

Logical Knowledge and Gettier Cases

Corine Besson

ABSTRACT: Knowledge of the basic rules of logic is often thought to be very distinctive, for it seems to be a case of non-inferential a priori knowledge. And so many philosophers take its source to be different from those of other types of knowledge, such as knowledge of empirical facts. The prominent account of knowledge of the basic rules of logic takes this source to be the understanding of logical expressions or concepts. On this account, that is, what explains that such knowledge is distinctive is that it is grounded on semantic or conceptual understanding. However, this paper shows that this cannot be the right sort of account of knowledge of the basic rules of logic because it is open to Gettier style counter-examples.

Introduction

A Gettier case is one in which one has justified true belief but fails to have knowledge. Such cases were first proposed by Edmund Gettier to show that the so-called ‘traditional’ or ‘tripartite’ definition of propositional knowledge, which is precisely that knowledge is justified true belief, is wrong.¹ Typically, one fails to have knowledge in such cases because there is an element of luck in the way in which the belief is formed which prevents it from counting as knowledge. My aim in this paper is to investigate Gettier cases for a priori knowledge of the basic rules of logic, and their significance for the epistemology of logic. More precisely, it is to show that the prominent account of knowledge of the basic rules of logic is open to Gettier cases. On this account, such knowledge is grounded on our understanding of logical expressions or concepts. This account, I call the ‘Understanding Account’. And so I will show how the possibility of constructing Gettier cases for a priori knowledge of the basic rules of logic refutes the Understanding Account.

Given how much literature has been generated by Gettier cases in other areas of epistemology, it is a surprising fact that there are virtually no discussions of such cases in the literature on the epistemology of logic, and on the epistemology of a priori knowledge, more generally. There are passing mentions of Gettier cases involving

¹ See E. Gettier, ‘Is Justified True Belief Knowledge?’, *Analysis*, 23 (1963), pp. 121-123.

inferential – rather than *basic* – logical or mathematical knowledge.² But there is no mention of such cases for basic or non-inferential logical knowledge, such as knowledge of the basic rules of logic – where, that is, such rules are not inferred using other rules.³ One possible explanation for this is that for some reason it is thought that in principle such cases are not possible for knowledge of a basic rule of logic: such knowledge is just immune from the sorts of vagaries that befall other types of knowledge. The reason for this cannot merely be that such knowledge is basic, since Gettier cases for non-inferential knowledge such as perceptual knowledge are widespread in the literature. Perhaps the reason connects with the fact that knowledge of the basic rules of logic is *both* basic and *a priori*. Another, more mundane, possible explanation for this absence is simply that, until fairly recently, discussions of logical and *a priori* knowledge had somewhat fallen out of fashion or that no account of such knowledge seemed to generate substantive debates. And it is certainly a virtue of the Understanding Account that it revived those debates. In any case, this paper shows that Gettier cases for knowledge of the basic rules of logic are possible, and that they have interesting consequences for the epistemology of logic – in particular for the Understanding Account.

The paper is organised as follows. In section I, I propose a Gettier case for *a priori* knowledge of a basic rule of logic. In section II, elements of the Understanding Account are outlined, and in section III it is shown how we have a conflict between my proposed Gettier case and the Understanding Account. Sections IV-IX consider several ways in which it could be argued that my Gettier case does not refute the Understanding Account. In particular sections IV-VII focus on the interaction between semantic or conceptual understanding, and knowledge of the basic rules of logic. Section VIII considers a possible weakening of the Understanding Account and finally section IX briefly considers the question of whether *a priori* knowledge of the basic rules of logic can be gained through testimony, since this is relevant both to the Understanding Account and to my proposed Gettier case.

² C. McGinn, 'The Concept of Knowledge', *Midwest Studies in Philosophy*, 9 (1984), pp. 529-554, gives examples of such cases.

³ Of course different systems of logic treat different rules as basic, and I am assuming here that the rules discussed in the paper – Conditional Proof and Modus Ponens – are treated as basic. But nothing substantive in the discussion hangs on the particular rules considered.

I. Brenda, Nate and the Rule (MP)

Suppose that Nate wants to learn the rules of logic. And suppose further that he does not know some logical vocabulary, for instance he does not know the word ‘if’ and does not have the concept of material implication. (If there are worries about how Nate could reason at all without knowing the word ‘if’ or having the concept of material implication, we can for instance suppose that, oddly, he got by with disjunctive syllogisms.) Brenda, who is a renowned expert in logic, agrees to teach him. Given that Nate does not know the word ‘if’, she intends to teach him first the rules for ‘if’ – that is to say the rules of Conditional Proof and Modus Ponens. (I assume here that ‘if’ is truth-functional, which is a standard thing to do in discussion concerning the nature of knowledge of the basic rules of logic.) As Brenda and Nate are lost on a desert island, the teaching is done orally; and to teach him, she plans on telling him the two rules and perhaps give him a few of their instances.

She starts with the rule of Conditional Proof (CP), which says that:

(CP) From an inference that Q from the assumption that P, one can infer that if P, then Q.⁴

The teaching goes well with (CP), which Nate understands just fine. But then, when she turns to the rule of Modus Ponens, Brenda gets tired and irritated. She decides to trick Nate and sets about teaching him a fallacy. She will not teach him the correct rule of Modus Ponens (MP), which says that:

(MP) From if P, then Q, and P, one may infer that Q.

An instance of (MP) is for example:

(MP_{day}) From if it is day, then it is light, and it is day, one may infer that it is light.

⁴ Here, nothing important here rests on this particular way of stating the rules (CP), (MP) and (AC), except for the fact that they should not be understood as metalinguistic or metapropositional – as being about expressions or concepts. To anticipate section II, defenders of the Understanding Account regard (CP) and (MP) as implicit definitions collectively defining the meaning of ‘if’; and this requires that ‘if’ is used rather than mentioned in the statements of those rules.

Rather than (MP), Brenda decides to teach Nate the incorrect rule known as ‘the fallacy of asserting the consequent’ (AC), which says that:

(AC) From if P, then Q, and Q, one may infer that P.

An instance of (AC) is:

(AC_{day}) From if it is day, then it is light, and it is light, one may infer that it is day.

Now, suppose that by sheer coincidence, just at the time at which Brenda is supposed to utter (MP), or one of its instances, there is a whirlwind of sorts between her mouth and Nate’s ear, which they both fail to detect. The effect of this whirlwind is that although Brenda utters the rule (AC), Nate hears the rule (MP). That is to say, although she utters the second occurrence of Q as the second conjoint of the main antecedent, and the second occurrence of P as the main consequent, as she gives samples of the incorrect rule (AC), Nate systematically hears them the other way round; i.e., he hears the second occurrence of P as the second conjoint of the main antecedent, and the second occurrence of Q as the main consequent, and thus he hears samples of the correct rule (MP). For instance, although she utters the second occurrence of ‘it is light’ before uttering the second occurrence ‘it is day’, in trying to convey (AC_{day}), Nate hears ‘it is day’ before hearing ‘it is light’ and thus hears (MP_{day}).

Intuitively, Nate acquires the justified true belief that from if P, then Q and P, one may infer that Q – that is to say, he acquires a justified true belief of rule (MP). But he does not know it. For the process through which he acquired this belief was not reliable. Nate was just lucky that the atmospheric conditions happened to be capricious in a favourable way at the time Brenda tried to deceive him; if this had not happened, he would not have acquired a justified true belief.⁵ If that is so, this example shows that one can have a

⁵ An issue arises here about how to describe the case in which Nate is unlucky – the case in which it is the rule (AC) that he hears. We should perhaps not say that Nate acquires a justified false belief of the rule (AC). Many philosophers would think that he could not understand (AC) and could not have a belief whose content is (AC) because, being a fallacious rule, (AC) cannot be regarded as characterising the meaning of the word ‘if’ or the concept of material implication – or indeed any such word or concept.

justified true belief of a basic rule of logic and fail to know it, i.e. that knowledge of a basic rule of logic can be gettiered.

Now, for convenience, let us call this proposed Gettier case for knowledge of a basic rule of logic ‘the (MP)-case’. The (MP)-case has many interesting implications both for the epistemology of logic and for epistemology more generally. I mention a few before turning to the implications that I take it to have for the Understanding Account.

The (MP)-case of course shows that the traditional analysis of propositional knowledge does not work for knowledge of a basic rule of logic.⁶ Also, many discussions in standard epistemology that connect in one way or the other with Gettier cases, such as discussions of epistemic luck, sensitivity or safety, often ignore or set aside the case of logical knowledge (and more generally the case of a priori knowledge of necessary truths): but the existence of Gettier cases for knowledge of the basic rules of logic should push us towards a more integrated approach and a broadening of those notions. For instance, it is common to characterise what goes wrong in Gettier cases in terms of a belief turning true because of mere luck, or in terms of a proposition believed that could easily have turned out to be false. The truths of logic are necessarily true, so they cannot turn out to be false. Characterising what goes wrong in Gettier cases in this way is just to stipulate that Gettier cases will not apply to knowledge of necessary truths. This bias in favour of the contingent is unjustified, for intuitively, the reason why knowledge fails in Gettier cases has nothing to do with the modal status of the proposition that is believed.⁷ Finally, the target of Gettier cases is usually thought to be propositional knowledge. Thus if there are Gettier cases for knowledge of the basic rules of logic, that is an indication that this type of knowledge is propositional. However, many philosophers think that knowledge of the basic rules of logic is non-propositional.⁸ So either we have to make sense of Gettier

⁶ Alvin Goldman once said the following, when putting forward his causal account of propositional knowledge: ‘My concern will be with knowledge of empirical propositions only, since I think that the traditional analysis is adequate for knowledge of nonempirical truths.’ See A. I. Goldman, ‘A Causal Theory of Knowing’, *The Journal of Philosophy*, 64 (1967), pp. 357-372, at p. 367). This view is probably not the orthodoxy, but then again it has not been the subject of many discussions.

⁷ See here McGinn, *The Concept of Knowledge*, for an excellent discussion.

⁸ Many hold this view because of regress arguments put forward by Lewis Carroll and Gilbert Ryle that seem to show that logical knowledge cannot be propositional. See L. Carroll, ‘What the Tortoise Said to Achilles’, *Mind*, 4 (1895), pp. 78-80 and G. Ryle, ‘Knowing How and Knowing That’, repr. in his *Collected Papers* (London: Hutchinson & Co, 1971), pp. 212-225.

cases for non-propositional knowledge or we have to reassess the arguments for thinking that knowledge of the basic rules of logic is non-propositional.

II. The Understanding Account

To consider the implications of the (MP)-case for the Understanding Account, let me first briefly sketch the account.⁹ The Understanding Account says that knowledge of the basic rules of logic is in some sense *grounded on* semantic or conceptual understanding – that is to say, it is grounded on our understanding of certain sorts of words or on our grasp of certain sorts of concepts. More precisely, the account says that the basic rules of logic fix the meanings of logical expressions (such as ‘not’, ‘if’, or ‘all’) or logical concepts (such as those of negation, material implication, or the universal quantifier). Typically, those principles of logic are the introduction and elimination rules for logical expressions or concepts – for the so-called ‘logical constants’. So the first claim made by defenders of the Understanding Account says that the introduction and elimination rules for the logical constants fix the meanings of those constants or determine the concepts that such constants express. This claim is often cashed out by saying that a basic rule of logic is an *implicit definition* or a *meaning-postulate* of sorts: it is a sentence or proposition in which a logical constant occurs (is used) that is such that it (at least partly) defines that constant. Now, in principle many things could characterise the meanings of the logical constants: sentences, propositions, inference-forms, inferences, sequents, etc. However, the Understanding Account is best formulated in terms sentences or propositions (items that are true or false) and not in terms of inferences or sequents (items that are valid or invalid), given that it is knowledge and justification that are at issue in the account.

Now, (CP) and (MP) of section I are paradigmatic examples of the sorts of logical rules that are regarded as implicit definitions by defenders of the Understanding Account: (CP) is the introduction rule for the word ‘if’ or the concept of material implication and (MP) is its elimination rule. And so they are taken to jointly define that logical constant.

⁹ For particular versions of the Understanding Account, see *inter alia* P.A. Boghossian, ‘Analyticity Reconsidered’, *Noûs*, 30 (1996), pp. 360-391, C. A. Peacocke, ‘The A priori’, in F. Jackson and M. Smith (eds), *The Oxford Handbook of Contemporary Philosophy* (Oxford: Oxford University Press, 2005), pp.739-763, and B. Hale and C. Wright, ‘Implicit Definitions and the A Priori’, in P. A Boghossian and C. A. Peacocke (eds), *New Essays on the A Priori*, (Oxford: Oxford University Press), pp. 739-763. Although these versions have common ground, there are important differences between them; these differences do not matter for our purposes – for the argument presented in this paper attacks their common ground.

The second claim made by defenders of the Understanding Account is that, given their special status as implicit definitions, these basic rules of logic are such that understanding them is *sufficient* for knowing them: if one understands these rules, one knows them; and in particular, one knows them *a priori* – one knows them not on the basis of evidence from sensory experience. So I take it that the core idea of the account is this: understanding a basic rule of logic is sufficient for knowing it *a priori*. In particular, for the case in hand, understanding (MP) is sufficient for knowing (MP) *a priori*.

My characterisation of the Understanding Account allows for it to be one of either logical expressions or logical concepts. That is to say, the rules that are the objects of both understanding and knowledge can be taken to be either linguistic items (sentences), which define logical expressions or non-linguistic items (propositions), which characterise logical concepts, thus yielding different versions of the account. One possibility is to take sentences as the objects of both understanding and knowledge; another is to take propositions as the objects of both. Another consists in somewhat bizarrely taking propositions as the objects of understanding and sentences as those of knowledge. The most natural one is to take the objects of understanding to be sentences and those of knowledge to be the propositions expressed by those sentences: in this case, the account says that understanding the sentence ‘From if P, then Q, and P, one may infer that Q’ is sufficient for knowing the proposition expressed by that sentence, i.e. for knowing that from if P, then Q, and P, one may infer that Q. Here, it is natural to say that understanding this sentence involves grasping the proposition that this sentence expresses. And it is also natural to take the proposition that is grasped in understanding this sentence to be the same as the proposition that is known as the result of understanding. This last assumption is natural but perhaps not mandatory: Paul Boghossian, for instance, distinguishes between an inferential and a constitutive model of the relation between understanding a basic rule of logic and knowing it. According to the former knowledge is inferred from understanding, and according to the latter, which he endorses, knowledge is constituted by understanding. The inferential model allows for different propositions to be the objects of understanding and of knowledge; it is unclear, however, why one would take this option which would make knowing a rule of logic inferential.¹⁰

¹⁰ See P. A. Boghossian ‘Epistemic Analyticity: A Defense’, *Gratzer Philosophische Studien*, 66 (2003), pp. 15-35.

These different ways of formulating the Understanding Account, in terms of sentences or in terms of propositions, are all equally affected by my Gettier case. For simplicity, I will continue to characterise the core slogan of the Understanding Account as saying that understanding a basic rule of logic is sufficient for knowing it a priori, and in particular as saying that understanding (MP) is sufficient for knowing (MP) a priori; but the possibility of these different ways of interpreting the claim should be borne in mind.

It is useful for the foregoing discussion to consider a particular version of the Understanding Account. Let us look at Paul Boghossian's, which has received a lot of attention in the literature on the epistemology of logic. Boghossian gives an explanation of a priori knowledge of the basic rules of logic in terms of what he calls 'epistemic analyticity'.¹¹ Boghossian's analyticity is *epistemic* in that a sentence is 'true by virtue of its meaning' provided that understanding its meaning alone suffices for knowing it a priori. According to him, basic logical knowledge is a priori in that it is grounded on understanding sentences that serve to stipulate the meanings of the logical constants: it is grounded on the fact that the introduction and elimination rules for the logical constants are epistemically analytic, since those rules are precisely the sentences that fix the meanings of the logical constants – they are implicit definitions of sorts.

The general thought here is that the introduction and elimination rules define the meanings of the logical constants; and it is because such sentences are implicit definitions of some of their ingredient expressions – i.e. are epistemically analytic – that merely understanding those sentences justifies or entitles one to believing what these sentences say a priori: that is to say, merely understanding such sentences suffices for someone to be justified or entitled in believing what they say a priori in such a way as to count as knowing them a priori.¹² Thus the rules (CP) and (MP) are paradigmatic examples of rules that are epistemically analytic according to Boghossian; and that

¹¹ See *inter alia* Boghossian, 'Analyticity Reconsidered', pp. 360-391, 'Analyticity', in B. Bale and C. Wright (eds), *A Companion to the Philosophy of Language* (Oxford: Blackwell), pp. 331-368. (1997) and 'Epistemic Analyticity: A Defense', *Gratzer Philosophische Studien*, 66 (2003), pp. 15-35. Although he rightly stresses that other possibilities would be available to him, Boghossian takes sentences to be the objects of both understanding and knowledge to accommodate those who, unlike him, want to avoid commitment to propositions.

¹² In sections VII and VIII, I say more about the notions of justification and entitlement.

means that understanding (CP) is sufficient for Nate to know (CP) a priori, and equally, understanding (MP) is sufficient for him to know (MP) a priori.

So much for an overview of the Understanding Account, and of Boghossian's version of the account in terms of epistemic analyticity. This overview is by no means exhaustive, but it suffices for our purposes, and further elements of the account will be introduced as we go along.

III. Conflict Between the (MP)-case and the Understanding Account

If the Understanding Account, and in particular Boghossian's version of it in terms of epistemic analyticity, is true, then so is the following:

(1) If Nate understands (MP), he knows (MP) a priori.

However, the (MP)-case seems to show that the following is true:

(2) Nate understands (MP) and he does not know (MP) – he merely truly and justifiably believes (MP).

Claims (1) and (2) cannot be true together. I think that we should abandon (1). However, in the next sections I consider ways in which it could be argued that (2) is false. Now, if (2) is false, that means that one of the three possibilities below is true.

- a) Nate understands (MP) and knows (MP).
- b) Nate does not understand (MP) and does not know (MP).
- c) Nate does not understand (MP) and knows (MP).

In principle, a defender of the Understanding Account could argue that, unlike (2), one of a), b), and c) adequately represents what is going on in the (MP)-case. However, I will not further consider c): this option is ruled out by the way in which I have set up the example. In the example, it is stipulated that the unique way in which Nate could acquire a belief of the rule (MP) is through the testimony of Brenda; and that would require him to understand (MP). Thus here it is not possible for Nate to fail to understand (MP) and yet know it. So I only discuss options a) and b) in the next sections (IV-VII); and in later

sections (VIII-IX) I will also suggest other ways in which defenders of the Understanding Account could try to argue against 2) – and in general against the claim that understanding a basic rule of logic is not sufficient for knowing it.

IV. Option a) — Nate understands (MP) and knows (MP)

According to option a), in the (MP)-case, Nate understands the rule (MP) and knows (MP). To make this suggestion at all plausible, we would need a special reason to grant knowledge to Nate and not merely justified true belief. As possible sources of justification, the (MP)-case involves testimony and, if the Understanding Account is right, understanding. Testimony is flawed, and a natural thing to say is that this means that Nate does not know (MP). Arguing to the contrary, requires showing that in the case of knowledge of the basic rules of logic, lack of reliability does not defeat knowledge in the way it does for empirical knowledge.

To defend option a) it needs to be argued that the fact that testimony is flawed is irrelevant to Nate's justification because if he understands (MP), his operative justification somehow has to come from *understanding alone*. If so, given that Nate understands (MP), he has to know it. The suggestion here is in effect to alter the Understanding Account so that it says that when someone understands a basic rule of logic, that means that her justification for believing that rule comes from understanding alone. More precisely, since (MP) is an implicit definition and since Nate understands (MP), all his justification for believing (MP) can only come from the very fact that it is an implicit definition that he understands. So according to this revised version of the Understanding Account, the fact that the process by which Nate acquired his belief that (MP) is unreliable is irrelevant to his justification for believing it; for this justification comes entirely from the fact that it is an implicit definition that he understands – it comes from understanding alone.

This defence of a) is implausible and question-begging. In my example, Nate accesses the rule (MP) through the testimony of Brenda, and although he has no reason to doubt that Brenda is a reliable informant, she is not and he just got lucky. In effect, what this defence of a) demands us to do when we are assessing Nate's belief is to totally ignore the status of the source of communication and the fact that luck was involved in the process leading him to believe (MP) – somehow, in this case, these do not make Nate

epistemologically worse-off. But it is unclear why these facts can be ignored here, when elsewhere, when knowledge of propositions that are not basic logical rules is considered, lies and luck are paradigmatic defeaters of knowledge.

Now, it may be the case that a basic rule of logic is an implicit definition – and thus perhaps has some sort of justificatory force in and of itself – but that does not mean that we can *stipulate* that someone cannot believe that implicit definition in such a way that their justification depends on something other than the justificatory force that the rule has because of it is an implicit definition. In our example, it may be argued that Nate's justification has little to do with the fact that it is an implicit definition that he understands, or with the fact that he judges the rule to be obvious; it may be that he has not (yet) reflected on the sort of sentence that (MP) is, or on whether he finds it independently compelling – independently, that is, from the fact that Brenda (apparently) told him (MP). That is to say, the case envisaged here is one in which Nate simply believes what Brenda (apparently) said: he believes (MP) because Brenda (apparently) told him so; as far as he is concerned, Brenda said (MP), and a natural thing to do for a student, is to believe what the teacher says. In this case, it is not at all plausible to say that Nate's justification for believing (MP) merely comes from the fact that it is an implicit definition that he understands: Nate's justification comes from Brenda, and this is why Nate does not know (MP). Of course, it may be the case that the rule would strike him as obvious if he was to reflect on it, and that if he reflected on the sort of proposition that (MP) is, his operative justification would change from testimony to understanding, but *as things stand*, he believes (MP) because Brenda (apparently) told him so.

So the suggested defence of a) is question-begging because it stipulates that when someone forms a belief whose content is an implicit definition, that person's justification for believing the proposition ought to come from the fact that it is an implicit definition that is understood, and not from something else – it ought to come from understanding alone. In the (MP)-case, the operative justification for Nate's belief is testimonial justification. And so the fact that Nate understands (MP) does not *as such* mean that his operative justification for believing (MP) has to be understanding. As a result of his interaction with Brenda, we should not count Nate as knowing rule (MP), but as only having a justified true belief in (MP). And thus the (MP)-case indeed shows that

understanding does not guarantee knowledge in the way required by the Understanding Account.

V. Option b) — Nate does not understand (MP) and does not know (MP)

According to option b), Nate does not understand (MP) – and the word ‘if’ in particular – and thus he does not know (MP). The thought here is that since Nate does not understand (MP), he cannot even believe (MP), let alone truly and justifiably believe it. The general reason for this is that one cannot believe what one does not understand. Thus if we interpret the (MP)-case in this way, there is no conflict with the Understanding Account and Boghossian’s view in particular, and indeed we do not have a genuine Gettier case here. Obviously, it is a necessary condition to have such a case that Nate believes (MP).

However, I do not think that option b) does not seem to provide a plausible interpretation of the (MP)-case: it is implausible to say that Nate does not understand (MP) and the word ‘if’ in particular; for there is nothing in the situation that should make him semantically deficient in the relevant respect. Indeed, it seems that he has acquired whatever it takes to be a competent user of the word ‘if’. For example, he is in a position in which he could successfully start interacting with other users of the word ‘if’, and in a position in which he could even start teaching someone else what it means. We should count him as understanding because *everything is in place* for him to use the word successfully. And the fact that he has acquired the ability to use the word ought to be good enough for understanding.

Well, at least, we should say that everything is in place for Nate to use the word ‘if’ successfully, if, like defenders of the Understanding Account, we think that there is nothing more to understanding it than understanding both its introduction rule (CP) and its elimination rule (MP). If this is indeed all that understanding the word requires, it seems that Nate should be counted as a competent user of the word.

Perhaps an analogy is appropriate here to help sharpen the intuition that in the (MP)-case Nate understands (MP). The analogy concerns empirical knowledge and a case of ostensive definition – where someone teaches someone else what a word means or what a concept is by pointing to a sample of something that satisfies the word or concept. Ostensive definitions give us a good analogy because they are contextual definitions, and

they can be similar to the sorts of implicit definitions that are involved in the (MP)-case, as when the word or concept to be defined is used in the sentence or proposition that defines it.

Thus suppose this time that Brenda wants to teach Nate words for vegetables. She starts with the words ‘aubergine’ and ‘courgette’, since she has some in her garden. But again, she has the urge to trick him and decides to teach him the wrong word for the wrong thing. Rather than saying ‘These are aubergines’ when pointing at aubergines and ‘These are courgettes’ when pointing at courgettes, she intends to say ‘These are aubergines’ when pointing at courgettes and ‘These are courgettes’ when pointing at aubergines.

But as she takes Nate to her garden, one of these strange atmospheric conditions results in her having an optical illusion such that when it seems to her that she is pointing at aubergines, she is actually pointing at courgettes and when it seems to her that she is pointing at courgettes she is actually pointing at aubergines. So unbeknownst to her she actually points at aubergines when she says ‘These are aubergines’ and she actually points at courgettes when she says ‘These are courgettes’. However, since Nate is not the victim of this illusion, he gets the right word for the right thing – the ostensive definition is thus successful for him, and he gets to understand the words ‘courgette’ and ‘aubergine’. But he cannot be said to have knowledge of the propositions that these (over here) are aubergines and those (over there) are courgettes, again, for the familiar reason that one cannot get knowledge by luck.

It is intuitively clear that in this case Nate understands the words ‘aubergine’ and ‘courgette’. But it is also equally clear that if we are happy to grant understanding to Nate in this case, we should be happy to grant him understanding in the logical case involving the rule (MP).¹³ Yet, although the definitions are successful in each case, in neither case does Nate seem to gain knowledge.

VI. Towards a Diagnosis

If the two Gettier cases that I have presented are indeed cases in which we have understanding without knowledge, it would be nice to have an explanation as to why we have this lack of fit between the two – as to why does one fail in my Gettier cases and not

¹³ I come back to these two cases when I consider ways of strengthening understanding in section VII.

the other. Here, I try to articulate some possible reasons as to why in the two cases that I have presented, understanding obtains, but knowledge does not. What follows is somewhat hand wavy and should not be taken as a final diagnosis.

In both the (MP)- and the courgette/aubergine-cases, Brenda's intentions to deceive Nate and the fact that he got lucky explain why Nate understands the relevant proposition but fails to know it. In particular, what seems to be the case is that the *intentions of the source of communication* interact in different ways with understanding and with knowledge: Brenda's intention to communicate something false, in this case, incorrect definitions, does not undermine the fact that Nate understands the relevant expressions. To take the (MP)-case again: it was not Brenda's intention to communicate (MP) to Nate. She wanted to communicate something else to him – namely (CP). And what he ended up understanding is not what she intended to communicate to him. It is rather what she actually communicated to him – or, if we do not want to impute that sort of agency to her, it is what was actually communicated to him, by chance. This indicates that for Nate to understand the relevant piece of language it is not necessary that he understands what Brenda intended to communicate to him or what she actually said.

However, the fact that it was Brenda's intention to deceive Nate partly explains why Nate does not know the rule (MP), although he understands it. To see this, suppose that the (MP)-case actually went like this: Brenda actually utters (MP) but with the intention of lying or somehow thinking that she was uttering (CP), there is no whirlwind, and Nate understands (MP). In this case, we would still not think that Nate knows (MP). He would not know it because his source is lying and one cannot gain knowledge through lies.

If these considerations are right, that means that understanding does not connect with intentions of communication in the way in which knowledge does: understanding through testimony does not depend on the intentions of the source to the same extent that the acquisition of knowledge through testimony does. That in turn indicates that justification does not work in the same way for understanding and knowledge. And indeed the question of how justification works for understanding – that is to say the

question of what kind of justification is required for understanding – is by no means a settled matter.¹⁴

Further, something else that might explain the asymmetry between understanding and knowledge is that the standards on the former may be *more lax* than those on the latter. Perhaps understanding is less of a cognitive achievement than knowledge. Indeed, we seem to be happy to count as understanding or linguistically competent people with very little or shallow exposure to an expression or concept, and very little capacity to form an accurate or precise conception of what the expression means or the concept is – and logic is no exception. And of course, nowadays many philosophers are externalists, and think that a substantive amount of the expressions that we understand or concepts we have, we do so *deferentially*, by relying on other, more competent, speakers – and logic is again no exception.¹⁵

To really tackle these questions about whether the standards on understanding are in some sense more lax than those on knowledge, and about how justification works for understanding, we would have to look at how Gettier cases (if there are any) work for understanding, and at how the concepts of epistemic luck or safety or what have you would work for understanding language or having concepts. But that would deserve another paper.

Now at this stage, a clarification about something that the (MP)-case does not show but that it might initially be thought to show is here needed. So far I have contrasted understanding and knowledge, and this might be taken to suggest that understanding is not a sort of knowledge – and in particular, not a sort of propositional knowledge. Here, I want to leave it open whether understanding is a sort of propositional knowledge. So when I contrast understanding and knowledge, this should not be read as a way of saying that understanding is not a sort of propositional knowledge. To be more explicit, the

¹⁴ For discussions of how justification connects with understanding, see: Dean Pettit, ‘Why Knowledge is Unnecessary for Understanding Language’, D. Hunter, ‘Understanding and Belief’, *Philosophy and Phenomenological Research*, 64 (1998): pp. 147-269, E. Fricker, ‘Understanding and Knowledge of What is Said’, in A. Barber (ed.), *The Epistemology of Language* (Oxford: Oxford University Press, 2003): pp. 325-366.

¹⁵ Recall for instance Putnam, who claims that he understands the expressions ‘elm’ and ‘beech’ but cannot tell elms apart from beeches. See H. Putnam, ‘The Meaning of ‘Meaning’’, repr. in *Mind, Language, and Reality: Philosophical Papers*, Vol. 2 (Cambridge: Cambridge University Press, (1975), pp. 215-271, at p. 228.

contrast is between understanding (which may or may not be a sort of propositional knowledge) and propositional knowledge which is not specifically about language or about concepts, i.e. non-semantic or non-conceptual knowledge.

However, it might be thought that the fact in the (MP)-case (non-semantic or non-conceptual) knowledge fails but understanding does shows that understanding is not a sort of propositional knowledge; for if it were, it would fail in the (MP) case too. I do not think that this is the case: we would perhaps be in a position to conclude that understanding is not a form of propositional knowledge from the claim that Gettier cases whose explicit target is understanding are not possible.¹⁶ But this requires more than I have done here in just giving a Gettier case in which (non-semantic or non-conceptual) knowledge does fail but understanding does not. For it may be that the cases that defeat the latter are different from those that defeat the former: that they are the same cannot simply be taken for granted.

VII. Option b) Again — Strengthening Understanding

Let us now consider another strategy to defend option b) and object to the claim that Nate understands (MP), and the word ‘if’ in particular, but does not know it. One could object to the idea that Nate understands (MP) on the ground that the Understanding Account, and Boghossian’s account in terms of epistemic analyticity in particular, would actually require Nate to understand (MP) in a *far more exacting* sense than that answering to the lax standards mentioned in section VI. The objection would be that Nate in my example does *not* understand (MP) (and for that matter does not understand (CP)) because he does not understand it as an implicit definition; and to be justified in believing (MP) just by understanding it, one needs to know that (MP) is an implicit definition. Thus if Nate does not know that (MP) is an implicit definition, he does not understand (MP) in the required sense – and the (MP)-case is not a counter-example to the Understanding Account.

This objection might be thought to be especially fitting because it echoes some of the things that Boghossian says concerning epistemic analyticity. Indeed he sometimes

¹⁶ See D. Pettit, ‘Why Knowledge is Unnecessary for Understanding Language’, *Mind*, 111 (2002), pp. 519-550, for arguments that there are no Gettier cases for understanding.

suggests that understanding (MP) amounts to being in a position to run through the following sort of argument:¹⁷

- (i) If ‘if’ is to mean what it does, (MP) must be valid (since ‘if’ is to mean whatever makes (MP) valid).
- (ii) ‘if’ means what it does.
- (iii) Therefore, (MP) is valid.

According to this explanation of what it takes to understand an implicit definition such as (MP), understanding (MP) requires knowing how the meaning of (MP) is fixed. And since, in our example, Nate does not know how the meaning of (MP) is fixed, it cannot be claimed that he understands (MP).

I think that this strategy to save the Understanding Account does not work. Firstly, the (MP)-case could easily be modified so that Nate is told in advance that the rules that he is about to hear are implicit definitions. We could of course suppose that Brenda tells him that he is going to learn a new word that philosophers and logicians call ‘logical constant’, that he will learn it by being given rules that have the status of implicit definitions; she could tell him what implicit definitions are, how they work, and what is special about them, she could even talk to him about Boghossian’s views on what understanding implicit definitions involves, and perhaps run through an example, using another logical constant, say, conjunction. In this case, Nate would understand (MP) in the more demanding sense required – but he would still not know it. For as long as his operative justification for believing (MP) does not come from understanding alone, he does not know it. It could be further argued that if Nate can run through argument (i)-(iii), his operative justification for (MP) *ought to be* understanding. However, for the very same reasons given in section IV, I think that that would be begging the question. So I do not think that a more exacting conception of understanding would create a real difficulty for my (MP)-case.

Another reason why this strategy to save b) does not work is that philosophers who hold one version or the other of the Understanding Account, and Boghossian included, want the person on the street to be granted knowledge of the rules of logic. However such a

¹⁷ See Boghossian, ‘Analyticity Reconsidered’, at p. 386 and ‘Analyticity’, at pp. 356-357.

person typically does not understand the rules of logic in the exacting sense just mentioned. She could probably not verbalise or articulate such knowledge, let alone run through the argument (i)-(iii). So if she is to have knowledge of the rules of logic, understanding has sometimes to be easily attainable.

And so what defenders of the Understanding Account typically say is that implicit definitions may be implicit – they may be *implicitly known*. In this case, to count as knowing a proposition one need not be able to verbalise or to explicitly articulate that proposition. The picture of understanding that is required by the idea that implicit definitions may be implicitly known is not one that requires people to be able to know how the meanings of the logical constants are fixed, but merely able to use these expressions correctly.

For instance, Boghossian thinks that someone who does not know that a rule such as (MP) is an implicit definition can still count as understanding it, and the word ‘if’ in particular. If she understands it that way – or, as Boghossian puts it, if she *merely grasps its meaning*¹⁸ – she is not as such *justified* in believing (MP), but merely *entitled* to believe (MP). Entitlement and justification are different in that the latter needs to be cognitively accessible to the subject, whereas the former does not; but crucially, both are sufficient for knowledge. Here is how Boghossian contrasts the two notions:

When someone is entitled, all the facts relevant to the person’s justification are already in place, so to say; what is missing is the reflection that would reveal them.¹⁹

The person on the street typically inherits her entitlement for (MP) by interacting with other people who know (MP). Again, a conception of understanding that fits well the idea that understanding suffices to merely generate entitlement (as opposed to justification) is one in which implicit definitions need not be understood as implicit definitions to be understood. And if implicit definitions may be implicit, that again means that Nate should count as understanding rule (MP).

¹⁸ See Boghossian, ‘Epistemic Analyticity: A Defense’, pp. 16-7) for a distinction between what he calls ‘mere grasp of meaning’ (which is the less exacting sense at issue here), ‘knowledge of meaning’, and ‘understanding of meaning’. The latter two are more demanding, and the last one would demand that one knows how the meaning of the meaning of the expression is fixed.

¹⁹ Boghossian, ‘Analyticity Reconsidered’, p. 387.

At this stage one could simply require implicit definitions to be explicitly known, and thus make the conditions on understanding more exacting. However, that would limit greatly the appeal of the Understanding Account, because very few people would count as understanding (MP). Also, unless we want to say that understanding works differently in the (MP) case and in the courgette/aubergine case of section V, which seems implausible given that they involve similar sorts of definitions, we will have to say that in the latter case too, it should be required that one understands how the meanings of ‘courgette’ and ‘aubergine’ are fixed in order to count as understanding these words. But this is clearly inappropriate in this case. So either we have to artificially claim that understanding definitions works differently for logical terms or concepts and for non-logical terms or concepts, or we have to make understanding definitions more easily attainable.

Finally, quite generally, making understanding more difficult to achieve would simply be to misconstrue what understanding is about. To put it loosely: understanding is intuitively just about getting certain contents in one’s head, and if these contents are the right ones, one just understands what is said. Brenda is not an ordinary speaker, because she intends to make Nate understand the wrong thing. Yet it seems that Nate acquires the very same linguistic or conceptual abilities that he would have acquired if Brenda had been an ordinary speaker.²⁰

This closes my discussion of a) and b), neither of which seem to be satisfactory ways of defending the Understanding Account, and Boghossian’s epistemic analyticity in particular. These were put forward to defend the Understanding Account against the natural conclusion to be drawn from the (MP)-case that Nate understands (MP) but does not know (MP) (c.f. section III). None of these options work, which means that, as it is, the Understanding Account is false. It is false because there are situations in which understanding a basic rule of logic is not sufficient for knowing it.

²⁰ See Pettit ‘Why Knowledge is Unnecessary for Understanding Language’, p. 528, for similar remarks about understanding.

VIII. Weakening the Understanding Account

I now turn to a different sort of move that a defender of the Understanding Account could make. In the way that I have characterised it, it says that merely understanding a basic rule of logic is sufficient for knowing such a rule a priori. However, a defender of the Understanding Account could argue that this is not an accurate characterisation of the account: it does not simply say that understanding a rule such as (MP) is sufficient for knowing it. Rather, it should be taken to say that understanding (MP) together with some extra conditions is sufficient for knowledge of (MP) or that understanding (MP) gives one the sort of justification that makes one in a position to know (MP), if some other condition is fulfilled. So understanding generates justification or entitlement but it is not strictly speaking sufficient for knowing; we need an extra condition – a condition that ensures that there are no defeaters, and that insulates the account from cases such as the (MP)-case.

Now this suggestion that understanding alone is not sufficient for knowledge might initially seem right, since some proponents of the Understanding Account seem, at times, to focus on justification (or entitlements or warrants) rather than directly on knowledge. For instance, Boghossian often expresses himself in terms of justification rather than in terms of knowledge – as when he says that a rule of logic is epistemically analytic if ‘grasp of its meaning alone suffices for *justified* belief in its truth.’²¹ That might in turn suggest that what the Understanding Account aims to be an account of is merely justification for believing the basic rules of logic, where that might not be sufficient for knowledge.

However, it would be surprising if defenders of the Understanding Account are just aiming to talk about a sort of justification that is not sufficient for knowledge. And I see no real evidence in the literature that the account is really meant to be of the following form, taking again (MP) as an illustration:

(1*) If Nate understands (MP) & condition C obtains, then Nate knows (MP) a priori.

²¹ Boghossian ‘Analyticity Reconsidered’, p. 363. My italics.

Again, this lack of evidence may be due to the fact that defenders of the Understanding Account are focusing on justification and not knowledge, and are thus not interested in looking at what this extra condition might be. But still it would be surprising if something like (1*) really represented the structure of the account of knowledge of the basic rules of logic that they want to put forward.²²

One major problem with this suggestion is that it would no longer be clear how understanding could play such a special or central role in an account of knowledge of the basic rules of logic, and a priori knowledge in general; understanding would only be one condition on knowledge amongst others, which is much less than we hoped for; indeed, the initial hope was that understanding the right sorts of sentences could carry enough justificatory force to endow people with knowledge of the basic rules of logic. But if understanding is only an aspect of what explains such knowledge, there is no longer anything about the account that is clearly distinctive. For understanding is obviously a condition on having any propositional knowledge.

Now, it may be still be that, in the case of the rules of logic, understanding is sufficient for something less than knowledge, which is not targeted by Gettier cases, such as justified true belief or simply true belief.²³ I have not tried to argue against this possibility. That would be an interesting and claim in and of itself. If indeed merely understanding (MP) suffices for having a true belief that (MP), that perhaps means that the rule (MP) is in some sense an implicit definition or even in some sense analytic. But still, the road from true belief to knowledge, which is the target of the Understanding Account, is a long one.

IX. Testimony

Finally, I briefly mention a worry of a different kind concerning the (MP)-case. Defenders of the Understanding Account claim that the justification or entitlement one gets when one understands a rule such as (MP) is of an a priori sort. This is required if the knowledge thus explained is to be a priori. It could be argued that, given the setting of the (MP)-case, Nate's belief in (MP) is not a priori, because it has been acquired

²² See Peacocke 'The A Priori', who makes it particularly clear that it is knowledge that is at issue in his version of the Understanding Account.

²³ Thanks to one of my referees for making this suggestion.

through testimony, and one cannot gain an a priori belief in this way: testimony involves sensory experience (e.g. hearing or reading), and that sort of sensory experience is really part of one's justification for the belief thus acquired. And if sensory experience is part of Nate's justification for believing (MP), the belief is empirical.

I cannot here do justice to the complex issue of whether a priori beliefs can be gained through testimony.²⁴ But also, I do not need to; at least I do not need to, if my aim is just to give a counter-example to the Understanding Account. For as far as the question of whether Nate's belief is a priori is concerned, the Understanding Account and the (MP)-case stand or fall together. The reason for this is that defenders of the Understanding Account typically allow for a priori knowledge to be gained through testimony.

The point here is dialectical. A defender of the Understanding Account claims that understanding generates a priori knowledge of basic logical rules. But it is to be expected that understanding will in many instances involve sensory experience, in some sense or other, and that in many cases it will involve testimony. As we have seen in section VII, this is especially the case if the Understanding Account aims to explain the person on the street's justification for believing the truths of logic; for the person on the street who is typically merely *entitled* to believe the truths of logic, will typically know these rules through testimony – in a way similar to that in which Nate acquired his belief of (MP).

So the question of whether one can gain a priori knowledge through testimony is a question for everybody, but it is in particular a question for people who hold the Understanding Account, where understanding a priori knowledge of the basic rules of logic; for understanding typically involves perceptual experience in the way in which we access those sentences or propositions we understand. As Peacocke, who defends a version of the Understanding Account, writes (in discussing basic a priori knowledge in general):

The means by which we come to know outright a priori [logical] propositions [m]ay involve any or all of the following: conversations and discussions with others;

²⁴ For arguments that a priori beliefs can be gained through testimony, see T. Burge, 'Content Preservation', *The Philosophical Review*, 102.4 (1993), pp. 457-488. To my knowledge, all the defenders of the Understanding Account endorse Burge's views on testimony.

reading books which archive knowledge achieved many generations ago; our own workings out on paper; musing and reflection on examples; computer simulations and computer proofs; and much else. Many in this array of methods involve perception at some stage or other.²⁵

Again, answering the question whether the fact that these methods of gaining knowledge involve perception at some stage really invalidates the a priori status of that knowledge is beyond the scope of this paper. But the dialectical point stands here: if a defender of the Understanding Account grants a priori knowledge to people who acquire logical knowledge through testimony, one cannot use the fact that Nate acquired his belief in (MP) by testimony to argue that the (MP)-case fails to threaten the Understanding Account. Thus, even if it turned out that one cannot acquire an a priori belief through testimony, and thus that Nate's belief is empirical, we would still have a counter-example to the Understanding Account, because in that case understanding would typically not generate the sorts of a priori justification or entitlement that the view requires.

Concluding Remarks

The aim of this paper was to argue against the Understanding Account of the basic rules of logic, according to which knowledge of a basic rule of logic is grounded in semantic or conceptual understanding. To do so, I showed that the account is open to Gettier style counter-examples such as the (MP)-case. What the (MP)-case suggests is that there is no safe route from understanding to logical knowledge, contrary to what the Understanding Account claims. If so, that means that the Understanding Account is false, and that we have to look for other ways to think about what knowledge of a basic rule of logic amounts to.²⁶

St Hugh's College, Oxford

²⁵ Peacocke 'The A Priori', p. 746.

²⁶ Thanks to Hemdat Lerman, Ofra Magidor, Anna-Sara Malmgren, Bruno Whittle and Timothy Williamson for helpful written comments on a draft of this paper. Versions of the material were presented at Aix-en-Provence, Oxford and at the Arché Basic Knowledge Pilot Workshop at St. Andrews. I am grateful to the audiences on all these occasions for discussion, and especially to my commentator at St. Andrews, Elia Zardini. Thanks also to the referees for this journal for their very useful comments.