

Epistemic Analyticity: A Defense

Paul A. Boghossian
NYU

Epistemic vs. Metaphysical Analyticity

In an earlier paper – “Analyticity Reconsidered” – I attempted to do two things: salvage a notion of analyticity from Quine’s widely accepted critique of that notion and show how it might be able to do serious work in the epistemology of a priori knowledge.¹

Salvaging analyticity, I argued, depends crucially on distinguishing between a metaphysical and an epistemic version of that concept. According to the *metaphysical* notion, a sentence is analytic if it owes its truth entirely to its meaning and without any contribution from the ‘facts.’ By contrast, I took a sentence to be *epistemically* analytic if grasp of its meaning can suffice for justified belief in the truth of the proposition it expresses.

I believe that Quine was deeply right to insist that there are no metaphysically analytic sentences. However, I argued, first, that the considerations that militate against metaphysical analyticity do not extend to epistemic analyticity and, second, that it is possible to provide a model of how some sentences might be epistemically analytic.

While I continue to believe that these claims are true, I have also come to think that “Analyticity Reconsidered” was less clear than it needed to be about the exact nature of the relation between meaning facts and entitlement. In particular, at the time of writing that paper, I did not delineate sufficiently clearly the difference between *inferential* and *constitutive* construals of the relation between meaning and entitlement. In this paper, I explain in what that difference consists and how it bears on the project of the analytic theory of the a priori. In the course of the exposition, I attempt to respond to various criticisms that have been made of that project and of my particular way of pursuing it.

Before proceeding, it is worth emphasizing that although I follow tradition in construing analyticity as a property of linguistic items, that feature of the presentation of the view is entirely optional. For example, I will talk of grasp of the meaning of a sentence as sufficing for justified belief in the proposition it expresses; but I could equally well simply have talked about grasp of a proposition *p* as sufficing for justified belief in *p*. Thus, too, I will talk about words being synonymous with each other; but I could equally well have talked about concepts being identical to one another. Finally, I will talk of holding some sentences true, as a condition of meaning some specific proposition by them; but I could equally well have talked of believing some propositions as a condition of having some of their ingredient concepts. Although I find it convenient for expository

¹ See “Analyticity Reconsidered,” *Nous*, 30: 360-391, 1996. A longer version containing some further discussion of Quine and an appendix on knowledge of meaning appears as “Analyticity,” in Hale and Wright, ed., *A Companion to the Philosophy of Language*, (Oxford: Blackwell, 1997), pp. 331-368.

purposes to assume that we think in a language (which I will pretend is in fact English), no central issue hangs on that assumption.

So, playing along with this linguistic picture, and assuming that T's believing that p is constituted by T's accepting a sentence S which means that p, let us ask: Why should it matter whether or not there are epistemically analytic sentences – sentences such that grasp of their meaning can suffice for justified belief in the truth of the proposition they express?

The interest of the epistemically analytic derives from the thought that it might help explain how there could be factual propositions known a priori. To see how this might work, let us first distinguish between three different things that one might mean by the phrase “grasp of meaning.”

In one sense, a thinker T grasps S's meaning just in case T means some determinate thing or other by his use of S. Call this notion of grasp “mere grasp of meaning.” In a second sense, T grasps S's meaning when he can correctly and knowledgeably *state* what S means. In this sense, it is not enough that T mean something by S – he must also have second-order knowledge of what S means. Call this notion of grasp “knowledge of meaning.” In the third and most demanding sense, T's grasping S's meaning implies not only that T is able to state S's meaning knowledgeably but also that he understands that meaning well enough to know whether or not it means the same as some other sentence. Call this “understanding of meaning.”

As we shall see, different models of epistemic analyticity employ one or more of these notions of grasp of meaning. And although these notions differ from each other in significant ways, they share the following property: each of them has been thought to be attainable without the benefit of empirical knowledge. Given this assumption, it becomes easy to see how the existence of epistemically analytic sentences might contribute to demystifying the phenomenon of a priori knowledge.

In a minute we shall look in detail about how such meaning-justification connections are to be construed, but I first want to take up an objection to the idea that there could be epistemically analytic sentences that are not metaphysically analytic. The objection has been pressed by Eric Margolis and Stephen Laurence:

After all, if p really is an independent fact that makes S true, then just knowing that S means that p couldn't suffice for the needed justification; one would also need to be justified in believing that p. In other words, so long as the truth of S isn't merely a matter of what it means, then grasping its meaning can only be (at best) part of the story about why one is justified in holding it to be true.²

If we try to turn the hunch that Margolis and Laurence are giving expression to here into an argument, we would need to rely on the following epistemic principle:

² Eric Margolis and Stephen Laurence, *Analysis*, 2001, 61:293-302, p. 294

So long as the truth of S isn't merely a matter of F, but is also a function of G, then being justified in believing F can only be (at best) part of the story about why one is justified in holding S to be true; one would also need to be justified in believing G.

But this is not in general a sound epistemic principle. The truth of the sentence "This is water" isn't merely a matter of how some substance looks or feels; it is also a matter of its being H₂O. However, it doesn't follow that I could never be justified in holding some stuff to be water without my first being justified in believing it to be H₂O.

At least as far as this particular argument is concerned, I see no reason to doubt that there could be epistemically analytic sentences that are not, per impossibile, also metaphysically analytic.

Explaining Epistemic Analyticity: Synonymy

The real problem with the epistemically analytic lies not in demonstrating its independence from the metaphysically analytic; it lies, rather, in explaining how any sentence *could be* epistemically analytic. How might grasp of the meaning of a sentence S suffice for justified belief in the truth of the proposition S expresses?

Let us start with the chestnut "All bachelors are unmarried males." How might grasp of the meaning of this sentence by a speaker suffice for his being justified in believing the proposition it expresses?

One route from meaning to justification that philosophers have often had in mind, without bothering to spell it out, would imagine our thinker reasoning as follows:

1. "All bachelors are unmarried males" means that All bachelors are unmarried males. (By knowledge of S's meaning)
2. Since "bachelor" just means "unmarried male," "All bachelors are unmarried males" is synonymous with "All unmarried males are unmarried males." (By understanding of meaning)
3. "All unmarried males are unmarried males" means that All unmarried males are unmarried males. (By knowledge of S's meaning)
4. If sentence F is synonymous with sentence G, then F is true iff G is true. (Conceptual knowledge of the link between meaning and truth.)
5. Therefore, "All bachelors are unmarried males" is true iff All unmarried males are unmarried males.
6. All unmarried males are unmarried males. (By knowledge of logic)
7. Therefore, "All bachelors are unmarried males" is true.
8. Therefore, All bachelors are unmarried males.

The template that this bit of reasoning instantiates may be represented as follows:

1. S means that P. (Knowledge of Meaning)
2. S is synonymous with S'. (Understanding of meaning)

3. S' means that Q, where Q is some logical truth. (Knowledge of Meaning)
4. If F is synonymous with G, then F is true iff G is true. (Conceptual link between meaning and truth)
5. Therefore, S is true iff Q. Q (Logic)
6. Therefore, S is true. (Deductive reasoning)
7. Therefore, P. (Deductive Reasoning)

Let us call this the “Synonymy Template” and the model of a priori knowledge that it presents the “Synonymy Model.” If someone in possession of justification for the premises of this argument actually runs through it, I will say that he is *justified* in believing P; if someone is merely *able* to run through such an argument without actually doing so, I will say that he is *entitled* to P.³

Quine, of course, objected that the notion of synonymy on which this model relies is not sufficiently well-defined to do serious explanatory work. However, as I argued in “Analyticity,” no Meaning Realist can afford to accept Quine’s argument.

I think that the Synonymy model is correct so far it goes: it correctly explains the structure of our knowledge of *some* a priori truths – of those ‘conceptual’ truths, namely, that are transformable into logical truths by the substitution of synonyms for synonyms (sentences that in “Analyticity Reconsidered” I called “Frege-analytic.”).

However, the Synonymy model is obviously not a *complete* story about how grasp of meaning might generate entitlement, and this in several respects.

First, it relies on a piece of a priori conceptual knowledge at premise 4 – connecting meaning and truth – that it cannot explain.

Second, it relies on a priori knowledge of logical principles that it cannot explain, and this in two ways. First, it relies on the thinker’s knowing that a certain principle is a truth of logic; second, it relies on the thinker’s being able justifiably to infer in accordance with deductive rules of logic, so that the thinker is able knowledgeably to move from the premises to the conclusion of the Template.

Finally, it seems quite clear that there will be *other* propositions, beyond the ones mentioned thus far, that the Synonymy Model will be unable to account for. For there do seem to be a considerable number of a priori propositions that are expressed by sentences that are not transformable into logical truths by the substitution of synonyms for synonyms. For example,

- Whatever is red all over is not blue.
- Whatever colored is extended.
- If x is warmer than y, then y is not warmer than x.

³ This invokes Tyler Burge’s well-known distinction between justification and entitlement. See his “Content Preservation,” *The Philosophical Review*, 102 : 457-88, 1993.

If meaning-based explanations of the a priori are to be complete, a meaning-based explanation must be found for the apriority of each of these other types of proposition.

Explaining Epistemic Analyticity: Implicit Definition

As I explained in “Analyticity Reconsidered,” one of the central ideas in this connection is that of Implicit Definition (ID). I will look in detail here at the case of logic, postponing to another occasion a description of how ID explanations of apriority might be extended to the other cases we described.

As a first approximation, we may explain the idea of ID, as applied to logic, thus:

It is by stipulating that certain sentences of logic are to be true, or that certain inferences are to be valid, that we attach a meaning to the logical constants. More specifically, a particular logical constant means that logical concept, if any, which makes valid a specific set of sentences and/or inferences involving it.

For example, one might think that:

It is by stipulating that

It is not the case that both P and not P

is to be true, that someone comes to mean negation by ‘not.’

Or, to pick another example that will feature further below:

It is by stipulating that all inferences of the form

p
If, p, then q
Therefore, q

are to be valid that someone comes to mean *if* by ‘if.’

In considering implicit definitions, we must bear in mind that they come in two varieties: explicit and implicit. An *explicit* implicit definition involves an explicit stipulation by a thinker that a given sentence S(f) is to be true if its ingredient term f is to mean what it does. In the *implicit* variety, it is somehow tacit in that person’s behavior with the term f that S(f) is to be true if f is to mean what it does.

Later on I will come back to what it might mean for an implicit definition to be itself implicit. For now, let us operate with the explicit version and ask: how would S(f)’s being an implicit definer for f help explain how grasp of S(f)’s meaning might suffice for justified belief in its truth?

One suggestion, though as we shall see, not ultimately the favored one, would mimic the Synonymy model by supplying an argument template that a thinker will be in a position to perform as a result of grasping the meaning of S(f).⁴

1. S(f) means that P (By knowledge of meaning)
2. If S(f) means that P, S(f) is true (By knowledge of the contents of one's stipulations)
3. Therefore, S(f) is true
4. If S(f) means that P, then S(f) is true iff P (By knowledge of the link between meaning and truth)
5. S(f) is true iff P
6. Therefore, P

Call this the "Implicit Definition Template" and the model of the a priori that it presents the "Implicit Definition" model.⁵

Knowledge of Meaning and Implicit Definition

Several philosophers have objected to this model of the a priori. The main recurring complaint has been nicely expressed by Kathrin Gluer:

... let P be the proposition that S(f) expresses. On Boghossian's account, it seems to me, being justified in believing that S(f) means P *presupposes* being justified in believing P. Of course, it then *follows* from being justified in believing that S(f) means P that I am justified in believing P. However, believing P cannot be justified in this way; for knowledge of the meaning of the sentence already requires such justification and, therefore, cannot provide it.

Laurence Bonjour makes a similar point when he says:

Thus, for example, one might stipulate that the sentence '40@8=5' is to count as a (partial) implicit definition of the symbol '@'. This, along with other stipulations of the same kind, might prove a useful way of conveying that '@' is to stand for the operation of long division (assuming that the other symbols in the sentence are already understood). But if this is the right account of implicit definition, then the justification of the proposition that 40 divided by 8 is equal to 5 (as opposed to that of the linguistic formula '40@8=5' is not a result of the implicit definition, but is rather presupposed by it: if I were not justified in advance, presumably a priori, in believing that forty divided by eight is equal to five, I would have no reason for interpreting '@' in the indicated way.⁶

⁴ This seems to me a better representation of the relevant template than what I offered in "Analyticity," but it is not materially different.

⁵ An example of an a priori that this model might plausibly be taken to explain is the proposition *Vixens are female foxes*.

⁶ Bonjour, *In Defense of Pure Reason*, (Cambridge: CUP, 1998), pp. 50-51.

According to Bonjour and Gluer, if S(f)'s expressing P is fixed via the stipulation that S(f) is to be true, one cannot be justified in believing that S(f) expresses P without *first* being justified in believing that P is true. If that is right, then the Implicit Definition Template could not explain how someone could acquire a priori warrant for the belief that P, since warranted belief in P would be presupposed by anyone running an argument of that form.

This is a surprising objection. As I noted towards the beginning of this paper, most philosophers simply assume that meaning facts are first-person accessible in some privileged way, regardless of what the supervenience base for meaning facts is taken to be. For example, many philosophers believe that even if facts about meaning and concept possession were to supervene on facts that are *external* to the mind that that would have no tendency to undermine our privileged access to first-person facts about meaning.

Now, it is an interesting question whether these many philosophers are right to make this assumption. I, for one, have worried about whether we genuinely understand how first-person access to meaning is possible, especially if meaning is externally determined. I am not a skeptic about such access; I simply believe that we don't really understand how it works.

Bonjour and Gluer, however, write neither as skeptics about privileged access, nor in the service of raising a question about how such access should ultimately be understood. They write as though there is reason to think that in the *special* case where meaning is fixed by implicit definition, there is a problem with the assumption of privileged access. In that case, they assert, someone cannot be said to know that S means that P without first knowing that

- a. if S means that P, then S has to be true
and that
- b. P is true.

But they have supplied us with no special reason to think that, if S's meaning is fixed via implicit definition, the usual assumption of privileged access must be suspended.

This point is only strengthened when we reflect that the only plausible version of ID is not the explicit one with which we have been working so far, but rather the implicit one.

It is rare for a term to be introduced via some explicit stipulation. And in the special case that interests us – the case of the logical constants – it is not only rare but incoherent. As Quine pointed out, if the logical constants are to be thought of as having their meaning fixed by implicit definitions, that meaning cannot be thought of as fixed by explicit implicit definitions, since the logical constants will have to be presupposed in any statement of the stipulations by which we might seek to fix their meanings.⁷ Rather, if

⁷ See his "Truth by Convention," in his *The Ways of Paradox*, (Cambridge: Harvard University Press, 1976).

there is to be anything at all to the idea of ID as applied to logic, it must be that the logical constants have their meaning fixed by our *tacitly* regarding some of the inferences involving them as valid, or by our *tacitly* regard some of the sentences involving them as true. It's a good question what this sort of tacit stipulation amounts to – many a conceptual role semantics has struggled with that question. But it is a question to which there has to be an answer if, as seems likely, our only hope of explaining how we come to grasp the concepts of the logical constants is through the idea of ID.

But do Bonjour and Gluer really wish to say that if the meaning of 'and' is fixed by a thinker's being disposed to use it according to its standard introduction and elimination rules that he cannot be said to know what 'and' means without first knowing that

A and B implies A?

That would clearly be absurd, since the thinker will have no hope of knowing that without having the concept of conjunction to begin with.

If this particular style of objection is to be sustained, we need to be given a special reason for thinking that where a conceptual role semantics is concerned, there the usual assumption of privileged access must be rejected. But I don't see that we have been given any such reason.

Implicit Definition and Entitlement: The Constitutive Model

However, there are several other difficulties with the suggestion that the ID Template gives a fundamental account of the entitlement that implicit definition is able to provide.

First, there continues to be a reliance on the link between meaning and truth at step 4. Even if we can correctly say that we know that

If S means that P then S is true iff P

because we have stipulated it to be true, we cannot hope to explain *why* the stipulation grounds the knowledge via the Implicit Definition Template, because that template relies on our knowledge of that link.

Second, no such argument template could possibly hope to explain in what our entitlement to *reason* according to certain deductive rules consists, since – once again – it presupposes such reasoning.

Finally, if we are operating with the more promising *tacit* version of implicit definition rather than with the version according to which our stipulations are explicit, there is the problem that we can no longer rely on its following from the fact that S(f) is an implicit definer of f that S(f) is true. On the view according to which implicit definitions are themselves implicit, there is a difficulty seeing how it would follow from S(f)'s being an implicit definer for f, that S(f) has to be true. For how is it going to be implicit in someone's behavior with S(f) that it is acting as an implicit definer? Presumably, by

S(f)'s being used in a certain way – most plausibly, by its being held to be true come what may. But the point is that there is all the difference in the world between saying that a certain sentence must be held true, if it is to mean this that or the other, and saying that it *is* true.⁸ But it's actually being true is what premise 2 of the ID Template requires.

What these problems suggest, therefore, if implicit definition is to be a genuine source of a priori entitlement, is that there must be a different way in which implicit definitions can generate entitlement other than by supplying premises from which the truth of various propositions may somehow be derived.

Call this premise-and-derivation model, on which we have been concentrating so far, the *inferential* conception of how meaning generates entitlement. On the contrasting *constitutive* model, the thinker doesn't start with some premise about some sentence S's meaning from which he deduces that S is true. Rather, the mere fact that the thinker grasps S's meaning entails that the thinker is justified in holding S to be true. Or, if we focus on inferences rather than sentences: the mere fact that the thinker grasps inference rule R's meaning entails that the thinker is justified in inferring according to R. How would this work?⁹

Look first at the case that is likely to be of most central interest, the case of an inference rule's being meaning-constituting. How might it turn out that, as a mere consequence of R's being meaning-constituting for a thinker T, T is justified in inferring according to R?

Suppose R is Modus Ponens (MPP) and suppose also -- and plausibly -- that being willing to infer according to MPP is constitutive of possession of the concept *if*. How could that fact explain how we might be entitled to reason according to MPP?

One initially helpful thought is this. Suppose it's true that my taking p and 'if p, then q' as a warrant for believing q is constitutive of my being able to have *if* thoughts in the first place. Then doesn't it follow that I could not have been epistemically blameworthy in taking p and 'if p, then q' as a reason for believing q even in the absence of any reason for taking those premises to be a reason for believing that conclusion? If inferring from those premises to that conclusion is required, if I am to have the ingredient propositions, then, as a matter of metaphysical necessity, I cannot so much as consider the question whether the inference is justified without being disposed to reason in that way. Under those circumstances, then, it looks as though inferring according to MPP cannot be held against me, even if the inference is, as I shall put it, *blind* – unsupported by any positive warrant.

⁸ This point was first made in my "Inferential Role Semantics and the Analytic/Synthetic Distinction," in *Philosophical Studies*, 73: 109-122, 1994.

⁹ In "Analyticity Regained?" Harman correctly emphasizes the importance of the distinction between inferential and "direct" ways in which meaning might justify a proposition. However, neither his distinction nor his way of explaining direct a priori justification line up exactly with the notions that I outline here.

*Problems for the Meaning-Entitlement Connection*¹⁰

This explanation, however, is flawed. If we spell out the principle underlying it, it would be this:

(Meaning-Entitlement Connection, or MEC): Any inferential transitions built into the possession conditions for a concept are *eo ipso* prima facie entitling.

And the trouble is that, at least as stated, there seem to be clear-cut counterexamples to the MEC: cases where the acceptance of some inference is written into the possession of a given concept but where it is also clear that the inference isn't one to which the thinker is entitled.

One famous illustrative case is Arthur Prior's connective 'tonk.'¹¹ To possess this concept, Prior stipulated, a thinker must be willing to infer according to the following introduction and elimination rules:

(Tonk)	A	$A \text{ tonk } B$
	-----	-----
	$A \text{ tonk } B$	B

Obviously, no one could be entitled to infer any B from any A; but that's exactly what is implied by the MEC.

A similar conclusion can be drawn from the case of racist or abusive concepts, provided that an inferential role semantics extends to them as well. (Because it makes life harder for me, I will here make that assumption, though I am not positively committed to it.) Take, for example the concept *boche* discussed by Dummett.¹² According to Dummett, a thinker possesses the concept *boche* just in case he is willing to infer according to the following introduction and elimination rules:

(Boche)	$x \text{ is German}$	$x \text{ is boche}$
	-----	-----
	$x \text{ is boche}$	$x \text{ is cruel.}$

Yet no one is entitled – let alone simply as the result of the introduction of a concept into the language – to the view that all Germans are cruel.

How should we think about such cases?

¹⁰ The next two sections draw on material from my 'Blind Reasoning' in *Proceedings of the Aristotelian Society*, Supplementary Volume, LXXVII, 2003, pp. 225-248.

¹¹ A. Prior, "The Runabout Inference Ticket," *Analysis*, 21 (December 1960), pp. 38-39.

¹² M. Dummett, *Frege's Philosophy of Language*, (New York, Harper and Row, 1973), p. 454.

Robert Brandom has this to say about ‘boche’-like concepts:

The use of any concept or expression involves commitment to an inference from its grounds to its consequences of application. Critical thinkers, or merely fastidious ones, must examine their idioms to be sure that they are prepared to endorse and so defend the appropriateness of the material inferential commitments implicit in the concepts they employ. ..The proper question to ask in evaluating the introduction and evolution of a concept is not whether the inference embodied is one that is already endorsed, so that no new content is really involved, but rather whether the inference is one that *ought* to be endorsed. The problem with ‘boche’ is not that once we explicitly confront the material inferential commitment that gives the term its content it turns out to be novel, but that it can then be seen to be indefensible and inappropriate – a commitment we cannot become entitled to.¹³

Unfortunately, Brandom’s observations cannot help defend the MEC against the threatening counterexamples. I am not sure where Brandom stands on the MEC, although I am inclined to think that he is committed to it. However that may be, it is not enough, it seems to me, to defend the MEC by saying that it is always open to us to reject a concept, once we recognize that it commits us to claims that we cannot become entitled to. For what is implausible in the case of either ‘tonk’ or ‘boche’ is that there is any entitlement there at all, defeasible or not. But defeasible prima facie entitlement is precisely what the MEC commits us to.

If we are to save the MEC, we must do one or both of two things: either restrict it to certain concepts from which entitlement really does flow, or restrict what we count as a genuine concept.

The latter strategy is suggested by the work of Christopher Peacocke who has long urged that we should require that the meaning-constituting rules of a genuine concept be truth-preserving.¹⁴

If we adopt this requirement, we can say that what’s wrong with both ‘tonk’ and ‘boche’ is precisely that there is no concept that those terms express, for there is no reference for ‘tonk’ and ‘boche’ that’s capable of making all of their constitutive rules truth-preserving.

While this might seem to yield the right result for ‘tonk’ it doesn’t yield the right result for ‘boche’: it’s hard to believe that racists who employ boche-like concepts fail to express complete thoughts.

¹³ R. Brandom, *The Articulation of Reasons*, (Cambridge, MA: Harvard University Press, 2000), pp. 70-72.

¹⁴ See Peacocke “Proof and Truth,” in *Reality, Representation and Projection* eds. J. Haldane and C. Wright (New York: Oxford University Press, 1993). I myself took this line in “How are Objective Epistemic Reasons Possible?” so the present paper represents a change of heart on this point.

And even if we were to put this complaint to one side, it seems clear that truth-preservation alone will not suffice for dealing with our problem about the MEC.

Consider the concept *flurg* individuated by the following introduction and elimination rules:

(Flurg)

<i>x is an elliptical equation</i>	<i>x is flurg</i>
<i>x is flurg</i>	<i>x can be correlated with a modular form</i>

It turns out to be a result that Wiles had to prove on the way to proving Fermat’s Last Theorem that every elliptical equation can be correlated with a modular form (the Taniyama-Shimura conjecture). Once again, therefore, we have no independent reason to think that these introduction and elimination rules are not necessarily truth-preserving. But it’s hard to see that one is a priori entitled, merely on the basis of introducing the term “flurg,” to the Taniyama-Shimura conjecture. So there is still a problem for the claim that entitlement flows from meaning-constitution, given only the requirement that a concept’s introduction and elimination rules be truth-preserving.

Defective Concepts and Blameless Inference

I want to propose a different diagnosis of what has gone wrong with concepts such as *flurg* and *boche*, one that doesn’t depend on denying that they constitute genuine thinkable contents. That diagnosis will allow us to demarcate a class of concepts that are intuitively epistemically *defective*, with the result that inference in accord with their constitutive rules is not entitling.

If our account of the possession conditions for *boche* is correct, a thinker who possesses that concept is in effect implicitly affirming a theory – call it *T(boche)* – that consists of the following pair of claims:

- (1) For any x, if x is a German then that person has the further property of being boche.
- (2) For any x, if x is boche, then that person has the property of being cruel.

Typically, of course, a thinker will not affirm this theory explicitly, but only implicitly, by being willing to find primitively compelling the constitutive rules for boche. But whether the affirmation is explicit or implicit, it is quite clearly a commitment to a very substantive pair of claims, claims to which one cannot be entitled simply by introducing a concept into one’s language.

The substantial nature of the commitments is illustrated by the fact that someone could choose to rationally dispute them. But the concept is set up in such a way as to preclude that sort of rational disagreement. For to question whether there are any boches just

amounts to questioning whether there are any Germans; and to question whether boches are necessarily cruel is to question a constitutive possession condition for the concept.

‘Boche,’ then, is problematic in that possessing it requires a thinker to believe a substantive theory about the world. And it does seem intuitively right to insist that, in so far as is possible, a concept should not be set up in such a way that only those who are prepared to affirm a controversial creed are allowed to possess it.

How are we to avoid that? Well, by conditionalizing, where possible, on the substantive claims and commitments of the unhedged concept. Thus, instead of operating with *boche*, with its substantive claims and inferences,

$$T(\textit{boche}) = Gx \rightarrow Bx \text{ and } Bx \rightarrow Cx$$

we could work with *Condboche* whose constitutive claims would be merely this:

$$T(\textit{Condboche}) = \textit{If there is a property } F \textit{ such that } T(F), \textit{ then } T(\textit{boche}).$$

Once again, we would need to imagine a thinker affirming this as implicitly as possible, and so perhaps as operating with a rule along the following lines:

$$\frac{\exists FT(F)}{\text{-----}} T(\textit{boche})$$

On this picture, the thinker’s willingness to move from immediately from Gx to Bx and Bx to Cx is conditional on his background belief that there is an instantiated property F satisfying the conditions on boche theory.

Ordinary scientific terms in good standing – “neutrino” for example – are held to have just this feature, of intelligibly allowing for disagreement about their governing theories. Thus, we don’t think of the rules which correspond to our possession of the concept *neutrino* as consisting in the propositions that would actually be believed by a proponent of neutrino theory, but rather as corresponding only to what someone would be willing to believe who was *conditionalizing* on the truth of neutrino theory.

If, following Russell, Ramsey, Carnap and Lewis, we represent neutrino theory

$$T(\textit{neutrino})$$

as the conjunction of the two propositions

$$(S) \exists FT(F)$$

and

(M) $\exists FT(F) \rightarrow T(\textit{Neutrino})$,

then the point is that we think of possession of the concept *neutrino* as requiring someone to affirm only M and not S as well.¹⁵

Now, someone could certainly introduce a concept that did not have the conditionalized structure that I've claimed is actually true of *neutrino*, but which consists rather in the inferences that are characteristic of neutrino theory unconditionalized. Call this *neutrino+*. Such a person would insist that it is a condition on having *his* concept of neutrino that one be willing to endorse the characteristic claims and inferences of neutrino theory, and not merely the conditionalized claim captured in (M). But, for the reasons previously articulated, there would be something epistemically *defective* about this concept, even if its constitutive rules turned out to be truth-preserving.

Flurg and *neutrino+*, and *boche* then, all suffer from the same problem: they are unconditional versions of a concept, when only its conditionalized version would be epistemically acceptable. I don't think we should put this by saying that they are not real concepts. Concepts are relatively cheap. But they are *defective* concepts. They are structured in such a way that perfectly reasonable questions about their extensions are foreclosed.

Under what conditions is only a conditionalized version of a concept acceptable? Here I want to make two claims, one sober, one bold.

(Bold) Whenever both a conditional and an unconditional version of a given concept are available, it is the conditional version that ought to be used. Given the availability of both versions, the unconditional version counts as epistemically defective.

(Sober) In the case of some concepts, only the unconditionalized version will be available.

We have just been examining the argument for Bold: You don't ever want the *possession conditions* for a concept to foreclose on the possible falsity of some particular set of claims about the world, if you can possibly avoid it. You want the possessor of the concept to be able coherently to ask whether there is anything that falls under it, and you want people to be able to disagree about whether there is. If in a certain range of cases, however, it is logically impossible to hold the governing theory at arm's length then, in

¹⁵ This paragraph follows Paul Horwich's discussion of the conditional nature of semantic stipulation in "Stipulation, Meaning and Apriority," in Boghossian and Peacocke: *New Essays on the A Priori*. I ignore various complexities that a thorough discussion of the representation of scientific theories would require and, in particular, the need to account for *uniqueness*. These further complexities would only have strengthened the case for the central claim that I will want to make later on, namely, that a considerable logical apparatus is presupposed in the possession of any conditionalized concept.

those cases, obviously, it can hardly be a requirement that one do so. But in all those cases where it is possible, it ought to be done.

What about Sober? I think it is clear, given what it means to conditionalize on the truth of an arbitrary theory, that not every meaningful term in a language can be thought of as expressing a concept that conditionalizes on the existence of an appropriate semantic value for it. In particular, a certain number of logical concepts will be presupposed in any conditionalization and those that are so presupposed will not themselves have conditionalized versions.

Timothy Williamson has objected to this line of reasoning as follows:

Although \exists and \rightarrow occur in the Carnap sentence $\exists FT(F)$, in place of that sentence Boghossian could have used the rule that allows one to infer $T(\textit{Neutrino})$ directly from any premise of the form $T(A)$. That rule is formulated without reference to the logical operators in the object-language, but is interderivable with the Carnap sentence once one has the standard rules for E and \rightarrow Logical operators may of course occur in the theory T itself, although Boghossian does not appeal to that point. In any case, it seems insufficiently general for his argument, since for some less highly theoretical concepts than *neutrino*, the analogue of the theory T for conditionalization may consist of some simple sentences free of logical operators.¹⁶

Of course, I did not mean to suggest that one could simply read off the Carnap sentence that existential quantification and conditional would be presupposed by any conditionalization, though no doubt my presentation was overly elliptical. In the cases of most central interest, the affirmation of the Carnap sentence would be *implicit* in the thinker's behavior and could not be supposed to amount to an explicit belief from which one could simply read off the ingredient conceptual materials.

To see whether we could have nothing but conditionalized concepts, we have to ask whether it is possible for someone to implicitly affirm the Carnap sentence for, e.g., boche, without possessing any of the logical concepts with which we would explicitly have conditionalized our concepts.

We have agreed that for someone to affirm $T(\textit{boche})$ implicitly is for them to be willing to infer according to the following introduction and elimination rules:

$Gx/Bx \quad Bx/Cx$.

Now, the question is: What would it be for a thinker to implicitly conditionalize his affirmation of $T(\textit{boche})$ on the existence of an appropriate semantic value for these rules.

¹⁶ "Understanding and Inference," in *Proceedings of the Aristotelian Society*, Supplementary Volume, LXXVII, 2003, pp. 249-293.

Williamson says that this could be adequately captured by picturing the conditionalizing thinker as operating according to the following rule:

$$\frac{T(A)}{\text{-----}} \\ T(\textit{boche})$$

But what this seems to me to say is something very different from what is needed. A thinker operating according to Williamson's rule is like someone who already has the concept *boche* but is now simply relabeling it with the word 'boche.' Whereas what I want to capture is the idea of someone who is only prepared to infer according to the *boche* rules because they antecedently believe that

There is a property F , such that $Gx \rightarrow Fx$ and $Fx \rightarrow Cx$

And I don't see how their reasoning could depend on that without their having, at a minimum, the conceptual materials that make up the antecedent of the Carnap sentence, including the quantificational apparatus and the conditionals that make up the statement of the theory.

If all of this is right, it follows that conditional counterparts for one's primitive logical constants will not be available and hence that one could hardly be blamed for employing their unconditionalized versions. In particular, you couldn't conditionalize on the existence of an appropriate truth function for the conditional, for you would need it in order to conditionalize on anything. In such a case, there is no alternative but to accept "conditional theory" – Modus Ponens and Conditional Proof, in effect – if you are to so much as have the conditional concept. It thus couldn't be epistemically irresponsible of you to just go ahead and infer according to MPP without conditionalizing on the existence of an appropriate truth function for it – that is simply not a coherent option in this case.

If we go back to the MEC, it seems clear how that principle needs to be modified:

(MEC*) Any rules that are written into the possession conditions for a *non-defective* concept are a fortiori entitling.

And with this principle in hand, we have the answer to our question: How could a thinker be entitled to reason according to MPP just in virtue of grasping the meaning of that rule? The answer is that he can be so entitled because MPP is a possession condition for the conditional, and the conditional is a non-defective concept.

Gilbert Harman has asserted that the only non-inferential route from grasp of meaning to entitlement must run through Conventionalism and, therefore, through the dubious doctrine of metaphysical analyticity.¹⁷ However, if the Constitutive Model I have offered

¹⁷ Gilbert Harman, "Analyticity Regained?" pp. 145-6.

is correct, his assertion is false, for the Constitutive Model is in no way committed to Conventionalism or to metaphysical analyticity.

Conclusion

The most fundamental relation between grasp of meaning and entitlement occurs when a thinker is entitled to reason in accord with a certain rule R simply by virtue of the fact that R is constitutive of a non-defective concept of his.

This Constitutive model could be extended to account for the thinker's entitlement to have certain beliefs, provided that those beliefs were similarly constitutive of the possession of a non-defective concept. I think that something like this story will explain how we get to be entitled to believe in the principle of non-contradiction and in the link between truth and meaning that featured in the templates described above.

With these two pieces in place, it becomes possible to invoke the Implicit Definition and Synonymy Templates to explain the a priori knowability of other, less fundamental, a priori truths.¹⁸

¹⁸ I am grateful to Stephen Schiffer for comments on an earlier version of this paper.