Warrant and accidentally true belief

ALVIN PLANTINGA

My gratitude to Richard Greene and Nancy Balmert for their perceptive discussion of my account of warrant ('Two notions of warrant and Plantinga's solution to the Gettier problem', ANALYSIS, this issue). Actually, they begin by noting *two* accounts: an initial description or location as 'that which when added to true belief yields knowledge' (k-warrant), and then the account I give of what I think the item so located really is (p-warrant). There is a minor glitch *ab initio*, however: I said that warrant is that quantity (whatever precisely it is) *enough* of which is what distinguishes knowledge from mere true belief. The point here is that warrant comes in degrees. A belief constitutes knowledge only if it has a certain degree of the quantity in question; a true belief that enjoyed at least some warrant could still fail to be knowledge if it didn't have a sufficient degree of warrant (Greene and Balmert seem to me to stumble over this in their footnote 5, p. 135).

1. *Accidentally true beliefs*

Greene and Balmert ask '... is it possible for a belief which meets Plantinga's conditions for p-warrant to still be true merely by accident?' (134) They go on to point out that 'if there can be a false p-warranted belief, and deductive inference is p-warrant preserving, then there can be p-warranted beliefs which are true by accident' as in the classic Gettier case they present. That would of course be a problem for my account; a belief 'true by accident' (at any rate in the relevant sense of that phrase) clearly fails to constitute knowledge.[1] They go on to insist that this *is* a problem for my account, because I hold that false belief *can* have warrant: 'Thus, assuming that p-warrant is preserved under deductive inference, it is evident that Plantinga's account is one that allows for false p-warranted beliefs, and as such, does not overcome the Gettier problem' (135).

By way of reply: first, it is important in this context to remember that (both in reality and according to me) warrant comes in *degrees*. I am committed to the thought that false beliefs can have *some* warrant, but not to the thought that they can have warrant *sufficient for knowledge*. As I'll

[1] They also suggest that there is a problem here for the initial rough and ready identification of warrant (k-warrant): 'if false k-warranted belief exists, then k-warrant, when added to true belief, could not possibly be sufficient for knowledge' (135). Here see Trenton Merricks 1995, and the reply by Sharon Ryan 1996 and Merricks's response, forthcoming.

ANALYSIS 57.2, April 1997, pp. 140–145.

argue later, they can't. And second, Greene and Balmert's assumption that warrant (either p- or k-) is closed under deductive inference seems to me to be false. More exactly, what seems to me to be false is the stronger assumption they need for their argument: that *degree* of warrant is closed under deductive inference. We could put this assumption as follows:

(a) For any person *S*, propositions *p* and *q*, and degree of warrant *d*, if *p* has *d* for *S* and *S* properly infers[2] *q* from *p*, then *q* also has *d* for *S*.

(a) seems clearly false. If it were true, then the conclusion of a complex and lengthy deductive inference would always have the same degree of warrant, for the inferrer, as do the premises; but we've known since Descartes that that's false.

Of course Greene and Balmert's counterexample doesn't really require anything quite as strong as (a). It requires only the possibility that (1) a false proposition *p* have warrant sufficient for knowledge, and (2) on some occasion a person *S* infer *q* from *p*, where *q* is true and has no source of warrant other than *p*, and where this particular inference (perhaps by virtue of its clarity and simplicity) preserves warrant. If (1) were possible, I'd have no inclination to dispute (2). But I doubt that (1) *is* possible, and in any event its possibility does not follow from my account of warrant.

So do I escape unscathed? By no means. Their counterexample as stated doesn't quite do the job; but there are others that do. Robert Shope (forthcoming), Richard Feldman (1996), and Peter Klein (1996) all propose counterexamples that really do show that a belief could meet my conditions and still be true just by accident. Here is a composite of their examples. I own a Chevrolet van, drive to Notre Dame on a football Saturday, and unthinkingly park in one of the many spaces reserved for the football coach. Naturally his minions tow my van away and, as befits such lese-majesty, destroy it. By a splendid piece of good luck, however, I have won the Varsity Club's Win-a-Chevrolet-Van contest, although I haven't yet heard the good news. You ask me what sort of automobile I own; I reply, both honestly and truthfully, 'A Chevrolet van'. My true belief that I own such a van is true just by accident (more accurately, it is only by accident that I happen to form a true belief), however, and hence does not constitute knowledge.

What is important about the example is this: if the coach's minions had been a bit less zealous and had *not* destroyed my van, the conditions for

[2] I.e., *q* does follow from *p* and *S* sees that it does; we may add, if we like, that *q* follows by some principle of logic with which *S* is acquainted. These would be necessary conditions for propriety; no doubt they aren't sufficient, but we needn't tarry now over the question just what propriety consists in.

warrant (p-warrant) would have obtained and I would have known that I owned a Chevrolet van. But in the actual situation, the one in which the van is destroyed, my belief is produced by the very same processes functioning the very same way in the same cognitive environment. Hence, on my account, either both of these situations are ones in which I know that I own a Chevrolet van, or neither is. But clearly one is and the other isn't. Therefore my account is defective.[3] (Greene and Balmert make substantially this point on p. 137)

2. *The resolution problem*

Consider another Gettier example, this one, oddly enough, antedating Gettier's birth (it was proposed by Bertrand Russell). I glance at a clock, forming the opinion that it is 3:43 p. m. As luck would have it, the clock stopped precisely 24 hours ago. The belief I form is indeed true, but it is true 'just by accident' (the clock could just as well have stopped an hour earlier or later); it does not constitute knowledge. Again, however, if the clock had been running properly and I had formed the same belief by the same exercise of cognitive powers, I would have known; here we have another example that apparently refutes my account. Another example: I am not aware that Paul's look-alike brother Peter is staying at his house; if I'm across the street, take a quick look and form the belief that Paul is emerging from his house, I don't know that it's Paul even if in fact it is (it could just as well have been Peter emerging); and again, if Peter hadn't been in the neighbourhood, I would have known.

What is crucial, in each of these cases, is that my cognitive faculties display a certain *lack of resolution*. I am unable, by a quick glance, to distinguish the state of affairs in which the clock is running properly and telling the right time from a state of affairs in which it stopped just 12 or 24 hours earlier. I cannot distinguish Paul from Peter just by a quick look from across the street. Of course this lack of resolution is in each case relative to the particular exercise of cognitive powers in question. If I had watched the clock for 10 minutes, say, I would have known that it isn't running, and if I had walked across the street and taken a good look, I'd have known that it wasn't Paul but Peter at the door.

3. *Mini- vs. maxi- cognitive environments*

What I can't distinguish by those exercises of my epistemic powers are different *cognitive mini-environments*. In my 'Respondeo' (1996), I develop the distinction between cognitive maxi-environments and cognitive mini-environments; I won't repeat that development here. Suffice it to

[3] For fuller development here, see my 1996: 314 ff.

say the following. First, a cognitive maxi-environment is more general and more global than a cognitive mini-environment. Our cognitive maxi-environment here on earth would include such macroscopic features as the presence and properties of light and air, the presence of visible objects, of other objects detectable by cognitive systems of our kind, of some objects not so detectable, of the regularities of nature, of the existence and general nature of other people, and so on. Our cognitive faculties are designed (by God or evolution) to function in *this* maxi-environment, or one like it. They are not designed for a maxi-environment in which, e.g., there is constant darkness, or where everything is in a state of constant random flux, or where the only food available contains a substance that destroys short term memory, or where there aren't any distinguishable objects, or no regularities of a kind we can detect; in such an environment they will not perform their function of providing us with true beliefs. Now a given cognitive maxi-environment can contain many different mini-environments – e.g., the one where the clock stops, but also one where it doesn't; the one where Peter is visiting Paul, but also one where he isn't; the one where the coach's minions destroy my van, but also one where they mercifully temper the punishment I so richly deserve, and magnanimously do no more than paint the windshield black.

And now here's the point: some cognitive mini-environments – e.g., those of the Notre Dame van case, the clock that stopped, Peter's visit to Paul – are *misleading* for some exercises of cognitive faculties, even when those faculties are functioning properly and even when the maxi-environment is favourable. The maxi-environment is right, but the mini-environment isn't; in those mini-environments the cognitive faculties in question (more exactly, particular exercises of the cognitive faculties in question) can't be counted on to produce true beliefs.

The basic idea is this: our cognitive faculties have been designed for a certain kind of maxi-environment. Even within that maxi-environment, however, they don't function perfectly (they sometimes produce false belief), although they do function reliably. (Perhaps perfectly functioning cognitive faculties would require too much brain size, thus interfering with other desiderata.) In some mini-environments, therefore, they can't be counted on to produce a true belief: if they do, it is just by accident and does not constitute knowledge. So even if the maxi-environment is favourable and the other conditions of warrant are met, a belief could still be true 'just by accident', thus not constituting knowledge.

It is clear, therefore, that *S* knows *p*, on a given occasion, only if *S*'s cognitive mini-environment, on that occasion, is not misleading – more exactly, not misleading with respect to the particular exercise of cognitive powers producing the belief that *p*. So the conditions of warrant need an

addition: the maxi-environment must indeed be favourable or appropriate, but so must the cognitive mini-environment. Just what does 'appropriateness' or 'favourability', or 'non-misleadingness', for a cognitive mini-environment, consist in: can we say anything more definite? Intuitively, a mini-environment is favourable, for an exercise of cognitive powers, if that exercise *can be counted on* to produce a true belief in that mini-environment. Perhaps this is as specific as we can sensibly get; but in 'Respondeo' I suggested something a bit more precise. Where *B* is a belief, *E* the exercise of cognitive powers that produces *B* and *MBE* the mini-environment with respect to *B* and *E*, say that

(F) *MBE* is *favourable* for *E* if and only if, if *S* were to form a belief by way of *E* in *MBE*, *S* would form a true belief.[4]

What I then added to the other conditions of warrant is the *Resolution Condition*:

(RC) A belief *B* produced by an exercise *E* of cognitive powers has warrant only if *MBE* is favourable for *E*.

But here I must add a *caveat*: (RC) is to be understood as stipulating that a belief *has a degree of warrant sufficient for knowledge* only if its cognitive mini-environment is favourable. The thought is not that a belief produced in an unfavourable mini-environment has no warrant at all, but only that it doesn't have a degree of warrant sufficient for knowledge. And now I can also discharge a promissory note issued on p. 140: (RC) guarantees that no *false* belief has a degree of warrant sufficient for knowledge. No mini-environment in which *S* forms a false belief is favourable for the exercise of cognitive powers producing it, because, of course, the relevant counterfactual would have a true antecedent and a false consequent.

The overall picture, then, is as follows. Our faculties are designed for a certain kind of cognitive maxi-environment, one that sufficiently resembles the one in which we do in fact find ourselves. And when a belief is formed by properly functioning faculties in an environment of that sort (and the bit of the design plan governing its production is successfully aimed at truth), then the belief in question has *some* degree of warrant, even if it happens to be false. (Perhaps there are different degrees of warrant available here, so that even false beliefs can display different degrees of warrant.) But our cognitive faculties are not maximally effective – not only in that there is much we aren't capable of coming to know, but also in that we are

[4] Here I am assuming (contrary to the usual semantics for counterfactuals) that truth of antecedent and consequent is not sufficient for truth of the counterfactual (a counterfactual can be false even if it has a true antecedent and a true consequent). What is also required is that there be no sufficiently close possible world in which the counterfactual has a true antecedent and false consequent.

sometimes prone to err, even when the maxi-environment is right and the relevant faculties are functioning properly. Another way to put the same point: within a favourable cognitive maxi-environment, there can be mini-environments for a given exercise of our faculties, in which it is just by accident that a true belief is formed, if one is indeed formed. A true belief formed in such a mini-environment doesn't have warrant sufficient for knowledge, even if it has some degree of warrant. To achieve that more exalted degree of warrant, the belief must be formed in a mini-environment such that the exercise of the cognitive powers producing it can be counted on to produce a true belief. Hence the resolution condition. Beliefs that meet all of the conditions will then constitute knowledge (provided they are accepted with sufficient firmness). But even beliefs that constitute knowledge can still differ with respect to their degree of warrant. (I know that 2 + 1 = 3 'better', or 'more firmly', or whatever, than that the population of China exceeds that of the US, even though I do know both.) So degree of warrant among propositions with warrant sufficient for knowledge can still vary: my proposal is that it varies with strength or firmness of belief.

University of Notre Dame
Notre Dame, IN 46556, USA
Alvin.Plantinga.1@nd.edu

References

Feldman, R. 1996. Plantinga, Gettier and warrant. In Kvanvig 1996.

Klein, P. 1996. Warrant, proper function, reliabilism, and defeasibility. In Kvanvig 1996.

Kvanvig, J., ed. 1996. *Warrant in Contemporary Epistemology; Essays in Honor of Plantinga's Theory of Knowledge*. New York: Roman and Littlefield.

Merricks, T. 1995. Warrant entails truth. *Philosophy and Phenomenological Research* 55: 841–55.

Merricks, T. Forthcoming. More on warrant entails truth. *Philosophy and Phenomenological Research*.

Plantinga, A. 1996. Respondeo. In Kvanvig 1996.

Ryan S. 1996, Does warrant entail truth? *Philosophy and Phenomenological Research* 56.

Shope, R. Forthcoming a. Gettier problems. *Routledge Encyclopedia of Philosophy*.

Shope, R. Forthcoming b. *Knowledge as Power*.