as the concept the researcher him or herself is employing in this or that situation, quasi-empirical methods involving memory and introspection can be used. If the concept is identified as the concept employed by this or that person distinct from the researcher, then the researcher has to resort to empirical methods, such as the observation of verbal or other behaviour in order to find out what the rules are that govern the concept. We can distinguish the case where the concept is identified as the concept employed by a group from the case where the concept is identified by the role it plays in the thought of an individual. Once the conceptual rules governing the concept are found, the situation is again the same as in the stipulative case.

Examining Properties, Kinds or Objects Denoted by Concepts:

Sometimes philosophers separate two questions, the question of whether a given concept applies to anything and the question of whether what the concept applies to a (natural, real etc) property, a (natural, real etc) kind, or an (natural, real etc) object. This sort of question makes sense, I believe, only on a certain kind of background of assumptions and theoretical or explanatory interests. Let me explain.

Suppose we have an irreducible concept, i.e. a concept that is not definable in terms of other concepts. Perhaps the conceptual norms involve certain perceptual criteria for the application of the concept as input conditions. Then the question of whether the concept

has application at all will not necessarily require the employment of other concepts. Suppose that we have a concept that is definable in terms of other concepts which can be independently employed, such as the concept of a meat-eating plant. Then we will employ the concepts *meat-eating* and *plant* in adjudicating whether the concept has application, and in which instances it can be correctly applied.

The question whether a given concept corresponds to a property, a kind, or an object (I omit the qualifiers "natural", "real", etc for ease of expression) is usually different. In this case, we have a concept $c1$, from a system of concepts C, which may be either irreducible or definable in terms of other concepts $c2$, $c3$, … from C, and we ask whether the extension of the concept constitutes or, corresponds to, a property, a kind or an object. In order meaningfully to raise this question, we need to have some interesting separate system of concepts $N = \{n1, n2, … \}$, in relation to which we are asking the question. The question is then whether $c1$ is coextensional (or perhaps necessarily coextensional) with some concept $n \in N$.

N may be interesting for various reasons. For example, we might believe that the concepts in N are epistemologically privileged in some sense, perhaps because the methods used in employing these concepts are particularly well-understood or reliable. In other cases, we might believe that the concepts in N are somehow ontologically or metaphysically privileged, perhaps in the sense that we accept the doctrine that there exist only properties, kinds or objects that constitute the extension of some concept in N. Or perhaps N is ontologically privileged in the sense that the facts describable in terms of the concepts in N are sufficient to determine all the facts (supervenience). But no such extravagant significance need be attributed to N. Even someone who regards N as neither epistemologically nor ontologically privileged may still have a certain interest in finding out about the interrelation between the concepts in C and the concepts in N. Thus, she may for example be interested merely in establishing whether a reduction of $c1$ in terms of the concepts in N is possible. Unlike a potential analytic reduction of $c1$ in terms of $c2$, $c3$, … , a reduction of $c1$ in terms of $n1$, $n2$, … $\in N$ might be an empirical (or at any rate non-analytic or otherwise non-obvious) reduction which adds to our genuine knowledge. For example, a reduction of the concept of *heat* to concepts from thermodynamics (heat =

mean molecular energy) or of the concept of *water* to concepts of elementary chemistry (water = H2O) represent genuine gains in knowledge. Such findings can be valuable even in the absence of any general ontological or epistemological priority of N. It may, for example, help us develop methodological shortcuts, or just satisfy our curiosity as to how things hang together.

One more specific concern that may spurn our interest in a possible reduction of c_1 to some $n \in N$ is explanation. This may be related to the case where N is epistemologically privileged. Suppose we find that c_1 is coextensional (or necessarily coextensional) with some $n \in N$ (or perhaps merely that the extension of some $n \in N$ is contained in the extension of c_1). Suppose further that we have some theory or set of general principles that explain/predict why things fall under n. Then this will provide an explanation for, or explanatory understanding of, why things fall under c_1.

To summarize: it is important to keep in mind that any investigation as to which property, kind or object, if any, a given concept corresponds to makes sense only on the background of some interesting set of reductive concepts which is used to individuate properties, kinds or objects. In any such investigation it will be useful to be clear about what the reductive set is, and what kind of interest (epistemological, metaphysical, explanatory) we are pursuing in exploring the question.

The Conceptual Mess We Are Often In:

Usually, the issues are not as neatly separated as the last pages suggest. Very often, it is hard to make out whether in a particular philosophical debate about some issue, we are starting by identifying concepts stipulatively, or whether we are identifying them in one of the empirical ways described above. Another difficulty is that the terrain often moves as we are exploring it. Suppose, for example, we are examining the concept expressed at some time t by an expression e of some public language. Empirical and quasi-empirical research suggests that at t, e expresses the concept specified by conceptual rules R. Further a priori research shows that that R are inconsistent. This result may now influence the way e is used subsequently; perhaps some aspects of R are now branded as mistaken. Thus after t, R no longer governs the concept expressed by e.

Sometimes in philosophical debates everyone assumes that everyone is talking about the same concept or property, but some of the disagreements suggest that the disputing parties are pursuing different projects.[2] Thus, the different parties may all be employing the expression *e* when identifying the concept or property which is the assumed subject of the dispute. Their dispute presupposes that *e* is used with the same sense by everyone. However, one party may be basing their considerations on one set of a priori conceptual rules supposedly governing the concept expressed by *e*, while another party is basing their consideration on another set. It is now unclear whether we should say one of at least two things: either that each party is making stipulative, or quasi-empirically correct assumptions about the conceptual rules governing the concept they each express by *e*, but they are simply wrong to assume that both are using *e* with the same sense (i.e. to express the same concept). Or we say that given they are using *e* to express the same concept (since they take themselves to be disagreeing), one of them must be making mistaken claims about the conceptual rules governing that concept.

Let me illustrate some of the difficulties with an example. There is a philosophical debate about personhood. Locke, in the *Essay* seems to start his examination from a certain view of what a person is. He says that a person is "a thinking intelligent being, that has reason and reflection, and can consider itself as itself, the same thinking thing, in different times and places" (*Essay Concerning Human Understanding*, II.27.9). He distinguishes persons from "men", i.e. human beings, and "thinking substances". Thus Locke seems to start directly from certain assumptions about the concept of a person, and he uses these assumptions to generate certain conclusions, for example about personal identity. Others, for example Derek Parfit (1984), use quasi-empirical evidence to establish what identity criterion the concept of a person involves. Thus he will use imaginary situations and our responses to them in order to tease out certain features of the concept of a person. Locke is not bothered by certain counterintuitive consequences of his definition of a person. For example, it seems to be a consequence of Locke's view that if a sober man doesn't have appropriate memory links or no psychological continuity with the same man at an earlier time, when he was drunk, then the sober man is not the same person as the drunk man, even though they are the same man. Locke accepts this

---

[2] Kirkham argues that this has happened in the case of debates about truth, see Kirkham 1992/2001.

consequence. Why would we then punish the sober man for what the drunk man did, if they are different people? Locke bites the bullet and says that this is for pragmatic reasons only: it would be too easy to fake the condition of drunkenness (*Essay*, II.27.22). By contrast Parfit, on his approach, has to take seriously our intuition that the drunk man and the sober man are the same person.

Now, should we say that Locke and Parfit are simply talking past one another? While Locke is talking about a concept of person as stipulatively defined, Parfit is talking about a concept of person that is expressed by the public language expression "person", and which is employed by thinkers in certain contexts (otherwise his quasi-empirical methods would make no sense). There are reasons why this would not be a perfectly adequate portrayal of the situation. First, both philosophers are usually taken to be addressing the same issue. To proclaim that they are strictly adressing different questions would be to make a nonsense of much philosophical debate. Secondly, what concept we employ under the label "person" is not independent of considerations, such as Locke's and Parfit's. For whatever consequences any of them derives may be taken as a motivation for revising our thinking and reasoning habits, or our use of language. When we do change the way we use the expression "person" as a result of philosophical reflection, are we changing the meaning of the expression? Are we adopting a new concept of a person, or are we merely revising a mistaken view of what that concept is and by what rules it is governed? Hard to say.

When we discover that certain a priori principles we employ are incoherent, are we discovering that some of those principles were after all erroneous? Or are we discovering that the concept we used is incoherent, and that we should make a fresh start by adopting a new system of concepts? I am sure there are several different ways in which we can make sense of episodes like the philosophical debate about personhood. In the present context, I merely want to illustrate the mess we are in as philosophers.

Yet another dimension of confusion is added once we start considering whether our concepts correspond to any natural, real etc kinds or properties. Suppose our theory of the structure of human societies, or perhaps a theory of practical rationality, suggests that there are certain natural or social kinds. Then we might insist that our concept of a person

must be construed in such a way that it can be taken to correspond to one of those kinds (perhaps we think that this is one of the conceptual rules governing the concept). In that case, we might regard certain a priori assumptions about the concept of a person, or certain habits of language use (or concept employment) as mistaken.

A well-known example is that of our concept of jade. The concept *jade* was well-entrenched when, some time in the 19th century, mineralogists discovered that there were two chemical (or mineralogical) kinds all and only instances of either of which fell under the concept *jade*: jadeite and nephrite. Did this show that our concept of *jade*, as previously employed, was somehow flawed? Not necessarily. Let's distinguish two cases: the concept *jade* might have been governed by the conceptual rule that its extension is to correspond to a unique mineralogically uniform kind, or it might not. In the former case, the discovery of jadeite and nephrite would seem to have shown that we needed to abandon the concept as employed until then, because the discovery showed, precisely, that the original concept did not have application: nothing is both a mineralogically uniform kind and also answers to the criteria associated with the original concept. A revision would then have been called for: adopt a new concept which either involves modified criteria or is not a concept of a mineralogically uniform kind. In the latter case, the discovery is just the empirical discovery that exactly two discreet mineralogical kinds constitute the extension of the concept jade, and no revisions are called for.

The concept we were employing prior to the discovery seems to have been of the sort that does not aspire to reflect mineralogical kinds, or at least to have been retrospectively reinterpreted in this way. For the concept *jade*, including both jadeite and nephrite, continues in use. Under different historical conditions, we may have responded differently, for example by deciding that jadeite, but not nephrite was the real jade, and that any conceptual rules that allowed nephrite to count as jade were mistaken. But as things are, we became *pluralists* (or more precisely: dualists) about jade: there are more than one (exactly two) different mineralogical kinds which count as jade. Our pluralism about jade is relative to mineralogical kinds. This is not pluralism about the concept *jade*, for there continues to be just one concept, one that applies to exactly two all and only the members of two different mineralogical kinds.

The concept of fish, as employed prior to modern biological classification, provides a contrasting example. This concept did not line up with a classification of kinds according to modern biology. As a result, we modified our conceptual framework. This can be characterized as the abandonment of the old concept *fish* in favour of a new concept. Alternatively, it can be characterized as a case of coming to recognize that certain conceptual norms that had been assumed to govern the concept *fish* were recognized as mistaken. The philosophical discussion of mental concepts such as *belief* or *pain* seems to constitute yet another type of case. The presumption that there is no reduction of mental concepts to physiological or other material concepts has led eliminativists to deny the existence of mental states. Others draw the conclusion that in the absence of reduction, there is at least supervenience. Yet others deny even supervenience and accept dualism. Neither of these positions corresponds to the actual point of view regarding jade.

As I said above, the only way in which a question can arise as to whether, which and how many properties or kinds a concept corresponds to is on the background of some alternative system of concepts that has some special significance. The answer to that question can range from empirical reduction (to one or more kinds) over supervenience to a claim of complete unrelatedness.

To summarize: we can examine concepts in a purely a priori manner by examining concepts identified in a stipulative manner. When we do this, we run the risk of examining concepts that no-one ever employs or concepts that are not expressed by any expression. Thus, when examining concepts such as the concept of a person, or that of truth, we may need, in order to avoid irrelevance, to pay attention to actual conceptual habits and language use. However, even clarity in principle about these issues does not always make things easier, for often it is very difficult to separate purely a priori considerations about concepts in the abstract from empirical issues concerning actual language use or actual concept employment. Moreover, sometimes we ask questions, about how a given concept interacts empirically or non-analytically with certain other concepts. We might assume that it must denote a natural, real, etc property or kind, or perhaps several natural properties or kinds. In this case we are basically expressing certain, possibly empirical, hypotheses or discoveries about how the concept we have

identified interacts with certain other concepts we have in our repertoire, and to which we attach special significance.

2. More Than One Concept of Truth

In the case of truth, the conceptual mess we confront is considerable. The philosophical debate about truth is so old and so extensive that it would probably be impossible to separate pre-theoretical from philosophically informed intuitions or judgements. We can start by postulating analytic principles concerning truth and then explore what follows from them. This is what some philosophers do. But of course if one philosopher postulates, say, a disquotational or equivalence schema as the basic conceptual rule concerning the concept of truth, and another philosopher starts from a different set of principles—say the idea that truth is what enquiry aims at, or that truth is correspondence with objective reality—, then who is to decide which of the two, if any, has started from the correct set of principles? There can only be genuine philosophical debate about truth where there is some common ground concerning the concept of truth.

I believe that one good way to avoid irrelevance or speaking cross-purposes and to make sure that we know what we are talking about is to start by identifying the object of enquiry as the concept expressed by the predicate "is true", a concept that plays a certain role in actual reasoning.[3] Restricting ourselves to the English language would seem to be exaggerated, and I suspect that our findings concerning "is true" and the concept thereby expressed will generalize to translations of the English predicate into other languages. However, we should in principle remain alert to the possibility that other languages may lack a predicate that works exactly like "is true" in English.

We should carefully distinguish, first of all, the predicate "is true" as used in ordinary discourse and the predicate "is true" as used in semantic theorizing. In semantics, "is true" usually figures as a predicate that applies to sentences, more specifically to ordered pairs of sentences and contexts, or to ordered pairs of sentences,

---

[3] Prompted by editorial comment, let me stress that I do not mean to denigrate other types of enquiry, such as the examination of broadly truth-related concepts expressed or employed elsewhere. All I am saying is that *one* way of avoiding the irrelevance of addressing concepts that no-one uses, and of speaking cross-purpose, is explicitly to address one of the several legitimate questions one might raise in the area, namely: what kind of concept (or concepts) is expressed by a certain range of uses of the English "is true".

contexts and circumstances of evaluation—often the latter are conceived of simply as possible worlds (see, for instance, Lewis 1970, 1980, Kaplan 1977). Not so in ordinary discourse, where "is true" seems to be a one-place predicate that applies to what people say or think, not to sentences. Even those who, like Davidson (e.g. 1967, 1990), believe that semantic theorizing relies on our pre-theoretical understanding of a primitive truth notion will have to admit that no-one has ever pre-theoretically thought of applying "is true" to, for example, sentence-context pairs or to sentences as uttered by so-and-so at such-and-such a time. At the very least the ordinary concept of truth undergoes some modification before it can figure in semantic theories in the way it does. There may well be some bridge principle that analytically links the semantic truth concept with our ordinary concept, and that therefore our understanding of the semantic truth concept does rely, as Davidson and others claim, on our understanding of ordinary truth. But it seems to me that we cannot avoid distinguishing the concept of truth employed in semantics from any ordinary concept, expressed in ordinary discourse by the predicate "is true".[4]

Let us focus now on the concept expressed by the predicate "is true" in ordinary non-theoretical discourse. Let us focus only on those uses where "true" is applied to what people say or think, i.e. the contents of thought and speech ("true" as in "true friend" or "true wheel" seem to me to express different concepts that we are not currently interested in). I have argued elsewhere that empirical and quasi-empirical considerations support the view that in fact the ordinary truth predicate expresses at least two different concepts and is therefore ambiguous (Kölbel 2008b). I shall here provide a brief summary of the view and its motivation.

We find two tendencies in ordinary usage of the truth predicate. On the one hand, we are quite happy to apply the truth predicate in any topic area. Whatever someone says, supposes, believes or suspects, it is the sort of thing that we are liable to predicate "true" of. Suppose someone uses an assertoric sentence, *any* assertoric sentence, to say something. We will always be prepared to say things like: "That's true.", "That's not true." or "I don't know whether that's true.". It seems also correct to say that whenever

---

[4] I have argued, in Kölbel 2001, that Davidson's claim that semantics relies on an understanding of the notion of truth is unfounded. See also Ludwig 2002 and Badici & Ludwig 2007. In Kölbel 2008a, I have also examined, in more detail than can be provided here, how the semantic truth predicate and the ordinary truth predicate are related.

someone has said that *p,* and someone else replies in one of the three ways just mentioned, then it is very hard to see any difference over and above style between these three replies and the following truth-free replies: "*p*.", "not-*p*." and "I don't know whether *p*.".[5] The two forms seem to have exactly the same consequences, commit the utterer to exactly the same things, etc.

On the other hand, it is quite common for people to be more selective in their choice of candidates for truth-ascription. Thus, they may refuse to apply "true" in certain topic areas, notably in the evaluative realm. They may well say things like: "Statements about matters of taste can't be true or false.". People quite commonly and pre-theoretically associate a connotation of objectivity with the truth-predicate, and in so far as they believe that in a certain area our beliefs do not admit of objective correctness they would deny that claims, beliefs, judgements etc in this area are evaluable in terms of truth.[6]

The two different tendencies are, I believe, accessible to quasi-empirical evidence: competent users will, on the whole, agree that this usage is within the range of competent use. But there are also properly empirical data to back it up. On a one-page questionnaire, bona fide competent users repeatedly called the claim that Ali G is funny true (or false), but also denied that judgements on what is funny can be true or false.[7] There may be disagreement as to the exact status of this usage—for example to what extent it reflects literal meaning. However, even if we came to be convinced that one of the two tendencies reflected a non-literal use of language to convey something pragmatically: there would still be two different concepts employed here.

The conclusion I draw from these observations is that "is true" as employed in ordinary non-philosophical contexts, is ambiguous, or at the very least is used to express

---

[5] One needs to be careful in the exact articulation of this principle: If someone says that she is hungry, and someone else answers "That's true.", the answer obviously not equivalent to the answer that would have resulted from uttering "She is hungry.".
[6] One may suspect that the second tendency is contaminated by the influence of philosophical theorizing. For example Ayer and the logical positivists might be perceived to have brought about this tendency. However, I do not believe that it is useful to try to separate philosophical from non-philosophical uses here. The first tendency could with equal right be suspected of being the result of philosophical influence, for wasn't it Aristotle who claimed that to say of what is that it is, and of what is not that it is not, is true? That seems to be a principle which applies across the board.
[7] The set-up is described in somewhat more detail in Kölbel 2008b.

two different concepts.[8] Let us call the concept with the wide range of application "truthD" or "deflationary truth", because it can easily be interpreted as a deflationary truth concept. And let us call the other concept, associated with the more restricted range of application "truthS" or "substantial truth". The concept *truthD* can easily be interpreted as a concept whose primary conceptual rule is that for any thought or speech content, i.e. proposition, p, the proposition that p is true is equivalent to the proposition that p.[9] Such a concept is a useful concept to have, for it allows us, for example, to think propositions such as the proposition that the first claim of section 3 is true, even when we do not know what that claim is. Similarly, it is useful to have a predicate that expresses that concept. In fact, if we didn't already have a concept like this, and a predicate that expresses it, it would be high time to introduce and start using them.

TruthS, the concept in play when we are reluctant to apply "true" across the full range of propositions, seems to be governed at least by some additional conceptual rules. It must be governed by an extra constraint of objectivity. Minimally, this might be the principle that whenever it is correct for anyone to apply *truthS* to some proposition, then it is a mistake for everyone to deny the concept of that proposition. This would explain why some people are reluctant to apply *truthS* to some evaluative propositions.[10] Now, it seems that even *truthS* is subject to the equivalence principle within its more restricted range of application. Competent users of *truthS* will infer the proposition that p from the proposition that the proposition p is trueS. And in the range of application of *truthS*, competent users will infer the proposition that the proposition p is trueS from the proposition p.

This opens up the path for regarding the concept *truthS* as subsumed under the concept *truthD*, or in other words, that truthS is a special case of truthD. I have suggested moreover, that truthS can be defined in terms of the concept *truthD* and some notion of objectivity:

---

[8] As I explain below, I believe this ambiguity not to be of the accidental kind exhibited, for example, by "bank", "bill" or "premises", but rather I believe the distinct concepts expressed to be systematically related.

[9] This conceptual rule will generate contradictions if we allow speech or thought contents such as the content that this very thought content is not true. Discovery of this problem will lead us to restrict the range of applicability of the rule, or even to abandon the concept and replace it by a new one.

[10] Compare Richard 1997.

(D)   For all propositions p, p is trueS iff p is objective and p is trueD.

The thesis that "is true" expresses two different concepts, *truthD* and *truthS*, in combination with the reductive definition (D), raises the question of why we are not aware of this ambiguity in our use of "is true"—why, for example, we often do not feel that understanding an utterance involving "true" requires disambiguation, or why the term just doesn't seem ambiguous the way, for example, "coach" seems ambiguous. However, there are examples of similar ambiguities. For example, most competent users would not at first have the impression that the nouns "dog" and "duck" are ambiguous.[11] However, when confronted with the following examples, they may change their minds:

(1a)   Dogs are not allowed in the playground.

(1b)   Mina is a bitch, not a dog.

(2a)   Ducks like old bread.

(2b)   Donald is a drake, not a duck.

It is clear that the intended meaning, and the favoured interpretation, of "dog" in (1b) and of "duck" in (2b) is *male dog* and *female duck* respectively. However, in (1a) and (2a), the intended meaning, and favoured interpretation, of the same words is, respectively, *dog of any sex* and *duck of any sex*. To be sure, we may argue about whether this is genuine lexical ambiguity, or whether the b-examples merely express the gender-specific concepts in a pragmatic way. But what is hard to dispute is that the words "dog" and "duck" are used to convey two distinct concepts in these examples. Distinct, but related: just as in my ambiguity thesis concerning "true", one of the concepts subsumes the other and one is definable in terms of the other in a way analogous to (D).

Thus, just as one does not notice the ambiguity of "dog" and "duck" until one considers examples like (1b) and (2b), one will not notice any ambiguity in "is true" until one considers examples like (3):

(3)   It's true that Ali G is funny, though, actually, it's not true, because judgements concerning matters of taste do not admit of truth or falsehood.[12]

---

[11] Compare Lewis 1989, p. 130.

[12] Or, to preserve the analogy with (1) and (2), consider:

    (3a) It's true that Ali G is funny if and only if Ali G is funny.

Those who regard (3) as incoherent will at least concede that one and the same person may, without change of mind and without irrationality or insincerity, utter "That's true." concerning the claim that Ali G is funny and also utter, perhaps in another context, "Claims regarding what is funny can't be true or false.". Once this is conceded, I believe, it is conceded that "true" is in ordinary discourse used to express two different concepts. My hypothesis is that these two concepts are *truthD* and *truthS*.

The advantage of the ambiguity thesis is that it allows us to make sense not only of ordinary thought and language, but also of the philosophical debate. Most philosophical debate about truth before the advent of deflationism seems to have concerned truthS. A full-fledged deflationist who claims that "true" only ever expresses truthD, will have to regard these debates as in some sense fundamentally misguided. However, acknowledging that "true" can be used to express these two distinct but related concepts allows one to make sense of the traditional debate while still honouring the insights of deflationism about the so-called "logical" point of a truth predicate and truth concept (see for example Horwich 1990/98 and Wright 1992).

Some have warned against the thesis that "is true" expresses distinct concepts in different topic areas on the basis that this would make it difficult to understand validity as truth-preservation, and to treat connectives like "and" or "or" as truth-functional. I believe that these worries have been satisfactorily addressed by Pedersen 2006. However, the current ambiguity thesis avoids these problems before they even arise: *truthD* applies across the full range of propositions, so validity and truth-functionality can simply be construed in terms of *truthD*, at least in so far as they need to be defined in pre-theoretical language. Formal treatments of truth-functionality and validity will presumably in any case be couched within a formal semantics, employing its own distinct semantic truth concept.

My observations concerning the use of the truth predicate (assuming they are correct) do not, of course, conclusively demonstrate my ambiguity thesis. Let me briefly

---

(3b) Ali G is funny, but it's not true that Ali G is funny (because that's not a matter of truth or falsehood)

It looks like both could be felicitously uttered, even by the same person. In that case, the best interpretation of (3a) will interpret "true" as expressing a different concept from the one it expresses on the best interpretation of "3b). Thanks to the editors for this suggestion.

mention two alternative strategies for dealing with the observations, and comment on one of them. One alternative, already touched upon, is to regard one of the two tendencies as mistaken, i.e. to condemn as misguided either the tendency to apply the truth predicate across the board or the tendency to restrict its range of application. I regard this approach as unsatisfactory because both tendencies are, in my view, well entrenched and also justified by the usefulness and fruitfulness of each of the two truth concepts *truthS* and *truthD*. Another alternative, one that I would like to comment upon briefly, is to adopt what has been called a "meaning-inconsistency approach" as it has been pursued in addressing liar and sorites paradoxes by, for example, Eklund 2002, 2007, Patterson 2007, Scharp 2007 as well as Badici & Ludwig 2007. According to this approach, there can be meaning constituting principles that are false. This approach is in principle also available in response to the inconsistencies I observed in the use of the truth predicate (which are entirely independent of the liar problematic). Accordingly, instead of being ambiguous, the truth predicate is governed by inconsistent meaning-constituting principles.

I have three reasons why I believe this approach is less well-suited to the current problematic. The first is that I do think that statements like (3) bring out a distinct feeling of disambiguation which is resembles cases like (1a)/(1b) and (2a)/(2b). There is no such resemblance in the case of the liar or sorites paradoxes.

Secondly, I regard the meaning-inconsistency approach as a kind of last-resort response, which may well be appropriate in the case of the liar and possibly even the sorites paradox. But since there is a fairly simple solution available in the form of a credible ambiguity thesis, underpinned by a reduction of the form of (D), there is no need to adopt a meaning-inconsistency approach.

Thirdly, it is not clear that a meaning inconsistency approach would avoid the conclusion that we are employing two distinct concepts of truth. For the meaning-inconsistency theorist needs to explain how we manage to use the predicate unproblematically despite the inconsistency in its meaning. This explanation will involve the thesis that we do not accept or follow the meaning constituting principles or inferences. If we want to take seriously the observation that both the across-the-range

uses and the restricted uses are well-entrenched and fruitful, then the explanation would amount to the claim that we follow the relevant meaning constituting principles selectively. That is: sometimes we follow the principles that accord with *truthS* and sometimes we follow the principles that accord with *truthD*. There are two possible interpretations of this: either we say that on all these occasions we are employing an inconsistent concept (constituted by the inconsistent principles), but avoid getting into a tangle by applying different principles case by case. This interpretation seems to me to be of dubious coherence, for what would justify the claim that it is the same inconsistent concept we are employing in different way on these different occasions? The second interpretation is that even though the meaning of "true" is inconsistently governed by conflicting principles, the thoughts we express by means of it always involve only one of the consistent concepts, i.e. either *truthS* or *truthD* (see Scharp 2007, who seems to argue for this). But this again seems difficult to accept, for what would justify the thesis that "true" has this uniform inconsistent meaning in all the contexts in question, rather than being consistently ambiguous?

3. More Than One Truth Property?

In recent years, a new brand of deflationism or minimalism about truth has become popular: pluralism about truth. Recent discussion of the view seems to take its departure from Wright 1992. It combines the claim that the ordinary truth-predicate expresses just one concept of minimal truth, with the claim that there are nevertheless several distinct truth *properties*. The idea is, roughly, that while *truth* is a concept that has application across the whole range of propositions[13], there are nevertheless different properties of propositions that explain why a proposition falls under the deflationary concept of truth. More specifically, it is often suggested that the variety of truth properties corresponds to a variety in "discourses" or domains of propositions, i.e. a variety of topic areas in which the concept of truth may be applied.

---

[13] Wright claims not only that all assertoric content is truth-apt, but he is also a minimalist about assertoric content in the sense that certain syntactic criteria, such as being capable of figuring as the antecedent of a conditional, suffice for assertoricity (1992, 36).

Thus, according to Michael Lynch, while the concept of truth expresses a single functional property, this property can be realized (2004), or manifested (2009) in a variety of different properties, such as the property of corresponding to reality and the property of being a member of a coherent system. Lynch compares his view with the functionalist view that the functional property of being in pain can be differently realized in different kinds of organism. According to Nicolaj Jang Pedersen (2006), the concept of truth may not even express any single property, but only a plurality of different properties in different domains of propositions. Gila Sher (2004) also argues that the concept of truth denotes a variety of different properties, though she argues that all these truth properties are correspondence properties.

Pluralism curiously reverses the direction of progress familiar from most of Plato's early dialogues. In these dialogues, Socrates' interlocutors often begin with a characterization of a plurality of kinds of virtue, of piety, of justice etc, and Socrates them asks them to provide a uniform account: What do all these things have in common? Contemporary pluralists about truth, by contrast, deem themselves already to be in possession of such a uniform account, and go on to look for a plurality. Thus, Meno, in Plato's *Meno* at one point proposes a form of pluralism about virtue. He says:

> First, if you want the virtue of a man, it is easy to say that a man's virtue consists of being able to manage public affairs and in so doing to benefit his friends and harm his enemies and to be careful that no harm comes to himself; if you want the virtue of a woman, it is not difficult to describe: she must manage the home well, preserve its possessions, and be submissive to her husband, the virtue of a child, whether male or female, is different again, and so is that of an elderly man, if you want that, or if you want that of a free man or a slave. And there are very many other virtues, so that one is not at a loss to say what virtue is. (*Meno* 71e)

This is, of course, the type of account Socrates will dismiss in his usual fashion. In principle, however, Meno is attempting to provide worthwhile information. Even if he already had a uniform account of virtue in general, it would still be an interesting project to try to specify for various different types of people what virtue consists in for each type. He might even try to derive this information from the uniform account of virtue. (The fact

that we are likely to disagree with the details of Meno's pluralism just goes to show that there can be non-trivial things a pluralist can say about different forms of virtue for different types of people.) I will therefore assume that in principle there is similar scope for an interesting pluralism about truth properties, even when one is already in possession of a uniform account of the concept of truth.

The four pluralists just mentioned seem to agree that there is a single unified truth concept which has application across all the domains, which is governed by principles such as compliance with the equivalence schema and perhaps normative principles such as the principle that one ought to believe or assert only truths. I have no quarrel with that (see §2 above). In addition, and this is what makes them pluralists, they claim that there is not just one, but several truth properties. As I pointed out in §1, the question of whether a concept denotes one or more properties (when this is understood to be separate from the question of whether the truth concept is instantiated at all) only makes sense against a background of some way of counting properties. Thus, in order to get a determinate question as to how many truth properties there are, we need a system of concepts N, which has some epistemological, ontological, explanatory or some other kind of interest for us, so that we can take the question to be the question whether the concept of truth is coextensional with some one concept $n \in N$, or whether its extension is perhaps the union of the extensions of several (non-coextensional) concepts in N. The schematic pluralist thesis is: The extension of the concept *truth* is the union of the (non-empty) extensions of several concepts n_1, n_2, ... $n_n \in N$, where N is an interesting class of concepts. The condition that N must be "interesting" is necessary because otherwise the pluralist thesis would be trivial. Take any concepts c with several instances: then there will always be an *uninteresting* set of concepts that subdivide the extension of n. Thus, in order to construe an interesting pluralism, we should identify an epistemological, ontological-reductive, explanatory, or whatever project that the pluralist is pursuing.

These constraints leave a lot of room for potentially interesting forms of pluralism about truth, and I am not in a position to be able to offer a comprehensive survey of all such forms. However, I would like to explore the possibilities for a certain restricted subclass of potential forms of pluralism about truth, one that is, I believe prominent in the

minds of many theorists. This is the subclass of pluralisms that construe the concept *truth* as a concept that applies to propositions and also construe propositions as structured propositions in the manner of, for example, Russell (19??). I shall argue that there is a certain tension between (a) the claim that there is an explanatorily more basic set of concepts *N*, such that the *truth* concept is nontrivially coextensional with the union of several of the concepts in *N* and (b) the claim that *truth* applies generally to structured propositions only. The lesson will be that presumably the pluralist about the truth of propositions will have to postulate a variety of propositions that corresponds to the variety of truth properties she postulates.

Let us remind ourselves that the pluralist holds that there is a single truth concept that is expressed by the truth predicate. This concept is such that it conforms to a series of platitudes, such as the platitude that the proposition that *p* is true just if *p*, and the platitude that believing or asserting a proposition is correct only if the proposition is true. Moreover, the pluralist is a minimalist about propositions: any declarative sentence expresses a proposition, i.e. there is no discourse the declarative sentences of which fail to express propositions (as held by classical expressivists). There are, the pluralist claims, several distinct properties (from an explanatorily interesting class) in virtue of which this concept of truth can apply to a given proposition, and these distinct properties correspond to different discourses, i.e. to different topic areas, such as mathematical discourse, discourse about medium sized objects, moral discourse, etc.

Suppose now, that such a pluralist also holds that the truth concept applies to propositions, and that she construes propositions as structured propositions. Structured propositions, let us say, are complex entities that are constituted by particulars and universals, where universals include properties and relations of varying adicity, both first-order and higher order. The construction of propositions from these constituents will be governed by certain rules, such as, for example: an *n*-ary first order universal *U*, applied to a sequence $<o_1, \dots o_n>$ of *n* particulars makes a proposition; an *n*-ary first order universal *U*, applied to a sequence $<o_1, \dots o_{n-m}>$ of *n-m* particulars makes an *m*-ary propositional function; an *n*-ary second-order universal *U*, applied to an *n*-ary

propositional function makes a proposition; an *n*-ary propositional connective *C* applied to a sequence <*p*₁, … *pₙ*> of *n* propositions makes a proposition; and so on.[14]

Now, there is one reductive explanation for which such theories of propositions were originally designed: a reduction of the truth of complex propositions to that of atomic propositions, and/or the reduction of truth in general to the notion of instantiation (as when a sequence of particulars or propositional functions instantiates a universal). Thus, the explanatory account explains what it is for a complex proposition to be true in terms of the more basic concept of atomic truth and/or explains what it is to be true in general in terms of the more basic concept of instantiation.

Now, the pluralists I am considering are interested in an explanation of what it is to be true in terms of several distinct truth properties, each of them corresponding to a certain topic area, such as propositions about medium size objects, propositions about moral matters, the propositions of mathematics, etc. However, as far as I can see, the only kinds of contrast of this sort that this framework provides for are two: the property of atomic truth vs the property of complex truth, and a subanalysis of the latter property into logical and non-logical truth. Thus, at first sight, the framework of structured propositions does not lend itself to a pluralism that corresponds to the differences between different "discourses" (compare Sher 1999, 2004). The set of concepts *N* with respect to which we are asking the question: how many properties of truth are there? offers us instantiation and atomic truth as explanatory basics. However, these basics to not provide the tools to make the distinctions that our pluralist wants to make, and for which she seeks and explanatory account.

In order to postulate further differences in the properties that realize truth in different discourses, the pluralist would need, presumably, to say that there are different kinds of atomic truth, or different kinds of instantiation. This, however, is in tension with the idea of the framework of structured propositions. For the motivation for the framework lies precisely in its capability of offering a comprehensive reductive account of truth, i.e. in the presumption that the basic notion (atomic truth and/or instantiation) is

---

[14] This sketch of a theory of structured propositions is supposed to be representative for a whole range of different such theories. I do not think that the details matter for the argument to come.

all we need to explain truth. Thus, a pluralist of the sort we are considering ought to develop a new or modified framework of propositions or other suitable truth bearers.[15]

References

Badici, Emil and Kirk Ludwig (2007). The Concept of Truth and the Semantics of the Truth Predicate. *Inquiry* 50, 622–638.

Davidson, Donald (1967). Truth and Meaning. *Synthese* 17, 304–323. Reprinted in his 1984, 17–36.

Davidson, Donald (1984) *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Davidson, Donald (1990). The Structure and Content of Truth. *Journal of Philosophy* 87, 279–328.

Dretske, Fred (1981). *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT-Press.

Dummett, M. (1993). *The Seas of Language*. Oxford: Oxford University Press.

Eklund, Matti (2002). Inconsistent Languages. *Philosophy and Phenomenological Research* 64, 251–275.

Fodor, J. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik* 100: 25–50.

Horwich, Paul (1990/98). *Truth*, Oxford: Oxford University Press.

Jackson, Frank (1998). *From Metaphysics to Ethics*. Oxford: Oxford University Press.

King, Jeffrey C. (2007). *The Nature and Structure of Content*. Oxford: Oxford University Press.

Kirkham, R. (1992/2001). *Theories of Truth*. Cambridge, MA: MIT-Press.

Kölbel, Max (2001). Two Dogmas of Davidsonian Semantics. *Journal of Philosophy* 98, 613–635

Kölbel, Max (2008a). Truth in Semantics. *Midwest Studies in Philosophy* 32, 242–257.

Kölbel, Max (2008b). 'True' as Ambiguous. *Philosophy and Phenomenological Research* 77, 359–384.

Laurence, S. & Margolis, E. (1999). Concepts and Cognitive Science. In E. Margolis & S. Laurence (eds.), *Concepts: Core Readings*. Cambridge, MA: MIT Press, pp. 3-81.

Ludwig, K. (2002). What is the Role of a Truth Theory in a Meaning Theory?. In D. Shier, J.K. Campbell & M. O'Rourke (eds), *Meaning and Truth: Investigations in Philosophical Semantics*, New York: Seven Bridges Press, 142–163.

Kaplan, David (1977/1989). Demonstratives. In Almog et al. (eds.), *Themes from Kaplan*, Oxford: Clarendon Press 1989.

Lewis, David (1970). General Semantics. *Synthese* 22, 18–67. Reprinted in Lewis 1983, 189–232.

Lewis, David (1983). *Philosophical Papers*, Volume 1. Oxford: Oxford University Press.

Lewis, David (1989). Dispositional Theories of Value. *Proceedings of the Aristotelian Society Supplementary Volume* 63, pp. 113–137.

Locke, John (1975). *An Essay Concerning Human Understanding*. P. Nidditch (ed.). Oxford: Clarendon Press.

Lynch, M. (2004). Truth and Multiple Realizability. *Australasian Journal of Philosophy* 82, 384–408.

Lynch, M. (2006). ReWrighting Pluralism. *Monist* 89, 63–85.

Lynch, M. (2009). *Truth as One and Many*. Oxford: Oxford University Press.

Matthews, Robert (1994). The Measure of Mind. *Mind* 103, 131–46.

Parfit, Derek (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Patterson, Douglas (2006). Tarski, the liar and inconsistent languages. *The Monist* 89, 150–177.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.

Peacocke, C. (2008). *Truly Understood*. Oxford: Oxford University Press.

Pedersen, Nicolaj Jang (2006). What can the Problem of Mixed Inferences Teach us about Alethic Pluralism?. *Monist* 89, 103–117.

Pedersen, Nikolaj Jang (2010). Stabilizing alethic pluralism. *PhilosophicalQuarterly* 60, 92–108.

Plato: *Meno*.

Richard, Mark (1997). Deflating Truth", *Philosophical Issues* 8, 57–78.

Russell, B. (1918). The Philosophy of Logical Atomism. In R.C. Marsh (ed.), *Logic and Knowledge*, London: Allen and Unwin 1956, 177–281.

Sher, Gila (1999). On the Possibility of a Substantive Theory of Truth. *Synthese* 117, 133–172.

Sher, Gila (2004). In Search of a Substantive Theory of Truth". *Journal of Philosophy* 101, 5–36.

Sher, Gila (2005). Functionalist Pluralism. *Philosophical Books* 46, 311–330.

Tappolet, Christine (1997). Mixed Inferences: a Problem for Pluralism about Truth Predicates. *Analysis* 57, 209–210.

Wright, Crispin (1992). *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press.

Zalta, E. (2001). Fregean Senses, Modes of Presentation, and Concepts. *Philosophical Perspectives* 15, 335–359.