Philos Stud (2014) 171:161–175
DOI 10.1007/s11098-013-0255-y

Metaphysical analyticity and the epistemology of logic

Gillian K. Russell

Published online: 11 December 2013

Abstract Recent work on analyticity distinguishes two kinds, metaphysical and epistemic. This paper argues that the distinction allows for a new view in the philosophy of logic according to which the claims of logic are metaphysically analytic and have distinctive modal profiles, even though their epistemology is holist and in many ways rather Quinean. It is argued that such a view combines some of the more attractive aspects of the Carnapian and Quinean approaches to logic, whilst avoiding some famous problems.

Keywords Epistemology of logic · Metaphysical analyticity · Truth in virtue of meaning · Two dogmas · Web of belief · Regress argument

1 Introduction

This paper argues for a view in the philosophy of logic. My title suggests that it is a view in the *epistemology* of logic, and so it is, but what is really striking about it is not the epistemology as such, but the way in which the epistemology is related to certain metaphysical and semantic features, such as necessity and analyticity. I will introduce the view by contrasting it with two well-known ones, associated with Carnap and Quine respectively. On the first, logic is analytic and this accounts for both its apriority and its modal status.[1] On the second, logic is not analytic, not

[1] I will assume throughout that the logical properties hold of (sets of) sentences and sentence schemas. Since necessity is usually attributed to propositions, it isn't quite accurate to say that the logical truths are necessary. Rather, they have a distinctive modal profile and in simple cases this amounts to having the

G. K. Russell (✉)
Department of Philosophy, Washington University in St Louis, One Brookings Drive,
St Louis, MO 63130, USA
e-mail: grussell@wustl.edu

necessary, and its epistemology is like epistemology elsewhere: claims are rationally given up or adopted depending on the role they play in an overall successful theory of the world.

On both the above views, analyticity, necessity and a priority stand or fall together: logic has them all, or it has none. But on the view I will recommend they come apart. Logic is analytic, and the analyticity explains its distinctive modal status. But the *epistemology* of logic is holist, rather Quinean, and not exclusively dependent on the epistemology of meaning, on semantic competence, or on acts of implicit definition.

Though such a view is an option once we have distinguished the three properties, it might not initially seem an attractive one. If one can establish that logic is analytic, the accompanying epistemology is supposed to be a *benefit*. Why fight for logic's analyticity only to reject the epistemological help it offers? Moreover, many will assume that if logic is analytic, then its epistemology is settled. Analytic truths are meant to be obvious, or undeniable, or perhaps justified for anyone who understands them. So how could their epistemology be holist?

What makes space for the view is a distinction of more recent provenance, between metaphysical and epistemic analyticity. (Boghossian 1996; Williamson 2008) A sentence is said to be *metaphysically analytic* if and only if it is true in virtue of meaning. By contrast a sentence is *epistemically analytic* if and only if anyone who understands it is justified in taking it to be true.[2] Elsewhere I have argued for the coherence of metaphysical analyticity and shown that on my theory of the property, sentences may be metaphysically analytic—true in virtue of meaning—even though competent speakers who understand them have no reason at all to think them true.[3]

But even independently of my account, there are ways to see that sentences could be true in virtue of meaning and yet not epistemically accessible to speakers who understand them. The first reason is the sheer difficulty of establishing semantic facts. Once we get beyond the most basic stages, investigating the meaning of words like 'all', 'the', 'water', 'Hesperus' and 'I', and determining whether 'snow is white' must be either true or false (or whether it could in some circumstances have no truth-value), requires the appraisal of sophisticated rival theories. Often such theories have areas in which they perform well and areas where they do not perform

Footnote 1 continued
property of expressing a necessary truth. Similarly, unsatisfiable sentences express impossible truths, and valid arguments have the property of being necessarily truth-preserving. Schemas don't express propositions, but they may have a distinctive, derivative modal profile: $\phi \vee \neg\phi$ will express a necessary truth whatever is substituted for ϕ. Things are complicated further by logics for indexicals (Kaplan 1989; Russel 2012), where we have logical truths like 'I am here now' which do not express necessities, but such sentences have a distinctive modal profile of their own, namely, being such that they express a truth in all contexts. Carnap and Ayer tasked themselves with explaining the "necessity" of logic. I think we now have a much finer picture of modality in logic, and so I will refer more clumsily to "distinctive modal profiles" rather than necessity, but I see this as a more refined take on the same problem. It isn't too far from the truth to describe me as "worried about the necessity of logic."

[2] In fact there are many variants on this definition, i.e. see chapter 4 of (Williamson 2008).

[3] E.g. see the example of "Ali is Clay" in (Russell 2008 or 2010, pp. 197–199).

so well. Establishing that one is superior to another requires demonstrating that it does better overall, taking into account performance on all the puzzles, evidence (including evidence from other languages), intuitions, or whatever other tests we can devise. The ability to understand a sentence by no means confers the knowledge, awareness and understanding of the puzzles, or the epistemic skill of assessing theories. So there is a gap between a sentence's being true in virtue of meaning, and someone who understands the sentence being able to tell that it is true.

Alternatively, consider that, given semantic externalism, there can be a divide between two kinds of meaning: (i) what a speaker has to know to understand an expression, and (ii) the criteria something in the world has to meet to fall in the extension of the expression. Since these things are different, a sentence could be true in virtue of meaning in the second sense of "meaning", without it being the case that someone who understands the sentence "knows what it means" in that second sense (though they would by definition "know what it means" in the first.) If so, then someone who understands the sentence needn't be in possession of the knowledge of meaning required to figure out that it must be true.

Hence logic could be metaphysically analytic without this settling questions about its epistemology. In this paper I will argue that this makes room for a view that combines the idea that logic is analytic, and that this accounts for its distinctive modal status, with the view that the epistemology of logic is holist. Section 2 presents the problems that any view in the epistemology of logic needs to solve or dissolve. Section 3 outlines the traditional analyticity-based approach, and points out two standard problems with it, while Sect. 4 outlines the Quinean view and some problems for *it*. Then Sect. 5 presents the new view and argues that it incorporates the most attractive features of the older views, while solving or avoiding the problems.

2 The epistemology of logic

The central question in the epistemology of logic is how we come to be justified in believing claims which attribute one of the logical properties—logical truth, logical consequence, equivalence, unsatisfiability etc.—to linguistic items like sentences and schemas. I will represent these kinds of claim with a double turnstile, as in[4]:

$$\vDash \phi \vee \neg\phi$$

$$\phi \mathrel{=\!\!|\!\vDash} \psi$$

$$\Gamma \nvDash$$

Such claims are metalinguistic, and they can also be made in ordinary English when we say that an argument is *valid*, that the premises *follow logically* from or *entail* the conclusion, or that the conclusion is a *logical consequence* of the premises. We

[4] As you can see, the left hand side may be empty, and the turnstile may be negated or reversed.

make them about sentences when we say things like 'snow is white or it is not the case that snow is white' is a logical truth.

The epistemology of such claims is problematic. One reason is that logical truths are thought to be necessary (and known to be such) and accounting for knowledge of necessity is itself a non-trivial epistemic problem. Another is that knowledge of the logical properties is thought to be apriori, and understanding how we can have apriori knowledge at all is a further non-trivial problem.

True, there are some famous attempts to demonstrate apriority in philosophy. Among these are Frege's work on arithmetic and Kant's argument concerning analytic judgements. But these proceed by deriving the claims in question from *logic*.[5] That approach would not work here, as it assumes that knowledge of logic is a priori, which—since our question is the epistemology of logic—would be the very thing we are trying to show.

In the background is a third problem: knowledge of logic is expected to be *basic*. When giving justifications for beliefs in other disciplines, it is common to assume that one already has justification for some beliefs. An engineer can argue that a bridge will stay up for 100 years using mathematics and physics, she is not expected to *also* show that mathematics and physics are correct. Similarly, justifications in physics often assume mathematics and logic, and justifications in mathematics assume logic. But when giving a justification for logic, to what are we allowed to appeal? We can't assume knowledge of engineering, since that would make our justifications in engineering circular, and we cannot assume knowledge of logic, since that too leads to circularity.

In truth, this is too fast. Some logical knowledge can be derived from more basic logical knowledge; it seems unlikely that knowledge of Peirce's law—$\vDash((\phi \rightarrow \psi) \rightarrow \phi) \rightarrow \phi$—counts as *basic* logical knowledge. But eventually we must be left with a set of basic logical truths. And it is tempting to think that this has to be bedrock. And that makes one wonder how a justification for the basic laws of logic is even possible. Where would we find our premises?[6]

[5] In Kant's case, from the principle of contradiction, "I have only to extract from it [my judgement], in accordance with the principle of contradiction, the required predicate, and in doing so can at the same time become conscious of the necessity of the judgement" (Kant 1999, A7/B11–A8/B12) and in Frege's case from 2nd order logic: "I hope in this monograph to have made it probable that arithmetic laws are analytic judgements, and therefore a priori." (Frege 1964, pp. 7–8 [§87–§89])

[6] Gentzen's work on natural deduction seems to help with the basicness problem. (Gentzen 1964) showed that a proof-system needn't always consist of a set of axioms and rules but can instead consist entirely of rules. This helps in two ways. First, it allows proofs of logical truths from no premises at all. Thus in a certain sense, a natural deduction system permits knowledge of logic 'from nothing.' That seems like a promising feature if we are looking for knowledge that is basic. And second, it helps because Gentzen's rules look like much better candidates for basic knowledge than the axioms of some Frege-Hilbert-style axiomatic proof-system. It is more plausible that in some sense, everyone who knows anything knows rules like these—even if only implicitly—than it is that in some sense, everyone knows Pierce's Law. But even with Gentzen's help, the basic knowledge problem remains. For how do we acquire knowledge of the rules, without which we would have no justification for the basic logical truths? We again find ourself with the problem concerning basic knowledge.

The presumed basicness of logic makes the problems of necessity and a priority more acute, because it is really basic apriority and basic knowledge of necessities which present the most serious difficulties: if there is to be a priori knowledge at all, it seems likely that some of it will be non-basic, i.e. obtained by a priori means from other things that we know a priori. But, again on the assumption that justification must be non-circular and must end somewhere, not all a priori knowledge can be derived—some of it must be basic. The fact that we expect knowledge of logic to be both a priori knowledge *and* basic knowledge suggests that it must be of this more troubling kind. (Harman 1999, p. 145) Similarly with necessity.[7]

3 The traditional analyticity-based approach

On one story, analyticity can solve all the problems. (Carnap 1937, 1958a,b; Ayer 1936) If logical truths are analytic, that means that they are true in virtue of meaning and so their meaning is sufficient for truth—they will be true no matter what the world is like—which is to say that they will be necessary. Moreover, (so the story goes) understanding a sentence is a matter of knowing what it means, so an analytic sentence has its truth determined by something that anyone who understands the sentence already knows about. It seems reasonable to hope that they will be able to figure out that the claim is true, and we might think that this is close enough to apriority as to make no serious difference.

It is not as obvious how analyticity can help with basicness. If anything, one might think that if logical truths are true in virtue of meaning, then *meaning* is more basic than logic. But here it is traditional to turn to the doctrine of implicit definition. The idea is that explicit, metalinguistic definition—as in 'Henceforth, let the name 'Sally' refer to this fish'—is not the only way to introduce a new expression to a language. One can also introduce one with a sentence in the object language, as in: 'Sally is my new fish.' If 'Sally' had no meaning prior to the utterance and it is clear from context that the speaker intends the sentence literally, then the best way to understand him is as intending to introduce 'Sally' as a name for the fish.

Now we could introduce logical constants in a similar way. Rather than saying explicitly: "Let '$\rightarrow$' express the binary truth-function which yields the value $\top$ just in case either the first argument is $\bot$ or the second argument is $\top$" instead we might think it is introduced via implicit definition, perhaps the assertion of something like:

$$\phi \rightarrow \phi$$

It is only later that we work out the consequences of such a definition for the meaning of '$\rightarrow$' (e.g. which function it expresses.) On this account, our knowledge

[7] One famous justification for the necessity of a claim is Kripke's argument concerning identity statements using (distinct) rigid designators in (Kripke 1980, p. 3) One of the premises in that argument is that $a = a$ is necessary, so that, epistemically speaking, the necessity of $a = b$ (where a and b are rigid designators for the same object) is non-basic, and derived from that of $a = a$. While this is a paradigm example of justifying necessity, we cannot take it as a model for justifying *basic* necessities.

of what '$\to$' means is actually derived from our knowledge that $\phi \to \phi$ is true. We knew the latter first, even though its truth depends metaphysically on the former. Hence, epistemically speaking, the logical claim would be the more basic.

3.1 Two problems for analyticity-based views

These days, analyticity is itself controversial, and one might complain that to add analyticity to the list of unusual properties that logic is supposed to have—a priority, necessity, basicness etc.—just adds a further problem. I have defended analyticity elsewhere (Russell 2008) and here I want to focus on two arguments against the analyticity of *logic* specifically: The Regress Argument (Quine 1936) and an argument from *The Philosophy of Philosophy* (Williamson 2008).

3.1.1 Problem 1: the regress argument

The Regress Argument is an argument against a very narrow target, namely the view that the basic logical constants get their meanings by implicit definition, with the result that certain sentences or rules containing them are analytic. Of course, this is exactly the view outlined above. Suppose, for the sake of concreteness, that our primitive logical constants are '$\forall$', '$\to$' and '$\neg$'. Quine's argument begins with the observation that there are an infinite number of logical truths containing these symbols, and similarly, an infinite number of rules of implication. Let's focus on rules since these seemed like our best hope. Suppose the rules used in the implicit definition of '$\to$' are standard $\to$-introduction and -elimination rules. There are an infinite number of instances of these rules and presumably each is supposed to be analytic. We cannot perform the baptism/stipulation by pointing to or describing each instance, because there are too many of them. So we will have to find some other way. An obvious approach is to identify the instances by their form, and say: *I hereby stipulate that any rule which takes the following* form *is truth-preserving:*

$$\begin{array}{|l} 1.\ \phi \to \psi \\ 2.\ \phi \\ \hline 3.\ \psi \end{array} \qquad\qquad \begin{array}{|l} \begin{array}{|l} 1.\ \phi \\ 2.\ \psi \end{array} \\ \hline 3.\ \phi \to \psi \end{array}$$

But how to do we justify taking a *particular* argument of this form to be truth-preserving? We note, first of all, that it is of that form, and then that if something is of that form then it is truth-preserving, thanks to the implicit definition. Then we draw the conclusion that the argument is truth-preserving. The problem is that this method of justification clearly *uses logic*—in particular the conditional 'if..then...' and modus ponens—to reach its conclusion, and so it cannot be the story about how we come to know *basic* logic. You need logic to get logic by this method. That is the regress problem.

3.1.2 Problem 2: Peter and the foxes

A further problem for traditional analyticity-based views can be found in (Williamson 2008). Analytic truths are thought to have a special epistemic status, for example one might accept one of the following:

(1) Anyone who understands an analytic truth accepts it.
(2) Anyone who understands an analytic truth is justified in accepting it.

Williamson argues that if analytic truth really provides such epistemic accessibility, then truths of logic are not analytic. His example uses the logical truth

(3) Every vixen is a vixen

and stars Peter, a man with two unusual views, one in linguistics, and one in socio-political zoology. The linguistic view is that claims of the form 'every F is a G,' entail the corresponding existential claim 'there are some Fs'. For example, on Peter's view, 'every giraffe is a mammal' entails 'there are some giraffes'. Of course, many people think that universally quantified conditionals have existential import of some kind, even if it is merely a matter of conversational implicature: if I were to say 'everyone who bought a ticket won a prize' you would be rather surprised if you learned later that no-one bought a ticket, and you might reasonably accuse me of attempting to mislead. But some would say that what I said was nonetheless strictly and literally true, just trivially so. Not Peter. He agrees that existential information is conveyed by universally quantified conditionals, but he locates the source of that information in the sentence's logical consequences, rather than in what is conversationally implicated. Perhaps his view is wrong, but it is not beyond the range of reasonable views on the topic for a non-expert.

His socio-political zoology is stranger. Peter has been reading a lot of conspiracy websites, and has become convinced that there are no foxes. So-called "foxes" were an invention of the Ministry of Defence who wanted to distract activists from the Iraq war by focusing their protests on fox-hunting instead. They have even gone so far as to create children's stories, nature films (using a lot of CGI) and arrange some drug-assisted hallucinations and false memories of foxes, in order to spread the belief that there are foxes, foxes who need your help.

Peter denies that (3) is true. He neither asserts it nor believes it, for, he reasons, if *every vixen is a vixen* were true, the existential claim that there are vixens would have to be true as well, and it isn't. Perhaps he is not even justified in believing or accepting that every vixen is a vixen; for even if there were, in some sense, a kind of semantic justification available to him in virtue of his linguistic understanding, it is unclear that this justification can survive corruption or undermining by his explicit views.

Williamson claims—plausibly, I think, though the view is controversial—that Peter understands (3) perfectly well. His false view about foxes does not threaten his understanding, since this is not a semantic view about the meaning of 'fox', but only one about how things are with foxes. His false view about the entailments of universalised conditionals is more likely to be seen as undermining his semantic competence with (3), but Williamson argues, and I agree, that it would be a mistake to see it this way. Peter's position is no different from that of many professional philosophers and linguists who hold false views about the meaning of words in English and "giving an incorrect theory of a meaning of a word is not the same as using the word with an idiosyncratic sense—linguists who work on the semantics of natural language often do the former without doing the latter." (Williamson 2008, p. 9) Giving a theoretical account of the meaning of 'every' is a difficult task and Peter has (let us assume) got it wrong. But his normal linguistic processing needn't be affected by this any more than his normal visual processing would be affected by his having a false theory of edge detection. Peter understands 'every vixen is a vixen', he just thinks it is false. Perhaps he lacks justification to accept it. But if he does, and if the principles like (1) or (2) above are correct, then 'every vixen is a vixen' is not analytic. So some logical truths are not analytic.

One final problem with the traditional approach based on analyticity is that, if it is right, disagreement in logic seems surprisingly *messy*. One attractive feature of logic is that questions can often be settled in a more decisive way than they can in other parts of philosophy; sometimes one proof is sufficient to settle an issue. But this isn't always so. A question in the dispute between classical and intuitionist logicians is whether the law of excluded middle is *really* a logical truth. But the classical logician cannot win the argument by presenting her proof, nor can the intuitionist logician win it by presenting their counterexample; the intuitionist rejects the validity of the classical logician's proof, and the classical logician denies that intuitionist (e.g. Kripke-style) countermodels are really such. Recognising this, both parties argue about foundational assumptions, such as the semantics of negation or the metaphysics of mathematics. In such disputes—even though what is at issue is whether or not a certain schema is a logical truth—things are every bit as messy as they are in the rest of philosophy. This does not sit especially well with the idea that the law of excluded middle is analytic on a traditional conception of analyticity. Shouldn't disagreements about analytic truth be easier to resolve?

4 Quine and the web of belief

In the face of these problems for analyticity-based views, one might be tempted by the rival empiricist view on which logical claims are simply nodes in the total web of our beliefs about the world (Quine 1951). When the entire web comes into conflict with experience, we adjust it as best we can but, in principle, any part of it could be given up, including beliefs in mathematics and logic. In practice, we should try to obey certain rational principles when we make changes: Quine explicitly mentions trying to keep the web simple overall, and trying to make changes as conservative as possible, but we could add other principles (perhaps overall unity and explanatory power are important too) and remain within the spirit of his approach.

Mathematics and logic are thought to be *central* to the web of belief. To say that they are central is to say that an adjustment to the web which did not retain such beliefs would tend to be non-conservative, or non-simple, so that it would be very unusual for an adjustment which complied with the principles to involve such a change. As a result, logic can seem as if it is immune to experience (though it is not really) and both necessary and unrevisable (though it is not really).

On Quine's view, logic is not really basic either. Rather, the principles that govern the revisions—conservativeness and simplicity—are non-logical things that we can use to justify overall theories, including claims in logic. Here and later it will be helpful to draw on Harman's distinction between implication and inference (Harman 1986). *Implication* or *entailment* is the central subject matter of logic. It is a relation on sets of sentences (or schemas) and sentences (or schemas), one which holds just in case the latter is a logical consequence of the former. A typical deliverance might be the principle of modus ponens: $\phi, \phi \rightarrow \psi$ *entails* ψ. Harman holds that implication has often been confused with *inference*, the study of what things a rational agent ought to believe, given what they already believe and any new data or arguments they have been presented with. The study of this second topic has been less thoroughly developed than that of logic, but perhaps a putative principle would be 'believe the simplest explanation of the data' or even 'if you believe that ϕ and also that $\phi \rightarrow \psi$, then you ought to believe ψ as well'—which is superficially similar to modus ponens. One feature of inference, as opposed to implication, is that it will sometimes counsel us to *give up* a belief in response to new evidence.

Even where the two weren't conflated, it has often been assumed that implication is a sub-discipline of inference—the part that deals with deductive reasoning—but it is really a different subject matter altogether. Logic has nothing to say about beliefs, oughts or reasons directly, only about sentences and truth—and it turns out to be quite difficult to formulate principles linking logic to principles of belief change. For example, it is not in general the case that if you believe the proposition expressed by ϕ and the proposition expressed by $\phi \rightarrow \psi$, then you ought to believe the one expressed by ψ, since ψ might be a contradiction, or you might already believe $\neg\psi$. (Harman 1986; Field 2009)[8]

4.1 Problems with the web of belief view

On Quine's view, the three problems I outlined for the epistemology of logic do not arise. Given the failings of the traditional, analyticity-based view, this is an important consideration in favour of Quine's view. But there are also some problems.

It can be hard to relinquish the thought that logical truths have distinctive modal characteristics, in part because this is embedded quite deeply in logical practice. For example, with the rule of necessitation in the proof systems for standard modal logics. Necessitation allows one to place logical truths within the scope of a '$\Box$'. Quine has a famous response to this: he says we should give up modal logic. (Quine 1966; Burgess 1997) But I see this—and I do not think I am alone in this—as a serious cost of the view. If there were a view which had the virtues of Quine's approach, but did *not* have such costs, that view would be better.

[8] Also Laura Celani's unpublished "Logical norms and epistemic paradoxes."

A second worry is this: while it is not quite true that mathematicians and logicians don't care whether their work is useful to empirical scientists—most people are happy to have extra arguments in support of their research getting funded, and their students getting jobs—they don't usually act like they regard themselves as the handmaidens of physicists and computer engineers. They pursue work in the areas they do because it is interesting in its own right. Quine might maintain that such motivations are bad, and that they should ultimately be motivated to help construct the best physical theory of the world, but in truth, it's not obvious that we should agree. What's wrong with being motivated to make discoveries in mathematics that isn't also wrong with being motivated to make discoveries about the insides of distant stars? Quine's view comes too far apart from the actual practices of formal scientists. It is surprisingly philosophy-first and insufficiently *logic*-first.

5 A new view

I am going to tell you a story about someone who comes to have a justified belief that the law of excluded middle, $\phi \vee \neg\phi$, is a logical truth. The story has four parts, but I hope you'll bear with me because it does a good job of illustrating the view I want to advocate, and because the story's reasonableness makes it plausible both that this is how people come to have beliefs in logic, and that beliefs so reached are justified.

Part 1: Pre-Logic We begin before our hero has studied any formal logic. If we were to ask her which claims were laws of logic, or which followed logically from others, it is not clear that she would have much to say, though she might repeat some lines from detective stories: "you can't prove a negative," "you can't prove that something doesn't exist" or "when you have eliminated the impossible, whatever remains, however improbable, must be the truth." (Conan Doyle 1890) If we were to ask her whether LEM is a logical truth, she might not understand the question. In conversation, when someone says "look, either it is or is isn't", she will often acquiesce, and sometimes she uses the gambit herself, but it's unclear that she thinks it a *logical* truth, as opposed to any other kind of truth. She also says "look, everything has a cause, even if we don't know what it is, things don't just *happen*." It is not clear that she distinguishes laws of logic from, say, laws of metaphysics. Should we say that she believes the law of excluded middle is a law of logic? I think probably not.

Part 2: Intro Logic This all changes when our hero takes a first class in formal logic. It is an introduction to classical first-order logic, and by the end of it she unequivocally believes that the LEM is a logical truth. But to put it this way is to miss much of what has gone on: it is not as if the professor put a list of the logical truths on the board, which the student learned. Rather, the professor taught her the *theory* of classical logic. Among other things she learned: a formal language and skills for translating between her natural language and that language, the definition of an *interpretation* for that language, definitions of logical truth, logical consequence and other logical properties in terms of such interpretations, and techniques for demonstrating that particular schemas and sentences are logical truths. For example, she can now prove that LEM is a logical truth as follows:

> Look, whatever ϕ is, we could be looking at an interpretation on which is true, or one on which it is false. Suppose it is true. Then, given the clause in the definition of truth for disjunction, $\phi \vee \neg\phi$ is true. Suppose it is false. Then given the clause for negation, $\neg\phi$ is true, and so again $\phi \vee \neg\phi$ is true. So either way, $\phi \vee \neg\phi$ is true. So it is a logical truth.

But this same theory and techniques also do more. They allow her to show that $\phi \rightarrow \phi$ and $\neg(\phi \wedge \neg\phi)$ are logical truths, that ψ is a logical consequence of $\{\phi, \phi \rightarrow \psi\}$, that $\phi \wedge \neg\phi$ is unsatisfiable and that $\neg(\phi \wedge \psi)$ is logically equivalent to $\neg\phi \vee \neg\psi$. The theory explains all these logical properties in a very unified and simple fashion. It also predicts that the properties hold in some cases that she wouldn't have expected; when first confronted with modus tollens, our hero thought it invalid, but the theory said otherwise and she has come to realise that it was right. It also gives her answers in complex cases which outrun her intuitions about logical properties.

To her surprise, logic is useful in other classes: sometimes her physics teacher proves conditionals by proving the contrapositive. This was confusing at first, but now she can follow the proof. It's also notable that the assumptions she is asked to make, in adopting classical logic, are minimal: that sentences have exactly one of two truth-values, that names have exactly one referent, that predicates have determinate extensions. Compared to some of the assumptions she is asked to make for theories in her other classes (that economic agents are perfectly rational and perfectly self-interested, that a plane is frictionless, that Achilles has a motivation which can be unearthed through Marxist analysis etc.) these assumptions seem so innocent that she barely notices them.

In short, the new theory is dramatically better than her disparate collection of dictums from detective stories. The student embraces it wholeheartedly, coming, in the process, to believe that the law of excluded middle is a logical truth.

Part 3: Heresies Our hero signs up for a new class. On the first day, the professor says that in this course he is going to show that classical logic is false, where that means that some of the schemas and sentences that it says are logical truths are not *really* logical truths. The student is stunned by this—she doesn't really see how classical logic could turn out to be false—and when she remarks on this to her friends from the previous class, they agree. "How could the LEM turn out not to be a logical truth, given that we *proved* it?" There would have to be something wrong with that proof. But all the proof rested on were some semantic assumptions about things like the truth-conditions of sentences containing '$\vee$'. Those were stipulated, how could they be wrong?

The professor turns out to challenge some of the minimal assumptions on which classical logic was based. Drawing on examples from vagueness, presupposition failure, reference failure and future contingency, he argues that some sentences can enter into logical relations, even though they are neither true nor false. Since classical logic does not allow for this possibility, it needs to be improved. He presents an alternative theory: instead of the usual two-valued truth-tables, they now draw tables with three values: true, false and undetermined. He shows them how to determine the value of a whole sentence based on the values of its parts, and he defines logical truth and the other logical properties in terms of the new three-valued interpretations: in particular a schema is a logical truth if and only if it is true on all interpretations. He argues that the LEM fails this test, and so it is not a logical truth.

The student is impressed. She studies the new theory, considers the examples and decides to give up classical logic for the new one (for definiteness let us say that it is strong Kleene logic, taking only *true* as designated.) She reasons as follows: the three-valued logic has all the virtues of classical logic: it explains the presence of the various logical properties, does so in a simple, unified fashion etc. but it *also* accounts for some difficult cases where classical logic says nothing. So on balance she thinks it better, and she switches, in the process losing her belief that the law of excluded middle is a logical truth.

One aspect of her new position turns out to be trying: the new logic is weaker than the old, and this means there are proofs that she can no longer accept. Sometimes this is frustrating in physics classes, or when she is doing metatheory. But she resigns herself to it; after all, she thinks, she had good reasons for making the switch and she ought to stick to her convictions.

Part 4: Regrets? Our hero has become more and more interested in logic, and more and more adept. Whilst researching for a paper, she reads a book that contains a new argument, also designed to make one give up classical logic for a three-valued logic, though this time the logic is Priest's **LP**. (Priest 1987) The argument goes by way of the unresolved paradoxes, and maintains that the best solution to these paradoxes is to accept that sentences may have a third truth-status: *both* true and false. As before, a replacement logic is proposed, only this time there are two designated values: True and Both.

The student feels the force of the arguments from paradox and she sees the similarity between the author's argument and the arguments that led her to adopt her own three-valued logic. She also feels some pressure to be consistent and treat like cases alike. But suppose she were indeed to accept that sentences may have the truth-status Both. Then, given *her* present views, she would need to accept that sentences may have any one of *four* truth-statuses: True, False, Undetermined and Both. And there would be two designated values: True and Both. She quickly realises that the resulting logic is not attractive at all, because it is very weak. A few months of working with a weak logic have sensitised her to the difficulties and sacrifice of elegance this entails, and she finds the idea of losing all of modus ponens, modus tollens, the LEM, PNC, disjunctive syllogism and double negation elimination intolerable. *That* is not an attractive view. There is one other way to treat the similar cases alike: she could reject both. Perhaps, she suspects, she moved too fast when she accepted the heretical arguments of her second professor. Perhaps she should have found some way to resist the move, by adjusting her theory of vagueness, or future contingents, or of the scope of logic. She has begun to feel as if her step away from classical logic was the first step on a slippery slope, one she now wishes to return to the top of. And so that is what she does: for these reasons she returns to classical logic, in the process regaining her belief that the LEM is a logical truth.

Many of the specifics of the story above are extraneous, so let me draw out the essentials. First, logical beliefs are given up as well as gained. This does not sit well with the idea that the only methodology in logic is deduction, since that would only allow us to add beliefs. Rather, the whole process reeks of inference, in Harman's sense.

Second, the overall virtues of logical theories were an important part of the justification for adopting or rejecting a theory. Here I have stressed simplicity, unification, elegance, strength, usefulness and explanatory power, but there could be others and I could be wrong about some of these. (Perhaps usefulness is not really a reason for *belief*?)

Third, once we got beyond the initial pretheoretic stages, beliefs about logical principles were accepted or rejected holistically, that is, based on their part in a theory that was justified *overall*. The student doesn't give up classical logic for three-valued logic by rejecting logical principles piecemeal until finally (surprise!) ending up with a strong Kleene logic. Rather, she eventually comes to think that the three-valued logic does a better job overall all, and switches her views regarding particular cases accordingly.

But, fourth, just because the overall justification of the LEM is holistic, this doesn't mean there is no place for atomistic proofs, even proofs based on claims about meaning; once we have adopted a particular theory—say classical logic—the LEM can be proved as the student proved it, based on the claims about the possible truth-values, and truth-clauses for the logical constants. Thus we are still able to explain some of the plausibility and attractiveness of Carnap's approach.

Fifth, this epistemological story is compatible with the necessity of the claims being justified. Something may be rationally revisable even if it expresses a necessary truth. This point is obvious for the standard examples of the necessary a posteriori, including sentences like 'water is H_2O' and 'Hesperus is Phosphorus.' Someone could live in a society with little developed chemistry or astronomy. Perhaps an early theory would suggest that Hesperus is Mars, or that water is XYZ. There might even be a time when those theories are the *best* theories and then someone might—with reason—reject the necessary claims (and perhaps later revise their opinion when a better theory, or less misleading evidence, came in.) But as soon as we accept that the justification of claims in logic is holistic and requires the weighing of theories, it seems clear that just the same thing could happen with the question of, say, whether or not the LEM is a logical truth. Hence we are not committed to the least palatable aspect of Quine's views in logic.

One might suspect however, that while it is possible for sentences like 'Hesperus is Phosphorus' to express necessities and still be rationally revisable, this is not a possibility in logic, perhaps because those claims are analytic, i.e. true in virtue of meaning. But this assumes the link between metaphysical and epistemic analyticity that we rejected earlier. Moreover, once one accepts that establishing claims about meaning—including claims about the possible truth-statuses, or the correct account of universal quantification—is itself an epistemically non-trivial task which requires us to weigh theories against evidence, it seems clear that our reasons for accepting or rejecting claims in logic can wax and wane with our reasons for accepting or rejecting claims in semantics. A reasonable person can hold a false—even a necessarily false—view in semantics, and so the same goes for logic. So we can agree with Williamson when he argues that logic is not epistemically analytic, since someone—like Peter—may understand a truth of logic—like 'All vixens are vixens'—without thereby being in a special epistemic position regarding the truth of that claim. Williamson's target is a view on which logic is *epistemically* analytic.

Seventh, while Quine took apriority to be lost with holism, this is not inevitable. It depends on what kinds of consideration one uses to assess a theory. Perhaps the best theory in logic is the one that provides the most elegant explanation of a set of a priori data. Is elegance a property that can be assessed independently of experience? I don't know, but *if* it is, and if the data is really a priori, then logic could retain its apriority in spite of epistemic holism. If, by contrast, usefulness in physics turned out to be key, then apriority would be lost. Ultimately, the thesis of this paper does not decide the question of apriority.[9]

Finally, on the picture above, logic is not basic. There is something outside of logic that we can use to provide justifications for our beliefs in logic, namely, *reasoning*, where this includes the kinds of principles that fall under the umbrella of inference to the best explanation. As a result, we do not need to rely on implicit definition of the basic logical constants and do not face the Regress Argument. It might be objected that this only "solves" a hard problem by relocating it: where once we wondered how we could possibly have a justification for our beliefs in logic, now we think we get them using inference to the best explanation and wonder how we could possibly have a justification for *that*. This, I think, is fair. But progress has still been made. Everyone has the problem of justifying the use of inference to the best explanation and other principles of belief revision, since no-one should think that we can justify them using deductive logic alone. But we used to also have another hard problem: the question of how we could ever be justified in our beliefs in *logic*. The proposed solution reduces the latter problem to the former, and this seems like progress.

Acknowledgments I'm grateful for the input from my commentors at the 2013 Pacific APA—Roberta Ballarin and Sinan Dogramaci—as well as to Richard Zach for organising the session. Thank you also to audiences there and at NIP in Aberdeen, Arché in St Andrews, Eiden in Edinburgh, and Logos in Barcelona, as well as at St Louis University, UMSL, and the Universities of Genova, Oslo, Miami, Missouri, Omaha, and British Columbia for helpful comments.

References

Ayer, A. J. (1990/1936). *Language, truth and logic*. Harmondsworth: Penguin Books.

Boghossian, P. A. (1996). Analyticity reconsidered. *Nous, 30*(3), 360–391.

Burgess, J. P. (1997). Quinus ab omni nævo vindicatus. *Canadian Journal of Philosophy, 23*, 25–65.

Carnap, R. (1937). *The Logical Syntax of Language*. London: Keagan Paul.

Carnap, R. (1958a). *Empiricism, semantics and ontology* (2nd ed., pp. 205–221). Chicago: The University of Chicago Press.

Carnap, R. (1958b). *Meaning and necessity: a study in semantics and modal logic* (2nd ed.). Chicago: University of Chicago Press.

Conan Doyle, S. A. (2005/1890). *The sign of the four*. Stillwell, KS: Digireads.com.

[9] I'm grateful to a discussion with Brit Brogaard for helping me to see this point.

Field, H. (2009). The normative role of logic. *Proceedings of the Aristotelian Society* (pp. 251–268).
Frege, G. (1964). The concept of number (excerpt from the *foundations of arithmetic*). In P. Benacerraf, & H. Putnam (Eds.), *Philosophy of mathematics: Selected readings* (pp. 85–112). Englewood Cliffs, NJ: Prentice-Hall Inc.
Gentzen, G. (1964). Investigations into logical deduction. *American Philosophical Quarterly, 1*(4), 288–306.
Harman, G. (1986). *Change in view*. Cambridge, MA: MIT.
Harman, G. (1999). Analyticity regained? In reasoning, meaning and mind, chapter 7. Oxford: Oxford University Press.
Kant, I. (1787/1999). *The critique of pure reason*. The Cambridge edition of the works of Immanuel Kant. Cambridge University Press, Cambridge. (Series editors) Paul Guyer and Allen Wood.
Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. New York: Oxford University Press.
Kripke, S. A. (1980). *Naming and necessity*. Oxford: Blackwell.
Priest, G. (1987). *In contradiction*. Dortrecht: Kluwer.
Quine, W. V. O. (1936). Truth by convention. In *The ways of paradox and other essays*, chapter 9, (pp. 70–99). New York: Random House.
Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review, 60*, 20–43.
Quine, W. V. O. (1966). *Three grades of modal involvement*, chapter 15. Cambridge, MA: Harvard University Press.
Russell, G. (2008). *Truth in virtue of meaning: a defence of the analytic/synthetic distinction*. Oxford: Oxford University Press.
Russell, G. (2010). Analyticity in externalist languages. In *New waves in philosophy of language*, (pp. 186–205). Basingstoke, UK: Palgrave MacMillan.
Russell, G. (2012). Lessons from the logic of demonstratives: what indexicality teaches us about logic and vice versa. In G. Restall, & G. Russell (eds.), *New waves in philosophical logic*. Basingstoke: Palgrave Macmillan.
Williamson, T. (2008). *The philosophy of philosophy*. Oxford: Blackwell.