Inferentialism and the Epistemology of Logic: Reflections on Casalegno and Williamson

1.

Paolo Casalegno's 'Logical Concepts and Logical Inferences' (Casalegno 2004) is a searching and insightful critique of my attempt to explain how someone could be entitled to infer according to a basic logical rule. I will say a bit about what I take the problem to be before considering his discussion in detail.

Let us agree that we reason according to logical rules. (There are various issues about this, most forcefully pressed by Gilbert Harman (see Harman 1986), which I propose to set aside for present purposes.) One of the most central of the rules by which we may be said to reason is *Modus Ponens*, which I will take to say (again ignoring many complexities which are irrelevant for present purposes):

(MP) Whenever both p and 'if p, then q', infer q.

Those of us who have learned to formulate this rule recognize it as a rule that we operate with and, putting aside some deviant logicians, consider it valid.

However, there are many perfectly rational persons – youngsters or undereducated adults – who are not aware that this is a rule that they operate with. Despite this, we think that when such people reason according to MP:

(Rain) It rained last night.

If it rained last night, then the streets are wet.

So,

The streets are wet.

they are perfectly *entitled* to do so; and that the justification that they have for their premises transmits smoothly to their conclusion.

In what does their entitlement to reason according to MP consist?

Obviously it cannot consist in some *argument* that they have formulated for the belief that this form of inference is valid. By assumption, they do not have the belief in question so could hardly have seen the need to formulate an argument for it. But if their entitlement does not consist in an explicit justification for the validity of MP, what does it consist in?

What about us sophisticated philosophers who have arrived at the knowledge that we reason according to the rule MP? In what does our entitlement to use this rule consist?

Have not we formulated an explicit justification for operating according to this rule? And if we have not actually bothered doing so, is not it clear that we are at least in a position to do so, if asked and given enough time for reflection?

It might seem as if the answer to the latter question is 'Yes'. For could not we offer something like the following argument (I will not bother with the niceties of semantic ascent and descent)?

(i) If 'p' is T and '$p{\to}q$' is T, then 'q' is T (by knowledge of the truth table)

(ii) 'p' is T and '$p{\rightarrow}q$' is T (by assumption)

Therefore,

(iii) 'q' is T (by MP)

Well, this particular argument obviously cannot do much for us by way of justifying our use of MP, since it relies on MP at its third step and I am assuming that we cannot explain in what our entitlement to reason with a certain rule, R, consists by showing that we have to hand an argument for R that employs R.

Of course, we might be able to offer other justifications for MP that rely not on MP itself but on certain other rules, say $R1$ and $R2$. But then the question will arise what entitles us to use $R1$ and $R2$?

We could now repeat the process of providing a justification for $R1$ and $R2$. Pretty soon, though, our justifications will end up appealing either to *Modus Ponens*, or to $R1$ or $R2$, or to some other rules for which we will owe a justification.

It seems obvious, then, that even the most sophisticated and powerful philosopher will face the following dilemma: with regard to her most basic logical rules, either she has *no entitlement* to them, or she has an entitlement that is *not* grounded in her ability to provide an explicit *argument* for them.

The skeptical alternative is dire. For if she has no entitlement to her most basic rules, then she has no entitlement to anything that is based upon them; and that means that she will have *no* entitlement to *any* of the rules of logic that she is inclined to use and therefore no entitlement to any of the beliefs that she will have based on them. This seems to me too fantastic to believe.

It also seems to me to tee up an extreme form of *relativism* about rationality, one that I find worrisome, both philosophically and socially. For if none of us is entitled to the particular set of logical rules that we operate with, then if others among us were to find it natural to operate with a *different* and incompatible set of logical rules, then they would have to be deemed as rational as we are, in so far as their use of logical rules is concerned. We could not say that such people were irrational, for they are surely no worse off in their entitlements to their logical rules than we are with respect to ours.

The skeptical alternative, then, is fraught with difficulty.

The non-skeptical alternative, however, requires us to explain how someone might be entitled to operate according to some basic logical rule, say MP, without his being able to provide anything like a cogent argument for MP. How could someone be so entitled? In particular, how could someone be so entitled in a way that did not imply an "anything goes' conception of 'logical rationality?"

I hope it's obvious that this is a highly non-trivial task. Indeed, it remains unclear to me even now that there is a way of executing it that is even remotely satisfying.

Large as the task may be, though, it is not quite as large as the task of providing an *overall* epistemology of logic. As I am thinking about it, refuting the skeptical threat involves showing that it is possible to get our entitlement to *some* logical rules off the ground, even if the means by which that is accomplished may not generalize to all the logical rules to which we feel intuitively entitled. What is needed, in other words, is a plausible answer to a Kantian-style "How possible?" question. We may worry about how to account for the rest of logic later. This relates to a point that shows up in an exchange that I have had with Timothy Williamson on these issues. Tim says:

> In previous work, Boghossian developed an epistemology of logic based on understanding-assent links corresponding to fundamental rules of logic. His paradigm was modus ponens: a necessary condition for understanding 'if' was supposed to be willingness to assent to

inferences by modus ponens involving 'if'. The book presents a series of counterexamples, some actual, some possible, to such putative understanding-assent links, for both modus ponens and other equally fundamental rules (85–121). The counterexamples concern native speakers of a natural language who come to understand the logical words at issue in the usual way but then go in for deviant logical theorizing without losing their linguistic competence; most philosophers know such people. In response, Boghossian picks what he regards as the clearest understanding–assent link, willingness to assent to 'and'-elimination (the inference from 'P and Q' to 'P' or to 'Q') as a condition for understanding 'and', and denies that the counterexamples I propose to it (95–6) make sense.

Strategically, Boghossian's response is not very promising. If he can rely on understanding-assent links only for 'and'-elimination and a few other equally banal rules, but not for modus ponens or other fundamental principles, then he is in no position to base either a *general* epistemology of logic or a *general* account of the understanding of logical constants on understanding–assent links. It is a little lame for him to claim in effect that not *every* fundamental rule of logic is a counterexample to his original account. A bolder strategy for him would be to seek a way of defending the claim that *no* fundamental rule of logic is a counterexample to his original account, and in particular of defending his original test case, modus ponens, as a putative understanding-assent link for 'if' against my counterexamples. In keeping away from the bolder strategy, Boghossian concedes so much ground that it is quite unclear what his fallback general epistemology of logic or his fallback general account of the understanding of logical constants could be. (Williamson 2011, 500)

Now, I do not, of course, deny the desirability of having a general epistemology of logic, but, as I have just been emphasizing, to my mind the fundamental difficulty in this area is to show that there might be even a *single* promising pathway for avoiding skepticism about our entitlement to the fundamental rules of logic. So, *pace* Tim, I would be happy if, in the first instance, I could come up with a plausible account just for one or two of the most 'banal' rules of logic.

2.

Well, what are the possible anti-skeptical alternatives? We have ruled out accounts that trace our entitlement to using a basic rule of logic in terms of our ability to provide any sort of argument. Could we plausibly say that it consists in some sort of non-inferential warrant?

A traditionally influential answer along these lines deploys the idea of an 'intuition.' A thinker is entitled to MP if he intuits its validity in some essentially non-discursive and non-inferential way.

I find this answer very problematic, for reasons that I have developed elsewhere (see Boghossian 2003), although I am now inclined to be more sympathetic to the notion of an intuition than I used to be.

If we put aside explanations in terms of intuition, then the anti-skeptical task before us becomes one of explaining how a thinker might be differentially entitled to MP *blindly*, without being in a position to point to *any* sort of justification for his use of MP, whether this be of an inferential or non-inferential variety.

How could we be blindly entitled to operate according to MP? There look to be two main options.

The first would consist in embracing a crudely reliabilist conception of inferential justification, according to which a thinker is entitled to the use of a rule if that rule is reliably truth-preserving, then

we would have an easy answer to our problem. Any thinker would be blindly entitled to MP, since (let us assume) MP is necessarily truth-preserving.

However, such a crude reliabilism is clearly false. There are lots of logically valid inferences, for example, from the Peano axioms to any instance of the inequality of Fermat's Last Theorem, which no one would be entitled to perform merely as a result of their reliability.

The second avenue for explaining blind justification involves deploying the classical notion of *analyticity*. For one important strand in that notion is the *epistemic* idea that *understanding* alone can sometimes suffice for entitlement: it is plausible, for example, that the mere understanding of the word 'bachelor' suffices for our knowing that all bachelors are male. If, on analogy with this, our understanding of 'if' could be shown to suffice for our being entitled to use MP, we would have the answer to our skeptic.

Taking such a notion of epistemic analyticity seriously, though, required showing that it could be detached from the much more dubious doctrine of metaphysical analyticity – of truth (or validity) in virtue of meaning – with which it had always been associated, but this seemed doable (see Boghossian 1996).

The problem then became one of explaining concretely how our understanding of 'if' might suffice for our entitlement to use MP.

An obvious starting point was a theory of our understanding of the logical constants that had always found favor among philosophers, quite independently of epistemological issues, according to which to grasp a logical constant necessarily involves being prepared to use it according to some inference rules and not others.

In its strongest form, such a theory is called an Inferential Role Semantics and says that it is *in virtue* of our using a constant, say 'if', according to some basic rule involving it, say MP, that 'if' means *if* in our idiolect.[1]

A weaker doctrine, which is all I will assume here, would simply have it that meaning *if* by 'if' *requires* inferring according to MP, without necessarily being sufficient for it.

As I say, many philosophers, among whom we may number Michael Dummett, Robert Brandom, Paul Horwich, Ned Block, Stephen Schiffer and Christopher Peacocke, have been partial to some version or other of an inferential role semantics. Even one of the harshest critics of this style of meaning theory, Jerry Fodor, has always maintained that when it came to the case of the logical constants, no other style of theory seemed to be in the running – certainly not causal or teleological or definitional theories.

Accordingly, I got interested in the question: Suppose we assume that

- (A) Inferring according to MP is necessary for someone to mean *if* by 'if'.

could we make it plausible that:

- (B) We are *blindly entitled* to infer according to MP.

I did not take it upon myself to argue for (A). I followed in the footsteps of the philosophers listed above and simply assumed (A). My main task was to try to show that if (A) is true (B) is true. In a series of papers, I experimented with a number of different ways of arguing for this conditional, none of which I am fully satisfied with.

3.

My critics, however, Paolo and Tim included, have largely concentrated not on the arguments I

provided for the conditional 'if (A), then (B)', but rather on the inferentialist assumption (A) itself.

This has thrust me into the role of defender of an inferential role account of the meaning of the logical constants, although, as I say, my main focus was elsewhere. Nevertheless, the inferential role theory is assumed by my account and my critics have brought up many interesting points. I am therefore happy to discuss them.

Now, the first two thirds of Paolo's paper consists in some very interesting observations about why we should not take inference according to some rule to be *sufficient* for possessing a particular logical constant. I find some of that discussion to be very interesting, but, as Paolo realizes, it is not relevant to the sort of account that I was exploring, which depends only on the necessity claim.

Paolo does eventually turn his attention to the necessity claim, and to the account of entitlement that is built upon it, and he has a number of telling objections to make against them both.

4.

One objection that Paolo makes is that my story about entitlement is bound to be incomplete. Paolo claims that we are blindly entitled to many more inferences than could plausibly be said to be necessary for concept possession. When should we say that those non-concept-constituting inferences are also blameless? (When I say "concept-constituting" in this paper, I shall just mean "is necessary for possession of the concept in question".)

> To provide an answer to this question, the analysis should be substantially supplemented. On the other hand, once we have found a satisfactory account of blamelessness for blind inferences which are not instances of rules belonging to the possession conditions for the logical constants, why shouldn't we apply this account also to the blind inferences which are instances of those rules, so making the initial analysis superfluous? Unless, of course, what Boghossian intends to suggest is that only inferences which are instances of rules belonging to the possession conditions for logical constants can be at the same time blind and blameless. But this would be hard to maintain. Take Ramanujan, great Indian mathematician. We are told that he had an astonishing capacity to draw immediately very remote and complex consequences from given premises. He was often unable to justify those conclusions by means of what most mathematicians would have regarded as an acceptable proof; in fact, we are told that he had only a very vague notion of what a proof is. Well, I think it would be wrong to deny that Ramanujan's blind inferences were (at least in many cases) blameless, i.e. that they transferred knowledge. At the same time, it would make no sense to say that being able to perform inferences like those of Ramanujan is part of the possession conditions for logical constants, i.e. that it is necessary in order to know what they mean. (406)

What I would say in response is that, in the sense I have in mind, it is plausible that Ramanujan *was* in a position to offer some justification for his feats of inference, even if he was not in a position to provide a rigorous proof of them. The kind of circumstance for which I coined the notion of "blind entitlement" was for the case of a logical rule so basic that no person, no matter how well informed and how good at rigorous proof, *could* provide any sort of justification for the use of the rule, because of the inevitable circularity that such a justification would entrain.

So, I do not see that we have in Paolo's description a clear example of a rule for which we are in principle not in a position to supply a justification, yet to which we are clearly entitled, and which could not plausibly be claimed to be concept-constituting.

5.

Paolo's next point, however, is more worrisome. He maintains that there is a general *recipe* for generating a counterexample to any claim of the form: Reasoning according to rule *R* is required in order to have logical constant *C*.

The idea that there is such a recipe strikingly anticipates an argumentative strategy developed by Timothy Williamson in his recent book ([Williamson 2007](#)), although the recipes that Paolo and Tim have in mind are different from one another (and were no doubt developed independently of one another).

Before looking at these recipes in greater detail, let me comment on the importance of the claim that there is a *recipe* of this kind, a general method for generating a counterexample to any particular concept constitution claim.

The point is that there is, even among the friends of inferentialism, considerable uncertainty about exactly *which* rules are meaning constituting for *which* constant. Of course, this is not meant to be a virtue and is raised by Paolo, as we shall see below, as a point of criticism of inferential theories. This uncertainty, though, can serve as a natural defense against proffered counterexamples. Faced with a counterexample to a particular concept constitution claim, the inferential theorist can always distance himself from that particular claim, while clinging to the claim that some rule or other will be constitutive.

Vann McGee, for example, has developed what he takes to be a set of counterexamples to *Modus Ponens* (see [McGee 1985](#)). These examples all involve cases in which a conditional is embedded in another conditional. Does this show that MP is not necessary for possession of *if?* It may be hard to answer this question precisely because we may not be sure if assent to *all* instances of MP is necessary for possession of *if*. Perhaps it is enough that a thinker assents to all instances of MP that do not involve conditionals embedded in bigger conditionals? Perhaps MP is not involved in the possession of *if* at all. When the focus is, as mine was, on the conditional.

> If *R* is concept constituting, then we are blindly entitled to *R*,

it can seem a matter of indifference that someone has come up with a counterexample to any particular concept constitution claim.

However, if someone can show that there is a general *recipe* for generating a counterexample to *any* pair of *C* and *R* such that reasoning with *R* is held to be necessary for *C*, then clearly that goes to the heart of inferentialism and of any epistemology that might be built upon it.

Here, then is Paolo's recipe for generating such counterexamples:

> Apart from this difficulty, the idea that, given a logical constant *C*, there is a well-defined set R of rules of inference such that a subject cannot be regarded as knowing the meaning of *C* unless she accepts the rules in R is intrinsically problematic. No matter how a rule of inference is chosen, it seems to me that we can imagine situations in which we would be disposed to say that a subject knows the meaning of *C* although the subject does not accept the rule in question. Suppose Mary suffers from a cognitive disability which makes her completely incapable of performing anything which could be counted as an inference. Nevertheless, she is able to use logically complex sentences to describe visually presented scenes. (I do not know whether this kind of cognitive disability has ever been observed; but it is no doubt conceivable.) For example, we can imagine that Mary, although unable to perform conjunction-introductions and conjunction-eliminations, would be able to assert, in appropriate circumstances, "The box is red *and* the book is blue", or things like that. It

seems to me that it would then be possible to say that, in spite of her disability, Mary knows the meaning of the word "and". At least: my intuition is that we would spontaneously adopt a homophonic translation for sentences such as "The box is red and the book is blue" uttered by Mary; and wouldn't that be a way of acknowledging that, in Mary's mouth, the word "and" has the same meaning it is has for us? (407)

So, the way Paolo's recipe is supposed to work is that, for any rule, R, and constant, C, we can cook up a counterexample to the claim that inferring according to R is required for possession of C, by imagining someone who lacks the ability to infer with R but who, we might make plausible, possesses C because she uses it competently in sentences held true.

I am not convinced that there is a good recipe here for generating counterexamples to inferentialism. There are two ways in which we can develop the Mary example between which Paolo's description does not distinguish and, however we develop the example, I do not see that we get a convincing counterexample to the necessity of either conjunction elimination or conjunction introduction for possession of *and*.

On the first way of developing the example, we may claim that Mary can think each of the atomic sentences "The box is red" and "The book is blue" separately, reliably asserting the first in the presence of a red box and the second in the presence of a blue book. Furthermore, we may claim that she can also think the compound sentence "The box is red and the book is blue", reliably asserting it in the presence of a state of affairs that contains both a red box and a blue book. What, then, is she lacking?

What we are told is that she cannot *infer* from the atomic sentences taken together to the compound sentence, or from the compound sentence to either of the atomics.

Confronted with a red box and blue book she is prepared to assert "The box is red and the book is blue", but if you ask her right after she has asserted that conjunction (perhaps this has to take place in a different room), "So, is the box red?" she might say "No". Similarly, for a question about the book.

Having shown Mary a red box and gotten her to assert "The box is red", and having shown her a blue book in a different room and gotten her to assert "The book is blue", she then refuses to assent to "The box is red and the book is blue", although, by hypothesis she is willing to assent to the compound sentence when she is shown the box and book together.

I find all this mystifying and certainly do not feel inclined to say: Clearly, Mary retains the ordinary concept of conjunction that ordinary people have, even as she fails to make various inferences that ordinary people would make.

We can also flesh out Paolo's example by stipulating that Mary can think only the compound sentence and cannot think the atomic ones separately. This, I think, would make it even less plausible that Mary means *and* by 'and'.

6.

Tim's recipe for generating a counterexample to understanding-assent links is interestingly different from Paolo's. Paolo's example turns on a *disability*. Tim goes in the other direction. His counterexamples consist of *experts* on logic and language who have allowed their linguistic behavior to be influenced by the somewhat kooky theories of logic and language that they have developed as adults. In some ways, this can seem a more promising strategy. It can seem easier to make it plausible that someone has retained a concept C by making him an expert on C than by giving him a disability with respect to C.

Here is how the recipe is supposed to work. Take any constant *C* and any rule *R*. Suppose it is maintained that *T*'s inferring according to the rule *R* is required for *T* to have *C*. Then we can always describe a case of an expert on *C* who becomes convinced, however incorrectly, by a complex theoretical argument, that *R* is invalid and so refuses to infer according to it but who, by any ordinary standards, still fully understands *C*. So there can be no *R* such that inferring according to it is necessary for *T* to have *C*.

Tim, too, uses his recipe to generate a putative counterexample to conjunction-elimination. He describes the case of Simon, an expert on the philosophy of language who has views on vagueness. Simon holds that borderline cases constitute truth–value gaps. He generalizes classical two-valued semantics by treating the gap as a third value and by conforming his practice to Kleene's weak three-valued tables. According to these tables, a conjunction is indefinite (neither true nor false) if at least one conjunct is, irrespective of the value of the other conjunct. Furthermore, Simon regards truth and indefiniteness as designated (acceptable) semantic values for an assertion: what matters to him is to avoid falsity. So he accepts sentences that are either true or indefinite.

It is easy to see that someone with Simon's semantic commitments would have reason to reject conjunction elimination as a rule of inference, for there could be cases where 'A' is simply false while 'B' is indefinite. In such cases 'A and B' would be indefinite, but 'A' false. Thus, the corresponding instance of conjunction elimination would have a designated premise and an undesignated conclusion, and so Simon would reject it. This, then, is the basis for Tim's confidence that not even something as seemingly safe as conjunction elimination is required for meaning *and* by 'and'.

Now, I find Tim's examples as hard to understand as Paolo's, although they raise different issues. I do not believe that Simon presents us with an intelligible counterexample to the analyticity of conjunction elimination; and I do not believe that Tim has provided us with a general recipe for dispatching any understanding–assent link that might be proposed.

To get a sense of how puzzling Tim's Simon would be, imagine that Simon has come to the view that someone other than John Wilkes Booth shot Lincoln. According to him, Booth had a co-conspirator, Schmidt, who was actually responsible for pulling the trigger. Both men were there, in Lincoln's box at Ford's theater, but it was Schmidt that shot Lincoln, not Booth. So Simon assents to:

>   (0) Schmidt, not Booth, shot Lincoln.

However, Simon is very willing to assent to the sentence

>   (1) Booth saw the balding Lincoln and shot him.

since he takes Booth to have been there and seen Lincoln, and, since he regards the first conjunct as indefinite, he regards the whole sentence as indefinite and so acceptable.

He is not willing to assent to:

>   (2) Booth shot Lincoln.

In fact he rejects (2). If we continue to spin the case out, we would have to say that Simon is also willing to assent to:

>   (3) Booth didn't see the balding Lincoln and shot him.

He would also assent to:

>   (4) Booth saw the balding Lincoln and didn't shoot him.

Also to:

(5) Booth didn't see the balding Lincoln and didn't shoot him.

As well as to:

(6) Lincoln is both bald and not bald.

When I look at the description of this case, I find myself with no clear intuitions about what Simon is saying or thinking. I certainly do not think: "Oh, he clearly means conjunction by 'and' ".

Tim has replied to this by saying:

> Obviously, if we have just met Simon, and know nothing about his background beliefs, we are likely to find his combined reactions to (1) and (2) utterly bewildering. We may reasonably wonder whether he knows what the word 'and' means. In practice, independently of his reaction to (1), since it is so well known that Booth shot Lincoln we may also find Simon's rejection of (2) initially puzzling, and wonder whether he is using the name 'Booth' to refer to the man we mean. Once we become aware of Simon's conspiracy theory of the assassination, we realize that there was no linguistic misunderstanding over (2); we simply disagree with him about the historical facts. Similarly, once we become aware of Simon's deviant theory of logic, an explanation of his unwillingness to deduce (2) from (1) in terms of linguistic incompetence looks much less attractive. On theoretical grounds, Simon holds that borderline cases for vague terms induce truth-value gaps, and that such gaps should be treated by Kleene's weak three-valued tables, which coincide with the classical two-valued tables when all the constituent sub-sentences are true or false but make the complex sentence gappy when at least one sub-sentence is gappy. Simon also thinks that it is legitimate to assent to gappy sentences as well as to true ones; what matters is to avoid falsity. Since he thinks that Booth saw Lincoln and regards Lincoln as a borderline case for the vague term 'bald', he thinks that 'Booth saw the balding Lincoln' is gappy, and that (1) inherits its gappiness. He concludes that it is legitimate to assent to (1). The gappiness does not infect (2). Simon rejects (2) as straightforwardly false.
>
> Of course, Simon would be quick to point out that in conversational terms it would be highly misleading to assert (1) on grounds of its gappiness when one's audience had no reason to suspect that one was doing so. In the absence of special background assumptions, asserting 'A(P)' leaves it open whether 'A(P)' is true or gappy, on Simon's view. If one knows that 'A(P)' is gappy because it has the gappy constituent 'P', one can therefore make a simpler and more informative assertion by simply asserting that 'P' is gappy, omitting the other material in 'A(P)' as irrelevant. On Simon's view, one can gain the effect of asserting that 'P' is gappy without going meta-linguistic by asserting 'P and not P'. Thus if Simon asserts (1), his audience is entitled for Gricean reasons to assume that he is not doing so merely on the grounds that 'Lincoln was bald' is gappy, since otherwise he is being conversationally uncooperative and should have said something like 'Was Lincoln bald? Well, he was and he wasn't' instead. The default conversational assumption is that one is not dealing with borderline cases; under that assumption one can defeasibly move from 'P and Q' to 'P' and to 'Q'. Nevertheless, according to Simon, the move is not deductively valid, and the case of (1) and (2) is a counterexample.
>
> Once Simon has explained his view, it is much less plausible that his unwillingness to infer

(2) from (1) manifests linguistic incompetence. It looks much more like a case of theoretical disagreement. ([Williamson 2011](), 502)

Tim raises many interesting points and there is a huge amount to be said in reply. Here I have space only to make a start.

First, a small terminological point. To call someone who is as sophisticated about logic and language as Simon is, "linguistically incompetent", would be obviously misleading, just as it would be to so label Vann McGee for doubting *Modus Ponens*. All that the inferential role theorist is committed to saying is that, if Simon succeeds in altering his behavior with 'and' and flouts a meaning-constituting rule for ordinary conjunction, then he necessarily means something different by 'and' than ordinary conjunction. It is better to call this "meaning change" rather than incompetence.

Second, it might not be such a big meaning change (assuming we know how to measure such things). The new concept might play many of the same roles we associate with ordinary conjunction. It just would not be ordinary conjunction.

(Can we always rely on there being a sharp fact of the matter whether there has or there has not been meaning change? It would be surprising if meaning facts were more determinate than facts in other domains, so a certain amount of indeterminacy about meaning change would have to be allowed for as well.)

Third, Tim says that once we know about Simon's deviant theory of logic, the explanation in terms of change of meaning "looks much less attractive". I disagree with this assessment. There are two large reasons. First, in deciding whether Simon is best described as expressing one meaning versus another by 'and', we cannot rely on the fact that he has arrived at his inferential role for 'and' on the basis of theorizing. Second, I believe that explanations in terms of meaning change, rather than theory change, can sometimes be the most attractive. I will develop each point in turn.

When you look at formulations of inferential role semantics, you find that theorists want to identify the concept-constituting inferences with those that are "primitively compelling" (Peacocke) or that incorporate "underived conceptual roles" (Schiffer). The inferences that are said to be concept-constituting for a thinker are those that the thinker finds compelling, is willing to engage in, without the benefit of any prior theory. As Peacocke puts it in connection with possession of the concept of conjunction: "On any theory, this possession-condition will entail that thinkers must find the transition from A and B to A compelling, and must do so without relying on any background information" ([Peacocke 2004](), 172; quoted in [Williamson 2007](), 125).

This is obviously a very important feature of inferential role theories. It would make no sense to identify a meaning-constituting inferential rule for a constant with a *derived* rule for that constant, arrived at on the basis of rationally optional theorizing involving inferences with that very constant. Such a procedure can lead you to all sorts of mistaken views about what the meaning-constituting inferences for that constant are.

Simon's deviant inferences, however, are obviously highly derived. He does not find them primitively compelling, but compelling only on the basis of lots of (bad) theorizing. No inferential role theorist would look to those derived inferences to say what concept Simon expresses by 'and'. They would look, rather, to the rules that Simon found primitively compelling, before engaging in all of that bad theorizing. By assumption, those rules are just the standard ones. So far, then, we have not yet got a counterexample to the necessity of conjunction elimination for possession of conjunction.

To get one, we would have to argue that even a non-theoretically minded analogue of Simon's, who did not have fancy views about vagueness and gappiness, but who exhibited the same pattern of behavior

with 'and' as has been stipulated for Simon, would clearly be credited with possession of ordinary conjunction.

I do not believe that many would be sympathetic to such a verdict.[2]

7.

Does not the mere fact that it is possible for Simon to intelligibly question whether conjunction elimination is valid, however, show that conjunction elimination is not meaning-constituting? This brings me to the second large point I signaled above. I believe that the inferential role theorist can explain what Simon is up to in a way that is consistent with the theorist's commitments.

Consider a different case. I think it is very plausible that our pre-Einsteinian ancestors worked with a notion of simultaneity of which it was analytic that it denoted a 2-place relation. They would not have understood how simultaneity – or time order more generally – could be relative to an observer's frame of reference, and so a 3-place relation.

Einstein, however, came along and claimed just that. Here we can mimic something that Tim might want to say: Surely, Einstein was not just committing some linguistic mistake. Surely, once we know about his Special Theory of Relativity, an explanation of his unwillingness to deduce "x stands in a 2-place relation to y" from "x is simultaneous with y" in terms of a change in the meaning of 'simultaneous' becomes much less attractive.

I do not agree that an explanation that invokes a change of meaning is far less attractive. At any rate, I think that there is a perfectly good story in terms of change of meaning that can preserve many of the features of the case that seem worth preserving. (Grice and Strawson made this point some time ago.)

On the story I have in mind, Einstein is proposing that we get a better theory of motion if we work with 3-place simultaneity relations rather than with 2-place ones – that is, if the explanatory role for which we need a notion of simultaneity is filled by a particular kind of 3-place relation, rather than the classical 2-place one. This accommodates the Einsteinian achievement without having to deny that 2-placedness was constitutive of the classical notion. Similarly, we do not have to deny that the Parallels Postulate is constitutive of Euclidean space just because we recognize that the best theory of physical space may involve spaces that are non-Euclidean.

Hence, I do not think it is true that once we become aware of someone's substantive reasons for preferring one theory of 'X' over another, that we can no longer think of the disagreement as involving a change in the meaning of 'X'.

It might be thought that in helping ourselves to the notion of meaning change, as opposed to mere change in belief, we are begging the question against Quine. Two points: first, the Quinean claim that there are no determinate facts about meaning change has yet to be earned; second, Paolo and Tim do not express the same general skepticism about determinate meaning facts as Quine does. Tim, in particular, is quite clear that he believes that there are determinate facts about meaning. It is just that he does not think they are constituted by facts about inference rules.

At this point, we come face to face with another question: perhaps we do not *have* to think that there has been no change in meaning. How do we know whether there has been one? How do we know which inference rule is, and which inference rule is not, constitutive of someone's having concept C?

8. Paolo presses this question, too

> Should we say that the acceptance of a logical rule is part of the possession condition for some logical constant if we find its instances easy and natural, if we apply it as a matter of

course and irreflectively, and if we expect that anybody else would do the same? Unfortunately, a criterion of this kind is unlikely to work. What inferences different people find easy and natural varies greatly. Presumably, what was easy and natural for Gödel in the early thirties was not what is easy and natural for me now. One might try to overcome this problem by saying that the logical rules to be taken into account are the rules that are found easy and natural by everybody, or almost everybody. But this too would give rise to obvious difficulties. First: being easy and natural is a matter of degree. How easy and how natural should a rule be to be taken into account? At least at first sight, there is no principled way to draw the line. For example: for most people, Modus Tollens is slightly less easy and less natural than Modus Ponens. But only slightly. Does this slight difference matter? Is acceptance of Modus Tollens necessary to have the concepts of conditional and negation? Second: not only do different people find different inferences easy and natural, but the same person may apply a certain rule with great ease when she is reasoning on a certain topic (holidays, for example) and have a lot of trouble with it when she is reasoning on some other topic (abstract algebra, say). Being easy and natural are topic-dependent. What topics are to be regarded as relevant? Again, no principled answer seems possible. Third: let us imagine we have settled in some way the previous two problems. We might discover that the inferences which we have decided to count as easy and natural for everybody do not suffice to determine uniquely the denotation of logical constants. For example, they might be insufficient to establish whether "or" stands for inclusive or exclusive disjunction, or which truth-value should be assigned to "Every *P* is *Q*" when there are no *P*. For Peacocke and Boghossian this would be a problem for the reason explained at the outset. (408)

Paolo insists that the question here is *not* epistemic.

Note that the problem is *not* that of actually deciding whether this or that specific rule is or is not constitutive of some logical concept. The problem with the syllogism in Barbara and Modus Ponens is *not* that we do not have all the data required to establish whether they are or are not constitutive of the logical concepts they involve. The problem is that we have not been told what sort of data would be relevant to establish this. Notice also that the problem is not one of vagueness. The case of Modus Tollens is not like the case of a man who is neither clearly bald nor clearly not bald. To establish whether a man is bald or not, I know that I must apply a certain criterion: then it may happen that, in a particular case, the criterion does not give a clear verdict and I remain uncertain. In the case of Modus Tollens, I have simply no idea of what the relevant criterion might be like. (409)

What I want to say in reply is that, in an important sense, the only legitimate residual question in this vicinity *is* epistemic.

We have already been told *what* we are looking for. We are looking for those conditions that are necessary for having concept *C*. If we are working within an inferential framework, we are asking which inference rules are necessary for possessing concept *C*.

Paolo says: "we have not been told what sort of data would be relevant to establish this". We have "no idea what the relevant criterion might be like".

Why, in addition to being told that we are looking to find out which inference rules are necessary for possession of *C*, must there be criteria by which this matter is to be decided?

It is, indeed, very common in discussions of this topic, to think that there have to be behavioral markers

by which concept-constituting inferences are to be recognized. I think it is a mistake to look for such markers.

The idea that there have to be such behavioral markers is encouraged by reflection on the case where a word is introduced via *explicit stipulation*. Suppose I introduce the word "flurg" as follows:

(7) By "flurg" I shall mean: "Any murder committed on a Tuesday".

If I am rational, then, as a mere result of this stipulation, I will exhibit a certain kind of assent behavior. For example, I will assent to

(8) A flurg always occurs on a Tuesday.

and I will do so without the benefit of any empirical evidence. I will not regard any empirical evidence as bearing upon its acceptability in any direct way, and so forth. Such assent behavior can be regarded as 'criterial' for (7)'s having a meaning-constituting status for me.

Now, of course, in the case of the basic logical constants, there is no question of having introduced them via explicit stipulation (one would need some constants in order to make any stipulations). When we think of an inference rule as having concept-constituting status for a particular constant, there is a natural tendency to think of it as a sort of *tacit* analogue of an explicit stipulation. Then it becomes natural to ask how to identify the *tacit analogues* of the assent behavior that's criterial in the explicit case.

There is a mistake here. The reason that we may expect my characteristic assent behavior in the explicit case depends on the fact that, in such a case, I *know* my definition and *know* that it has definitional status for me. In the explicit case, all such facts are open to view.

However, matters are different in the tacit case. A rule R can be concept-constituting for C in S's idiolect without S knowing that it is. As a result, S can come rationally to question R on the sorts of highly theoretical grounds that Tim describes.

Such questioning by S need not mean that R is not concept-constituting for S's having C, but it does mean that ordinary speakers, who are not trained to think about such matters, can change their concepts without knowing that they have.

How can we tell whether that has happened? In the usual way, clearly being used both by Paolo and Tim – via intuitive judgments about possible cases.*

Footnotes

- 1

  Such theories are best run over a mental language rather than a public language. I will not worry about this distinction in the present paper.

- 2

  This is especially true if one finds it natural, as I do, to operate with an idiolectic rather than a social conception of meaning, although I do not think that the force of my point rests exclusively on that important divide.

- *

  This paper was presented at a conference in honor of the late Paolo Casalegno in Milan in April

Boghossian, P. 1996, 'Analyticity Reconsidered', Nous 30, pp. 360–391.

Boghossian, P. 2003, 'Blind Reasoning', Proceedings of the Aristotelian Society 77 (Suppl.), pp. 225–248.

Casalegno, P. 2004, 'Logical Concepts and Logical Inferences', Dialectica 58, pp. 395–411.

Harman, G. 1986, Change in View, Cambridge, MA: MIT Press.

McGee, V. 1985, 'A Counterexample to Modus Ponens', The Journal of Philosophy 82, pp. 462–471.

Peacocke, C. 2004, The Realm of Reason, Oxford: Oxford University Press.

Williamson, T. 2007, The Philosophy of Philosophy, Oxford: Blackwell.

Williamson, T. 2011, 'Reply to Boghossian', Philosophy and Phenomenological Research 82, pp. 498–506.