

What Should a Theory of Knowledge Do?

Elijah CHUDNOFF[†]

ABSTRACT

The Gettier Problem is the problem of revising the view that knowledge is justified true belief in a way that is immune to Gettier counter-examples. The “Gettier Problem problem”, according to Lycan, is the problem of saying what is misguided about trying to solve the Gettier Problem. In this paper I take up the Gettier Problem problem. I distinguish giving conditions that are necessary and sufficient for knowledge from giving conditions that explain why one knows when one does know. I argue that the problem with the Gettier Problem is that it requires us to articulate conditions that suffice for knowledge even if those conditions are non-explanatory. After defending this view, I take up two related methodological issues, one about the evidence that can be given in favor of an account of knowledge, and one about the role that investigating justification might play in investigating knowledge.

Everyone is familiar with the Gettier Problem. It is the problem of revising the view that knowledge is justified true belief in a way that is immune to Gettier counter-examples. William Lycan has recently introduced a new, related problem: “What I shall call the ‘Gettier Problem problem’ is that of explaining what is distinctively wrong with the Gettier project,” i.e. the project of trying to solve the Gettier Problem. (Lycan 2006, 150).

One approach to the Gettier Problem problem is to attack the analytical ambition at the core of the Gettier Problem. The Gettier Problem has a solution only if there is an analysis of knowledge, but there is no analysis of knowledge, so there is no solution to the Gettier Problem. This approach is workable only if the sense in which there is no analysis of knowledge is the same sense in which a solution to the Gettier Problem demands an analysis of knowledge. Lycan gives some reason to doubt that this is so.¹

Another approach to the Gettier Problem problem is to attack the intuitions that generate it. The idea is that while Gettier cases intuitively appear to be counter-examples to the view that knowledge is justified true belief, these intuitive

[†] Department of Philosophy, University of Miami, Coral Gables, FL, USA; Email: echudnoff@gmail.com.

¹ Quineans attack the analytical ambition by attacking the notion of analytic truth; see Quine (1951). As Lycan (2006, 151) points out, an ‘analysis’ of knowledge need not be analytic. Williamson (2000) attacks the analytical ambition by attacking the view that the state of knowing is composed of other more basic states. As Lycan (2006, 161) points out, an ‘analysis’ of knowledge need not be a dissection of it into components.

appearances are misleading. The answer to the Gettier Problem problem, then, is that there is no Gettier Problem to try to solve.²

Lycan himself defends the view that on reflection there really isn't anything wrong with trying to solve the Gettier Problem. The proper response to the Gettier Problem problem is to reject it. Lycan's optimistic conclusion, however, depends on his rejection of at least some of the intuitions that have driven the Gettier literature. For example, Lycan rejects the intuition that people in fake barn cases fail to have knowledge. (Lycan 2006, 158, 161).

In this paper I develop an alternative to these three approaches. It has the following virtues: it both vindicates and illuminates the widespread sense that there is something wrong with the Gettier Problem; it does not require us to reject any of our intuitions about Gettier cases, including fake barn cases; and it does not depend on any controversial claims about analysis, and in particular the claim that the sense in which there might be no analysis of knowledge is the same sense in which a solution to the Gettier Problem demands an analysis of knowledge.³

Here is a brief sketch of the main idea. A condition of success on a solution to the Gettier Problem is that it provides a sufficient condition for knowledge. But, I will argue, there are cases in which the set of conditions in virtue of which S knows that p, say, are not conditions sufficient for S to know that p. If we were to expand the set of conditions in virtue of which S knows that p into a set of conditions sufficient for S to know that p the result would be a set of conditions that does not explain why S knows that p. A solution to the Gettier Problem requires precisely such a non-explanatory set of sufficient conditions for knowledge, and so to that extent trying to develop such a solution is misguided.

I proceed as follows. In section 1, I discuss the 'in virtue of' relation and distinguish between saying in virtue of what a fact obtains and giving necessary and sufficient conditions for that fact to obtain. I argue that Q can obtain in virtue of P even though P does not suffice for Q. In section 2, I discuss perceptual knowledge. I set out a simple view of what it is in virtue of which one can have perceptual knowledge. In section 3, I discuss fake barn cases. These show that the view of the grounds of perceptual knowledge explored in section 2 fails to identify sufficient conditions for perceptual knowledge. I explain why I do not think that this is a good reason to reject it, and so why embarking on the Gettier project of seeking sufficient conditions for perceptual knowledge is misguided. This explanation yields an approach to the general Gettier Problem problem. In section 4, I conclude by taking up two methodological issues. The first is about what evidence can be given in favor

² For an example of this strategy see Weatherson (2003).

³ The published literature leaves it a bit unclear just how widespread disillusionment with the Gettier Problem is. See Lycan (2006) for a few references. My impression that it is fairly widespread derives more from personal communication.

of an account of knowledge, or knowledge of some specific kind. The second is about the role of investigating justification in investigating knowledge.

1. 'In virtue of'

Some facts obtain in virtue of other facts. For example, the fact that I am illegally parked might obtain in virtue of the fact that I am parked next to a fire hydrant, or the fact that this ball is red might obtain in virtue of the fact that it is vermilion. The 'in virtue of' relation is a *sui generis*, non-causal, dependence relation.⁴

There are two widely recognized differences between saying 'Q obtains in virtue of P' and saying 'P is a necessary and sufficient condition for Q to obtain'. First, P might be a necessary and sufficient condition for Q to obtain even though Q does not obtain in virtue of P. The fact that my car is illegally parked is a necessary and sufficient condition for the fact that my car is illegally parked to obtain, but the fact that my car is illegally parked does not obtain in virtue of the fact that my car is illegally parked. Second, Q might obtain in virtue of P even though P is not a necessary condition for Q to obtain. The fact that this ball is red might obtain in virtue of the fact that it is vermilion, but the fact that it is vermilion is not a necessary condition for the fact that it is red to obtain. In this section I argue that there is yet a third difference, which I will call Defeasibility:

Defeasibility: Q might obtain in virtue of P even though P is not a sufficient condition for Q to obtain.

This claim would be trivial if it meant that Q might obtain *partly* in virtue of P even though P is not a sufficient condition for Q to obtain. But what I mean by it is that Q might obtain *wholly* in virtue of P even though P is not a sufficient condition for Q to obtain.

Consider first a *prima facie* example illustrating this. Suppose you have a visual experience as of a red light ahead, and that it justifies you in believing that there is a red light ahead. An 'in virtue of' claim follows: you are justified in believing that there is a red light ahead in virtue of having a visual experience as of a red light ahead. But it is not the case that having a visual experience as of a red light ahead suffices for you to be justified in believing there is a red light ahead. Suppose you have the visual experience and you know you have taken a pill that will make green things look red. In this case your visual experience does not justify you in believing that there is a red light ahead.

⁴ For recent discussion of it see Schaffer (2009), Rosen (2010), Audi (ms), Fine (ms). The thrust of these discussions is that the 'in virtue of' or 'grounding' relation is as clear as any fundamental notion in philosophy and can play a role in structuring philosophical inquiry for which common alternatives such as entailment and supervenience are unsuited.

One might challenge the example. Perhaps in the initial setup you have your justification partly in virtue of having a visual experience and partly in virtue of *not* knowing that you have taken a pill that will make green things look red. This is a theoretical option we might pursue. But it is neither mandatory nor advisable. To explain why I do not think it is advisable, I will consider some more general principles about the ‘in virtue of’ relation.⁵

The first two are widely appreciated:

1. ‘In virtue of’ is an explanatory relation. So, if Q obtains in virtue of P, then: P explains Q, or Q obtains because P obtains.
2. ‘In virtue of’, like other explanatory relations, is non-monotonic. So, suppose Q obtains in virtue of P; it doesn’t follow that Q obtains in virtue of P and R.

Why is ‘in virtue of’ non-monotonic? Consider an example that illustrates its non-monotony:

The ball is red in virtue of being vermilion, but it is not the case that the ball is red in virtue of the fact that the ball is vermilion and the fact that Socrates is a philosopher.

The fact that the ball is red does not obtain in virtue of both the fact that the ball is vermilion and the fact that Socrates is a philosopher because the fact that Socrates is a philosopher is explanatorily irrelevant to the fact that the ball is red. Explanatory relations in general are non-monotonic because a relation of explanatory relevance restricts them.⁶ That is, in general, if P explains Q, or if Q obtains because P obtains, then P is explanatorily relevant to Q. And more specifically:

3. If Q obtains in virtue of P, then P is explanatorily relevant to Q

In light of (3), the fourth point is trivial:

4. If there are P, Q, and C such that Q obtains in virtue of P, P does not entail C, Q obtains in virtue of P only if C, but C is not explanatorily relevant to Q, then Q obtains in virtue of P even though P is not a sufficient condition for Q to obtain.⁷

The fifth and final point is substantive:

⁵ I explain why (Chudnoff, ms) contrary to a line of argument suggested by some of the recent literature on the ‘in virtue of’ cited above, I do not think that it is mandatory.

⁶ For further discussion of this point see Salmon (1984, 1989), and Woodward (2003).

⁷ C might be called an enabling condition. See Dancy (2006) for further sympathetic discussion of the distinction between explanatory and enabling conditions.

5. There are P, Q, and C such that Q obtains in virtue of P, P does not entail C, Q obtains in virtue of P only if C, but C is not explanatorily relevant to Q.⁸

What these considerations do is connect the question of whether Defeasibility is true – i.e. whether Q might obtain in virtue of P even though P is not a sufficient condition for Q to obtain – to questions about explanatory relevance. More specifically, they show that if there is good reason to believe (5), then there is good reason to believe Defeasibility.

To appreciate the motivation for (5), consider a general point James Woodward (2003, 179) makes about explanatory relevance:

To say that certain information is ‘part’ of an explanation or contributes to its explanatory import is to say that this information contributes to the understanding provided by the explanation. This in turn imposes an epistemic constraint on what information can be part of an explanation and can contribute to its explanatory import: such information must be epistemically accessible to those who use the explanation.

Woodward argues from the claim that parts of an explanation – i.e. explanatorily relevant conditions – must contribute to our understanding of our explanandum to the claim that there are epistemic accessibility constraints on explanatory relevance. Perhaps this is correct. But I want to consider two constraints on explanatory relevance that are more modest than epistemic accessibility constraints.

These are non-vacuity and naturalness constraints:

C is explanatorily relevant to Q only if C is not vacuous in a way that renders it unable to contribute to understanding why Q obtains.

The Cs are explanatorily relevant to Q only if the set of the Cs is not unnatural in a way that renders it unable to contribute to understanding why Q obtains.

Neither vacuity nor naturalness admits of simple definition. But there are clear examples of both, and we seem largely to agree on how to apply the notions. *Virtus dormitiva* explanations illustrate vacuity. The explanation for why opium puts you to sleep should not contain the condition that it has the capacity to put you to sleep.

⁸ The formulations of (4) and (5) benefited from helpful advice from Kirk Ludwig, Anne Meylan, and Richard Woodward. Consider a simpler version of (5): There are P, Q, and C such that P grounds Q only if C, but C is not explanatorily relevant to Q. One might worry that this is trivial: let P be “the ball is vermilion,” Q be “the ball is red,” and C be “ $2 + 2 = 4$.” One way to reply to the worry is to complicate (5), as I have in the text. Another way to reply to the worry is to reject the reading of “only if” which it presupposes. It presupposes that “only if” is a material conditional or some other conditional for which a principle such as $\Box B \vdash A \supset B$ holds. But – setting aside issues about the proper interpretation of English conditionals – it is certainly possible to reject any such reading in favor of another for which no such principle holds. Though I am attracted to this option – since it would allow us to explore claims such as (4) and (5) with respect to the grounding of necessary truths – to examine the issue further would take us off topic.

Why? Precisely because this condition is vacuous given our explanandum.⁹ There are physical conditions that determine the fluctuations in the stock market. But they do not compose an explanation of the fluctuations in the stock market. Why think this? Precisely because they are so wildly heterogeneous that the set of them is too unnatural to provide us with any understanding of why the stock market behaves as it does.¹⁰

One might grant that vacuity and naturalness should be taken as primitive notions, but still worry about my deployment of them in characterizing explanatory relevance.¹¹ Let me put the worry in the form of an argument:

- (A) What counts as vacuous or natural is relative to us – our concepts, our interests, our background knowledge, etc.
- (B) What counts as explanatorily relevant is not relative to us.
- (C) So non-vacuity and naturalness are not constraints on explanatory relevance.

One way to resist the argument is to deny premise (B). I will not do this however. In my view it is premise (A) that should be denied.

This might not seem like an option for me since I have characterized vacuity and naturalness in part by their relation to our understanding. But I do not intend to *define* vacuity and naturalness in part by their relation to our understanding. I think it is true that vacuous conditions and unnatural sets of conditions do not contribute to our understanding of an explanandum. But this might be because what counts as understanding is partly defined in terms of vacuity and naturalness. Or it might be a contingent relation that holds independently of definitions. This is a large issue that I must set aside for now. The point I want to emphasize is that nothing I have said about vacuity and naturalness requires that what counts as vacuous or natural is relative to us.

Of course, we choose what our explanandum will be. And vacuity is relative to that. But once the explanandum is fixed, then what counts as a vacuous condition in an explanation of it is fixed. Heuristically: a condition is vacuous to the extent that it provides no more information than is already in the explanandum itself.¹² Further, I follow Lewis in taking naturalness to be a mind-independent character-

⁹ For more on vacuity see Strevens (2008). Strevens discusses the vacuity of *virtus dormitiva* explanations on pages 131–132.

¹⁰ For more on naturalness see Lewis (1983, 1986). Weatherston (2003) discusses the notion further and makes different use of it in developing his own approach to the Gettier Problem problem.

¹¹ I thank Kevin Morris and an anonymous reviewer for pressing me to address the following worry.

¹² I am thinking of an explanandum as a fact – a structured whole whose constituents are items such as properties, which are individuated independently of their relation to us. I am not thinking of an explanandum as a proposition or a Fregean Thought – a structured whole whose

istic of certain sets of individuals, conditions, truths, etc.¹³ Heuristically: a set is natural to the extent that its members are intrinsically similar. Following Lewis commits me to a substantive metaphysical claim, but this claim must remain an assumption in the present context.¹⁴

I illustrated vacuity and unnaturalness using examples drawn from the domain of causal explanation. But they have application in the domain of ‘in virtue of’ explanation as well. Consider the example discussed above. You are justified in believing that there is a red light ahead in virtue of having a visual experience as of a red light ahead. Suppose we supplement the condition that you have the visual experience so that the result is a complex condition that suffices for your having the justification. What would this complex condition look like? I propose that either it will be unnatural or contain a vacuous condition.

We might start by adding the condition that you not know that you have taken a pill that makes green things look red. But if we go down that path, then we also have to add conditions such as the condition that you not think the lighting conditions are off, that you not think that you are hallucinating, that you not think that there is an invisible color inverting film between you and the light, etc. The set of conditions will be wildly heterogeneous.

Alternatively, we might note that though having a visual experience as of a red light ahead does not suffice for having justification for believing there is a red light ahead it does suffice for having *prima facie* justification for believing that there is a red light ahead. If this *prima facie* justification is undefeated, then you do have justification for believing that there is a red light ahead. Suppose, then, we just add to the condition that you have the visual experience the condition that the *prima facie* justification for which it suffices is undefeated. Call this extra condition the no defeaters condition. Adding it gives us a complex condition that suffices for your having the justification. But, it seems to me, the no defeaters condition suffers from vacuity. We already know from the explanandum that you have the justification, and all the no defeaters condition tells us is that nothing makes it the case that you fail to have the justification, without providing us with any illuminating conception of the range of things, the defeaters, that might do this.

The foregoing gives us some reason to think that (5) is true. But if (5) is true, then it is possible for Q to obtain in virtue of P even though P is not a sufficient condition for Q to obtain – i.e. Defeasibility is true. This is what I set out to show in this section.

constituents are items such as modes of presentation of properties, which are individuated by their role in reasoning. This second view of what an explanandum is might re-introduce worries about the relativity of vacuity.

¹³ See Lewis (1983, 1986).

¹⁴ For further discussion, both critical and favorable, see, for example, Taylor (1993) and Weatherston (2003).

2. *Perceptual knowledge*

By a piece of perceptual knowledge I mean a belief that amounts to knowledge because it is based on a perception. Suppose you see that there is a red light ahead. Suppose you take your visual perception at face value, and so believe that there is a red light ahead. And suppose, finally, that your belief that there is a red light ahead amounts to knowledge because it is based on your visual perception that there is a red light ahead. Then this belief is a piece of perceptual knowledge.

I will assume that we have some perceptual knowledge. It is not my aim here to address skeptical challenges to this claim. My aim is to take up an explanatory question. If a perception makes a belief based on it amount to knowledge, then there must be something about it in virtue of which it does so. The explanatory question I want to address is this:

In virtue of what does a perception make a belief based on it amount to knowledge, when it does so?

Your visual perception that there is a red light ahead, for example, makes your belief that there is a red light ahead amount to knowledge. This is the sort of fact that I want to explain.

Elements of the account of perceptual knowledge that I will explore here can be found in other views developed in the recent literature. Consider the following:

[McDowell] Suppose someone is presented with an appearance that it is raining. It seems unproblematic that if his experience is in a suitable way the upshot of the fact that it is raining, then the fact itself can make it the case that he knows that it is raining. But that seems unproblematic precisely because the content of the appearance is the content of the knowledge. And it is arguable that we find that match in content intelligible only because we do not conceive the objects of such experiences as in general falling short of the meteorological facts. (McDowell 1982, 213–214 in Dancy 1988)

[Fumerton] The suggestion, again, is that one has noninferential justification for believing P when one has the thought that P and when one is acquainted with the correspondence between the thought that P and the fact that P. The idea is that when one has immediately before consciousness both the truth bearer and the truth maker one has all one needs, all one could ever want, by way of justification. (Fumerton 2006, pg 65)

[Johnston] Sensory awareness discloses the truthmakers of our immediate perceptual judgments. Those truthmakers are external spatio-temporal particulars, which sensory awareness makes available for immediate demonstration. The structural elements (objects, stuff, their qualities, and the relations in which they stand) in those truthmakers are then recombined in immediate judgment . . . if I am seeing a spoon on the table, and judge accordingly, then I typically know that there is a spoon on the table . . . I typically know these things because the judgments in question are reliably formed from their respective truthmakers, which awareness makes manifest. (Johnston 2006, 282, 289).

There are significant differences among the views presented in these passages. McDowell thinks that experiences both have propositional contents and make us aware of truth-makers. Johnston seems to think that experiences just make us aware of truth-makers, and that it is only our judgments that have propositional contents. Further, McDowell and Fumerton think the truth-makers we are aware of are facts, and Johnston thinks that they are items, such as objects, stuffs, states, and events. While McDowell and Johnston focus on knowledge, Fumerton presents his view as an account of noninferential justification.

Though there are these differences, I expect that the family resemblances among the three views are apparent. The view that I will explore also bears family resemblances to them. I will call it the Simple View, since it is a simple answer to our explanatory question:

Simple View: If a perceptual experience makes a belief that p based on it amount to knowledge, it does so in virtue of 1) being an experience in which it perceptually appears to you that p, and 2) being an experience in which you are aware of something x, such that 3) x makes p true.

Some brief remarks intended to clarify the aims and components of the Simple View follow.

First, it is an account of knowledge, not justification. I doubt that you need to be aware of a truth-maker for p, say, in order to acquire perceptual justification for believing p. You might acquire perceptual justification for believing p from a convincing hallucination that p, in which you are not aware of a truth-maker for p. Maybe you need to *seem* to be aware of a truth-maker for p. This is an issue best dealt with elsewhere.¹⁵ My focus in this paper is on perceptual knowledge and the veridical, non-hallucinatory, perceptual experiences that give us this knowledge.

Second, it is an account of knowledge that is wholly, not just partly perceptual. Suppose you perceive that Jack is at the party and receive reliable testimony that if Jack is there, then Jill is too. You might thereby come to know that Jill is at the party, without, let us stipulate, being aware of a truth-maker for the proposition that Jill is at the party. In this case you have knowledge that p that is partly perceptual, though you have not experienced a truth-maker for p. This is not a counter-example to the Simple View, though, since again, your knowledge is only partly, not wholly perceptual.

Third, it makes two assumptions about three notions – perceptual content, perceptual awareness, and truth-making. The first assumption is that these notions are sound enough to be put to useful theoretical work. The second assumption is that in at least some cases the sorts of things of which we can be perceptually aware are the sorts of things that can be truth-makers for the sorts of propositions

¹⁵ I discuss perceptual justification and seeming awareness of truth-makers in Chudnoff (2011).

that we perceptually represent. The assumptions generate a theoretical burden: the accounts of perceptual content, perceptual awareness, and truth-making that underwrite the first assumption should also underwrite the second. I cannot discharge this theoretical burden here, but developments in the literatures addressing the three notions make me hopeful about the prospects of doing so.¹⁶

Fourth, it allows flexibility about the scope of perceptual knowledge. We have perceptual knowledge about the sensible properties of items in our immediate environment. But perceptual knowledge might extend beyond this. Suppose, for example, you can ‘see’ that your friend is angry, say because of the angry expression on his face. This case poses a challenge: you seem to have perceptual knowledge that your friend is angry without being aware of a truth-maker for the proposition that your friend is angry. One response to the challenge is to argue that your knowledge is only partly, not wholly perceptual, since it depends on background skills in or knowledge about reading facial expressions. Another response to the challenge is to argue that as a mature adult the reach of your perceptual awareness has extended beyond the merely sensible to include mental states of others. Both responses seem to me workable, and both are compatible with the Simple View.

So that is the Simple View of perceptual knowledge. It is *prima facie* plausible.¹⁷ In addressing the Gettier Problem problem in the balance of this paper, I will argue that further reflection on the Simple View does not undermine it. It has a good chance to be the correct view of perceptual knowledge.

3. Fake barns¹⁸

Goldman (1976) introduced fake barn cases into the Gettier literature in “Discrimination and Perceptual Knowledge”.¹⁹ Driving in the country-side with his son, Henry identifies various items that come into view: “That’s a tractor,” “That’s a cow,” and eventually, “That’s a barn”. The barn is clearly visible; Henry has good eyesight; etc. On the face of it, Henry knows that what he’s spotted is a barn. But, Goldman explains, Henry has unknowingly entered fake barn county, where the locals erect papier-mâché barn facsimiles, which invariably fool travelers. Though Henry in fact spots a barn, had he seen a barn facsimile, he would have mistaken it for a barn. In this case, we are strongly inclined to think Henry does not know that what he’s spotted is a barn.

¹⁶ Mulligan, Simons, and Smith (1984), and Johnston (2006) argue for views of truth-making and the objects of perceptual awareness that partially support my two assumptions.

¹⁷ I discuss what more can be said in its favor in section 4.

¹⁸ Though I cannot discuss them here, I believe that the points I make in this section carry over to Gettier cases that exploit social aspects of knowledge, such as those Harman (1973) discusses.

¹⁹ Carl Ginet created fake barn cases; Goldman’s paper is the first published discussion of them.

The facts of the situation seem to be this. Henry believes that *that's a barn* on the basis of a perceptual experience in which it perceptually appears to him that *that's a barn*, and in which he is aware of an item, say a barn-hood trope, that is a truth-maker for the proposition that *that's a barn*, but still his belief does not amount to knowledge that *that's a barn*. So the situation is one in which the conditions that the Simple View identifies as those in virtue of which Henry might know that *that's a barn* obtain, but in which Henry does not know that *that's a barn*.

If Q obtains in virtue of P only if P is a sufficient condition for Q to obtain, then we have two options: we can give up the Simple View, or we can give up our intuition that Henry does not know. But, if Defeasibility from section 1 is correct, then we have more options. We should distinguish between the following:

- A. The claim that the Simple View identifies a sufficient condition for perceptual knowledge.
- B. The claim that the Simple View identifies those conditions in virtue of which one has perceptual knowledge, when one does so.

If Defeasibility is correct, then we can accept (B) without (A), and, further, though we cannot accept both (A) and our intuition that Henry does not know, we can accept both (B) and our intuition that Henry does not know. This is the response to the fake barn cases that seems to me the most reasonable. The Simple View is attractive; our intuition that Henry does not know is strong; and Defeasibility, i.e. the claim that Q might obtain in virtue of P even though P does not suffice for Q to obtain, is defensible.

Two questions remain. First, even though in some cases Q might obtain in virtue of P even though P does not suffice for Q to obtain, why think that this is one of those cases? Second, even if we give up (A) and accept (B), why stop here; that is, why not think that even though the Simple View identifies those conditions in virtue of which one has perceptual knowledge when one does so, we should still take up the Gettier project and look for an illuminating account of those conditions that suffice for one to have perceptual knowledge?

I think both questions have the same answer: the conditions that one would have to add to those that the Simple View identifies in order to generate a set of conditions that suffices for perceptual knowledge are not explanatorily conditions. The reason why is that the extra conditions we might add are either vacuous or unnatural. For example, we might add the condition that you are not in a Gettier case. But this is surely vacuous since it does not provide us with any understanding of what makes a case into a Gettier case. Alternatively, we might start on constructing a set of sufficient conditions by adding the condition that you are not in fake barn county and proceeding from there. But there is reason to think that this is the beginning of an unnatural set of conditions since at this level of specificity

the sorts of situations we would have to rule out are just too heterogeneous. Finally, we might start to construct a set of conditions at a higher level of abstraction than one that mentions fake barn county, but which is not vacuous.²⁰ But there is reason to think that no such set of sufficient conditions exists. I do not have an argument from first principles that there is no such set of conditions. But the history of the Gettier literature provides good inductive evidence that there is none.²¹ If there is none, then we only have recourse to conditions that are not explanatory because they are either vacuous or unnatural.

If this is so, then it answers both questions. Recall (4) from the argument in section 1: If there are P, Q, and C such that Q obtains in virtue of P, P does not entail C, Q obtains in virtue of P only if C, but C is not explanatorily relevant to Q, then Q obtains in virtue of P even though P is not a sufficient condition for Q to obtain. If the only conditions C that can be added to those identified by the Simple View in order to arrive at conditions sufficient for perceptual knowledge are either the vacuous or unnatural ones mentioned above, then these conditions C are precisely such that one has perceptual knowledge in virtue of the conditions the Simple View identifies, the conditions the Simple View identifies do not entail C, one has perceptual knowledge in virtue of the conditions the Simple View identifies only if C, but C is not explanatorily relevant to one's having perceptual knowledge. And that means we do have a case in which Q (= having perceptual knowledge) obtains in virtue of P (= the conditions identified by the Simple View) even though P does not suffice for Q to obtain.

Now for the second question. I assume that the following maxims are plausible. First, given that we want to understand Q, we are interested in conditions that explain Q. Second, given that we want to understand Q, we need not be interested in conditions that suffice for Q. That is, just because a condition suffices for Q doesn't make it interesting to those who want to understand Q; there are lots of boring sufficient conditions for Q. These two maxims suggest that if the conditions that explain Q are sufficient for Q then we are indeed interested in some sufficient conditions for Q, but because they are explanatory, not because they are sufficient. If it turns out that the conditions that explain Q are insufficient for Q, then, it seems, there is no reason for us to be interested in some expanded, non-explanatory set of conditions that suffices for Q. But seeking such an expanded set of conditions is precisely what the Gettier project involves. And to that extent it is

²⁰ A natural idea is to add a reliability condition to conditions (1) through (3) in the Simple View. For example, we might add: 4) being a reliable indicator that p.

²¹ Consider, for example, the reliabilist proposal mentioned in note 27. It faces what Conee and Feldman have called the Generality Problem for Reliabilism; see Conee and Feldman (1998). One might try to complicate it in a way that avoids the problem. The danger, then, is that the complications will lead to unnaturalness. In my view, the history of the Gettier literature provides good inductive evidence that this will happen.

misguided. Of course, to see this we first have to see that the explanatory conditions are insufficient. And to see this we first have to think through various Gettier cases. So it is not as if the Gettier literature has not furthered our understanding. It has – not by providing conditions that explain why we have knowledge when we do, however, but by showing that those conditions that explain why we have knowledge when we do are not sufficient conditions for knowledge.

The foregoing observations constitute a reason to reject the task of supplementing the conditions identified by the Simple View so as to arrive at a set of conditions sufficient for perceptual knowledge. And to that extent they constitute an approach to the Gettier Problem problem restricted to perceptual knowledge. Since knowledge in general includes perceptual knowledge, however, if it is a mistake to look for Gettier-proof conditions for perceptual knowledge, then it is a mistake to look for Gettier-proof conditions for knowledge in general. So the approach to the Gettier Problem problem restricted to perceptual knowledge also addresses the general Gettier Problem problem.

Still, we might wonder about other specific types of knowledge, such as inferential knowledge, introspective knowledge, and knowledge by testimony. Perhaps there are Gettier-proof conditions for inferential knowledge, introspective knowledge, or knowledge by testimony.

It seems to me that the best way to proceed here is to examine each of these types of knowledge in turn. That is, we should ask, “In virtue of what does one have inferential knowledge, when one does?” or “In virtue of what does an inference make a belief based on it amount to knowledge, when it does so?” and similar questions about introspective knowledge, knowledge by testimony, and whatever other types of knowledge there are. Once we have articulated an attractive theory that answers the ‘in virtue of’ question for knowledge of type X, say, then we can go on to see whether it also identifies conditions that suffice for knowledge of type X. If it does not, and it really is an attractive explanatory theory, then we have reason to forego searching for Gettier-proof conditions for knowledge of type X. This piecemeal exploration of the various restricted Gettier Problem problems is suggested by but well beyond the scope of my present inquiry.

4. *Methodological issues*

In this section I conclude my discussion by taking up two methodological issues.

The first methodological issue is about what evidence can be given in favor of a theory of knowledge, or knowledge of some kind. There are two worries about my proposed solution to the Gettier Problem problem that make this issue pressing. The first is a general worry. It is that when we give up looking for necessary and sufficient conditions in favor of looking for explanatory conditions we are left at

methodological sea, with little to rationally constrain our choice of theory of knowledge. The second worry is more specific. It is that in defending the Simple View I have started down a slippery slope.²² Consider the *Simpler View* that if a perceptual experience makes a belief that p based on it amount to knowledge, it does so in virtue of being an experience in which it perceptually appears to you that p – and that’s it. Why couldn’t we defend the *Simpler View* against obvious counter-examples using the same strategies I have used to defend the Simple View?

The typical way to support a theory of knowledge, or knowledge of some kind, is to check it against intuitions about cases, i.e. intuitions about the presence or absence of knowledge in real or imagined situations. If the theory aims to give necessary and sufficient conditions, then this methodology makes sense. But if the theory aims to give explanatory conditions, this methodology is at best incomplete. The reason why is that if our aim is explanation then we are not just interested in whether knowledge is present or absent in a situation, but also in what explains its presence when it is present. So we need to supplement intuitions about cases with additional checks on theory.²³

General theoretical virtues, such as simplicity, comprehensiveness, naturalness, etc, suggest additional checks on theory. Our account of knowledge should maximize its general theoretical virtues. These additional checks are not enough, however. They will help us to pick between rival theories that accommodate the data. But what remains unclear is what, if not intuitions about cases, the data are.

The answer, I suggest, is that the data still consists of intuitions, just not intuitions about cases. These are not the only sorts of intuitions that we have. We also have what I will call intuitions about grounds. Let us say that if Q obtains in virtue of P , then P grounds Q . Now consider a property – e.g., knowing that p . Intuitions about cases are intuitions about its application in real or imagined situations. Intuitions about grounds are intuitions about what grounds its application when it does apply in a situation. They include intuitions about what conditions ground the application of a property in a situation. And more importantly, it seems to me, they include intuitions about what features the conditions in virtue of which a property applies in a situation should have, i.e. what those conditions must be like. Accounts that aim to articulate conditions that explain why one knows when one does should be checked primarily against intuitions about grounds.

Consider the case of perceptual knowledge. The Simple View identifies conditions in virtue of which a perception makes a belief based on it amount to knowledge, when it does so. What reason do we have to believe the Simple View? It is not that it fits our intuitions about cases. It doesn’t, and it doesn’t aim to. It is

²² I thank an anonymous reviewer for first drawing my attention to this worry.

²³ By a “check on theory” I mean a consideration that should be taken into account in assessing the plausibility of a theory. The traditional checks on theories of knowledge are the propositional contents of our intuitions about cases.

not that it possesses many general theoretical virtues. Lots of theories completely irrelevant to perceptual knowledge do as well. The main reason we have to believe it is that it is supported by our intuitions about the grounds of perceptual knowledge.

What *are* our intuitions about the grounds of perceptual knowledge? We might have many. The following seem central:

- A. If a perceptual experience makes a belief based on it amount to knowledge, it does so partly in virtue of inviting you to have that very belief – i.e., partly in virtue of representing the environment as being the way you come to believe it to be.
- B. If a perceptual experience makes a belief based on it amount to knowledge, it does so partly in virtue of successfully fulfilling its function of making you aware of items in your environment.²⁴
- C. If a perceptual experience makes a belief based on it amount to knowledge, it does so partly in virtue of guaranteeing the belief's truth.²⁵

Once articulated, these claims seem intuitive. At least they seem true, and they don't seem true just because they follow from more basic claims about the grounds of perceptual knowledge.

Further, they lend support to the Simple View. Here it is for reference:

Simple View: If a perceptual experience makes a belief that *p* based on it amount to knowledge, it does so in virtue of 1) being an experience in which it perceptually appears to you that *p*, and 2) being an experience in which you are aware of something *x*, such that 3) *x* makes *p* true.

Intuition (A) supports condition (1); intuition (B) supports condition (2); intuition (C) supports condition (3). The intuitions do not *entail* the Simple View. They are intuitions about partial grounds, and the Simple View purports to identify the full ground of perceptual knowledge. In order to further support the Simple View one would at least need to compare its theoretical virtues with other views that fit the intuitions about grounds, and, additionally, explore the Gettier Problem problem restricted to perceptual knowledge to make sure that the fact that the grounds the Simple View cites are not sufficient conditions for knowledge is acceptable. I believe I have made progress on the second task. The first lies beyond the scope of my present inquiry. The present goal, anyway, isn't to demonstrate the Simple View; it is to illustrate how intuitions about grounds might be appealed to in support of it.

²⁴ For a fuller discussion of the function of perceptual experience, and the importance this might have in thinking about the epistemology of perception see Johnston (2006).

²⁵ One might challenge this intuition and argue that the perceptual experience need only reliably correlate with truth, or be sensitive to the truth, or have some other property that falls short of guaranteeing truth. This is an area for further exploration.

So the resources that we can appeal to in defending an account of knowledge, or knowledge of some kind, include (perhaps least importantly) intuitions about cases, measures of general theoretical virtue, and (perhaps most importantly) intuitions about grounds. Nothing ensures that these will determine a unique theory. But recognizing the extent of these resources should defuse the two worries mentioned above. The first worry was that when we give up looking for necessary and sufficient conditions in favor of looking for explanatory conditions we are left at methodological sea, with little to rationally constrain our theorizing. This isn't true, since we can appeal to our intuitions about grounds. The second worry was that in defending the Simple View I have appealed to strategies that could just as well be appealed to in defending patently false theories, such as what I called the Simpler View that if a perceptual experience makes a belief that p based on it amount to knowledge, it does so in virtue of being an experience in which it perceptually appears to you that p – and that's it. This isn't so, however. The reason why is that the Simpler View does not just violate our intuitions about cases; it also violates our intuitions about grounds.

The second methodological issue that I want to discuss is posed by this question. Is investigating the nature of justification an important part of investigating the nature of knowledge? To address it, let us consider another. Why might reflection on Gettier cases suggest that investigating the nature of justification is *not* an important part of investigating the nature of knowledge? Gettier cases, let us say, show that knowledge is not correctly analyzed as justified true belief – i.e. $K \neq JTB$. Here are three different conclusions one might draw from this.

1. Knowledge is correctly analysed as justified true belief and something else – i.e. $K = JTB + X$.²⁶
2. Knowledge is correctly analysed as true belief and something else, not justification – i.e. $K = TB + X$.²⁷
3. Knowledge cannot be correctly analysed – i.e. K is basic.²⁸

Assume that investigating the natures of things involves giving analyses, or necessary and sufficient conditions. Then it is clear why conclusion (2) would motivate shifting focus away from justification in investigating knowledge. A natural response to conclusion (3) is to stop trying to analyse knowledge and instead use it in the analysis of other things, i.e. to engage in 'knowledge first' epistemology.²⁹ Even conclusion (1) might play a role in shifting focus away from justification.

²⁶ There are many views of this sort. Lycan's preferred analysis of knowledge is an example; see Lycan (2006).

²⁷ Again, there are many views of this sort. Nozick's analysis of knowledge is an example; see Nozick (1981).

²⁸ Cf. Williamson (2000).

²⁹ Again, cf. Williamson (2000).

One might reason that since the necessity of justification is clear enough, what we have to do now to finish our investigation of knowledge is explore what condition X might be.

So these are some ways that reflection on Gettier cases might shift focus away from justification in investigations of knowledge. The approach to the Gettier Problem problem developed in the previous sections raises the prospects of alternative understandings of Gettier cases, on which investigation of justification has a central role to play in investigating knowledge. Set aside the assumption that investigating the natures of things involves giving analyses, or necessary and sufficient conditions. Suppose instead that to investigate knowledge, for example, is to seek those conditions that *explain* why one knows, when one does so. Given what I have argued in the previous sections, Gettier cases leave open the following theoretical option: though knowledge is not correctly analysed as justified true belief, instances of knowledge are explained by instances of justified true belief. That is, whenever S knows that p, S knows that p in virtue of the fact that S has a justified true belief that p. This view places justification back in the center of things.

My own view is that the connection between knowledge and justification is more complicated. Suppose, for example, that the Simple View of perceptual knowledge is correct. Then if S has perceptual knowledge that p, S has perceptual knowledge that p in virtue of basing his belief that p on a perceptual experience in which it perceptually appears to S that p, and in which S is aware of something x, such that x makes p true. So there is more to the explanation of S's perceptual knowledge than that S has a justified true belief that p. Still, there might be a connection to justification. S's perceptual experience also justifies him in believing that p. In virtue of what? I think that it does so in virtue of *being* an experience in which it perceptually appears to S that p, and in which S *seems* to be aware of a truth-maker for p.³⁰ If this is correct, then the conditions in virtue of which S's perceptual experience make S's belief amount to knowledge stand in a special relation to the conditions in virtue of which S's perceptual experience justifies S's belief: the latter conditions make things seem to S a certain way, and the former conditions obtain when things *are* that way. In my view this connection is an instance of a more general pattern of connections among justification, knowledge, things seeming a certain way, and things being the way they seem.³¹

Even if neither view proves correct – the view that whenever S knows that p, S knows that p in virtue of the fact that S has a justified true belief that p, or the more complicated view gestured toward in the previous paragraph – it should be clear that the existence of Gettier Cases itself provides no grounds for shifting

³⁰ I defend this view in Chudnoff (2011).

³¹ I explore this more general pattern in (Chudnoff forthcoming).

focus away from justification when exploring the nature of knowledge. To motivate such a shift one would need to rule out the possibility of justification playing an explanatory role in accounting for knowledge.*

REFERENCES

- AUDI, P. ms, 'Grounding: Toward a Theory of the In-Virtue-of Relation'.
- CHUDNOFF, E. 2011, 'The Nature of Intuitive Justification', *Philosophical Studies* **153**, 2, pp. 313–333.
- CHUDNOFF, E. forthcoming, 'Presentational Phenomenology', *Protosociology*.
- CHUDNOFF, E. ms, 'Grounding and Entailment'.
- CONEE, E. and FELDMAN, R. 1998, 'The Generality Problem for Reliabilism', *Philosophical Studies* **89**, 1, pp. 1–29.
- DANCY, J. 1988, *Perceptual Knowledge*, Oxford: Oxford University Press.
- DANCY, J. 2006, *Ethics Without Principles*, Oxford: Oxford University Press.
- FINE, K. ms, 'Guide to Ground'.
- FUMERTON, R. A. 2006, *Epistemology*, Oxford: Wiley-Blackwell.
- GOLDMAN, A. I. 1976, 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* **73**, 20, pp. 771–791.
- HARMAN, G. 1973, *Thought*, Princeton: Princeton University Press.
- JOHNSTON, M. 2006, 'Better Than Mere Knowledge? The Function of Sensory Awareness', in: T. S. Gendler and J. Hawthorne, eds, *Perceptual Experience*, Oxford: Oxford University Press, pp. 260–290.
- LEWIS, D. 1983, 'New Work for a Theory of Universals', *Australasian Journal of Philosophy* **61**, 4, pp. 343–377.
- LEWIS, D. 1986, *On the Plurality of Worlds*, Oxford: Blackwell.
- LYCAN, W. G. 2006, 'On the Gettier Problem Problem', in: S. Hetherington, ed., *Epistemology Futures*, Oxford: Oxford University Press, pp. 148–168.
- McDOWELL, J. 1982, 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, **68**, pp. 455–479.
- MULLIGAN, K., SIMONS, P. and SMITH, B. 1984, 'Truth-makers', *Philosophy and Phenomenological Research* **44**, 3, pp. 287–321.
- NOZICK, R. 1981, *Philosophical Explanations*. Cambridge, MA: Belknap Press.
- QUINE, W. V. 1951, 'Two Dogmas of Empiricism', *Philosophical Review* **60**, pp. 20–43.
- ROSEN, G. 2010, 'Metaphysical Dependence: Grounding and Reduction', in: B. Hale and A. Hoffman, eds, *Modality: Metaphysics, Logic and Epistemology*, Oxford: Oxford University Press, pp. 109–136.
- SALMON, W. 1984, *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.
- SALMON, W. 1989, *Four Decades of Scientific Explanation*, Minneapolis: University of Minnesota Press.
- SCHAFFER, J. 2009, 'On What Grounds What', in: D. Chalmers, D. Manley and R. Wasserman, eds, *Metametaphysics: New Essays on the Foundation of Ontology*, Oxford: Clarendon Press, pp. 347–383.

* I thank Avram Hiller, Uriah Kriegel, Kirk Ludwig, Anne Meylan, Kevin Morris, Greg Ray, and an anonymous reviewer for helpful comments on earlier drafts of this paper. I also thank Troy Cross, Kit Fine, Tatjana von Solodkoff, Nick Stang, and Richard Woodward for helpful feedback on (Chudnoff ms), which overlaps section 1.

- STREVENS, M. 2008, *Depth: An Account of Scientific Explanation*, Cambridge, MA: Harvard University Press.
- TAYLOR, B. 1993, 'On Natural Properties in Metaphysics', *Mind* **102**, 405, pp. 81–100.
- WEATHERSON, B. 2003, 'What Good are Counterexamples?', *Philosophical Studies* **115**, 1, pp. 1–31.
- WILLIAMSON, T. 2000, *Knowledge and Its Limits*, New York: Oxford University Press.
- WOODWARD, J. 2003, *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press.