

Knowledge is Normal Belief¹

Abstract

In this paper I offer a new analysis of knowledge: knowledge, I claim, is normal belief. I begin with what I take to be the conceptual truth that knowledge is epistemically justified, or permissible, belief. I then argue that this in turn is simply doxastically normal belief, first clarifying what is meant by this claim, and then providing reasons to think that normal belief, so understood, must be true and safe from error, making it a good candidate for knowledge.

Keywords

Analysis of Knowledge, Norm of Belief, Epistemic Justification, Normal Function

1. Introduction

Does knowledge have an analysis? According to an increasingly widespread view, it does not.

Williamson, for instance, claims that

On quite general grounds, one would not expect the concept knows to have a non-trivial analysis in somehow more basic terms. Not all concepts have such analyses, on pain of infinite regress; the history of analytic philosophy suggests that those of most philosophical interest do not. (2000: 31)

Williamson also offers more specific reasons for thinking that the concept of knowledge in particular cannot be analysed,² before suggesting that, in light of these considerations, we have no rational basis for optimism about the prospects of articulating conditions which are even necessarily co-extensive with the concept of knowledge (2000: 30). Nevertheless, he concedes, '[c]orrect deep analyses would doubtless still be interesting if they existed; what has gone is the reason to believe that they do exist' (2000: 32).³

In this short note I offer a new analysis of knowledge: knowledge, I claim, is normal belief. In what follows I elucidate this claim and offer some support for it. It should be stressed at the outset, however, that I do not so much as attempt to provide an analysis of the *concept* of knowledge; for

¹ Thanks to Corine Besson, Patrycja Kaszynska, Stephen Kearns, and an anonymous referee, for valuable comments on earlier drafts of this paper.

² More precisely, Williamson provides two such reasons: first, he argues that the concept of knowledge is a mental concept, while any purported analysis will employ non-mental concepts – notably *truth* and (perhaps) *justification* (2000: 30); and second, he claims that '[e]xperience confirms inductively... that no analysis of the concept *knows* of the standard kind [in terms of belief] is correct' (2000: 30).

³ This despite his suggestions that '[t]he pursuit of analyses is a degenerating research programme' (2000: 31), and that 'by abandoning that fruitless search [for an analysis of knowledge] we can gain new insight into epistemological problems' (2000: 5).

all that I will say, it is primitive, and is acquired by example through comparison with paradigms (Williamson 2000: 3, 22), and not by definition through the combination of simpler concepts. Rather, my analysis is presented as a metaphysical account of the underlying nature of knowledge itself.

I begin by defending what I take to be the conceptual truth that knowledge is epistemically justified, or permissible, belief. Then, appealing to Williamson's own thesis that knowledge is the norm of belief (2000: 47), I suggest that necessarily, a belief is epistemically permissible just in case it is doxastically 'normal' in a sense adapted from Millikan's (1984) work. I argue that beliefs meeting this condition are both true and safe from error, thus providing support for the thought that my proposed analysis is correct. Moreover, since the analysis is framed in naturalistic terms it promises robust philosophical rewards, and is therefore both 'deep' and 'interesting'.

2. Knowledge is Permissible Belief

Is knowledge justified true belief? Most philosophers will consider this issue to have been settled - negatively - some time ago, by Gettier's influential (1963) contribution to the debate. But we must distinguish two understandings of the thesis in contention; more specifically, we must consider two readings of the word 'justified' occurring in it.

On one way of thinking, a belief is justified just in case it has some evidential support – that is, just in case there are some propositions to which the subject of the belief has access and which provide evidence, in some degree, for the truth of the proposition believed. On this understanding of justification, Gettier refuted the thesis that knowledge is justified true belief: his counterexamples clearly showed that it is possible to have beliefs which are true, and for which one has some (high degree of) evidence, and yet fail to know.

But Gettier subjects do not have, in Ayer's memorable (1956) phrase, 'the right to be sure'. Indeed, this expression suggests an altogether different understanding of the term 'justified': to be justified is, in general, to be within one's rights; and to be justified in believing is, on this way of thinking, to believe rightly, permissibly, in an acceptable or admissible manner. The Gettier cases do not refute the thesis that knowledge is justified true belief when 'justified' is understood in this second way; for the beliefs of Gettier subjects are not justified in this sense. More specifically, Gettier subjects have acted (doxastically) in a way which fails to ensure the truth of their beliefs; in this sense, they have acted irresponsibly, unjustifiably, in believing as they do.

I think that, when the claim is properly understood, with 'justified' interpreted in this second way, knowledge is justified true belief. But justification, so understood, entails truth. Indeed, Zagzebski (1994) has shown that any attempt to avoid the Gettier counterexamples must concede this much: if

knowledge is justified true belief, then it is justified belief; the appeal to truth is redundant.⁴ Of course, redundancy is not so harmful as error. Nonetheless, it serves no one's interests to use contentious language, which can be easily misunderstood; confusion may arise. In light of what has been said so far we may accordingly reformulate the thesis under consideration thus: knowledge is permissible belief.

If there is something wrong with the traditional analysis of knowledge, it is not that its central thesis is false: indeed, I think this thesis is not only true, it is a conceptual truth; more specifically, it is *a priori* that it is necessary that a belief is permissible if and only if it constitutes knowledge.⁵ Still, one might be concerned that an account of knowledge as permissible belief cannot serve to ground an analysis: for the notion of *epistemic* permissibility tacitly involves that of knowledge; the account is therefore not reductive.⁶ While this concern is, I think, essentially correct, it will not affect my final proposal: for the claim that knowledge is permissible belief is only the starting point.

3. Knowledge is the Norm of Belief

Knowledge is permissible belief. How is this claim to be understood? The key notion, of course, is that of permissibility. It seems to involve two notable features: first, it comprises a modal element; and second, it has a normative character. Some clarification of each of these aspects of the permissibility in question must be sought.

The recognitions, first, that knowledge is permissible belief, and second, that permissibility is a normative modal notion, may suggest that *s knows p* is equivalent to *s may believe p, or it is permissible for s to believe p*. But the proposed equivalence fails. The reason is simple: that one knows *p* entails that one believes *p*, whereas that it is permissible for one to believe *p* does not. We may put our point another way: knowing is not merely a matter of *having* the right to be sure - to have this right is, as Williamson sometimes puts it, to be 'in a position to know' (2000: 93); but knowing requires more - it requires *exercising* that right, *being* rightly sure.

Nor can this problem be easily overcome; though it might be thought that it could. Consider an analogy: for a proposition to be contingently true is for it to be both true and possibly untrue. Similarly, one might suppose, to know - that is, to permissibly believe - is to believe, and for it to be permissible for one to do so. However, while knowledge entails both permission to believe and

⁴ Zagzebski assumes, however, that justification does not entail truth; she concludes that the Gettier counterexamples can't be avoided. By contrast, Sutton (2007) acknowledges that knowledge is justified belief.

⁵ See Williamson (2007: 51; especially footnote 3) for an account of analyticity along these lines.

⁶ Williamson says, 'if someone insists that knowledge is justified true belief... the problem is likely to be that no standard of justification is supplied independent of knowledge itself' (2000: 4).

belief, it cannot be equated with their conjunction: for one might have evidence such that, were one to believe that p on its basis one would thereby come to know that p, and yet believe that p on some other, illegitimate basis; in such a case one would believe that p, and it would be permissible for one to believe that p, and yet one would not know that p. We cannot reduce knowing to a conjunction of permission and belief.

We may, however, begin to make headway by reflecting on the thesis that knowledge is the norm of belief; or, as Williamson has put it, '[k]nowledge sets the standard of appropriateness for belief' (2000: 47).⁷ This suggests that belief ought, in some sense, to be knowledge. But how are we to spell this out? Comparing the claim that knowledge is the norm of belief with the claim that it is the norm of assertion offers an insight. Williamson articulates this latter thesis as the claim that '[o]ne must: assert p only if one knows p' (2000: 243). This suggests by analogy, then, that in the case of belief what one must do is believe p only if one knows p; that is, that one ought to either not believe p or else know p. And this is plausible. Take a case in which one falsely believes a proposition p. There is no good sense to be attached to the claim that one ought, in that case, to know p - no sense on which it is true. If 'ought' implies 'can' then it is simply false that one ought to know p; knowledge is factive, and so one cannot (in any relevant sense, keeping fixed such factors as the truth value of p) know p. Rather, one ought not to believe p in such a case. Our explication of the claim that knowledge is the norm of belief provides a good account of this otherwise difficult case; for we may conform to the norm so understood, not by knowing the proposition p, but by failing to believe it.

One might fear, however, that such an understanding of the knowledge norm for belief will be of no service in the project of analysing knowledge. For clearly, both the notions of belief and knowledge occur within the scope of the normative modal operator: what is obligatory is that a certain relation between belief and knowledge should obtain; and it is therefore far from obvious that we may explain what one of them is in terms of the other. Nonetheless, this is exactly what I think can be done; but seeing this requires a particular understanding of the normativity involved. In the next section I aim to provide such an understanding, and to substantiate not only the claim that knowledge is the norm of belief, but that it is normal belief.

4. Knowledge is Normal Belief

The claim that knowledge is the norm of belief is the claim that one ought to believe p only if one knows p. I take this to be (necessarily) equivalent to the claim that, normally, one believes p only if one knows p. Of course, I do not understand 'normally' here in any simple statistical sense: I am not

⁷ Williamson also says, 'the point of belief is knowledge' (2000: 1); '[k]nowing is... the best kind of believing' (2000: 47); and 'belief aims at knowledge' (2000: 47).

claiming that typically, or usually, one believes *p* only if one knows *p*; one may have wandered into an environment in which one is usually mistaken.⁸ Rather, what I mean is that, in certain normal cases, one believes *p* only if one knows *p*. And I submit if that's true, the best explanation of this fact is that, in those normal cases, one's beliefs constitute knowledge.

How are the normal cases to be defined? I follow Millikan (1984) in taking a device (type) to be functioning normally (or properly) if, and only if, it behaves as it did in those cases for which it was selected - that is, as it did in those cases in which it proved beneficial, and as a result of which it was reproduced.⁹ We may then say that a case, or circumstance, in which a device is employed is normal just in case that device functions normally in that case.

Why, then, should we believe, or even take seriously the hypothesis, that our two claims, articulated above, are (necessarily) equivalent? To begin with, we should reflect on the fact that both 'it ought to be the case that' and 'normally' are modal operators which may be understood as universal quantifiers over circumstances: in particular, the former quantifies over epistemically permissible circumstances, while the latter quantifies over the cases in which some device type functions normally. If the circumstances meeting these two distinct conditions should turn out to be (necessarily) the same, then our claims will be (necessarily) equivalent. Moreover, the hypothesis that these sets of circumstances are (necessarily) the same promises to be hugely explanatory: in particular, if the permissible circumstances are the normal ones, then we may achieve some understanding of the source of epistemic normativity. The claim that a circumstance is normal in Millikan's sense is the claim that it involves the proper functioning of a certain sort of device, namely one that is beneficial in some respect when it is operative in such cases. This suggests that those for whom the normal functioning of that device is beneficial should be motivated to ensure that the circumstances in which they find themselves are normal with respect to it. Thus, it is plausible that the normal cases and the permissible ones are (necessarily) the same.

So our hypothesis is that in cases in which some device type functions normally, beliefs constitute knowledge. This leaves two questions to be answered: What is the device (type) at issue? And under what conditions does it function normally? The second of these questions will be addressed in the next section; but the answer to the first is, I think, *the mechanism of belief formation* - that is, we are

⁸ For instance, one might be in fake barn country (Goldman 1976), or in a region riddled with a gas which induces false beliefs (Hawthorne 2007).

⁹ The selection in question might be natural, in which case it will have been some evolutionary advantage which accrued to the organism possessing the device; but equally, it could be artificial. For instance, a hammer might be thought to be behaving normally just in case it is driving nails: for it is because of the fact that it is useful to tool users that hammers drive nails that devices of this type - those possessing a certain characteristic shape - are selected by tool users and reproduced.

concerned with the normal cases of the production of belief.¹⁰ Thus, the claim that knowledge is the norm of belief reduces to the claim that, in normal cases of belief formation, one believes that p only if one knows that p. And this allows us to define knowledge as that which is produced in the normal cases in which a belief is formed. In other words, knowledge is normally formed belief; or, more tersely yet, knowledge is normal belief.

Whereas the notion of epistemic permissibility tacitly involves the notion of knowledge, that of doxastic normalcy, as explicated above, does not. My proposal is that knowledge is belief which is normal *for belief*: this definition is not circular, and can therefore serve as an analysis; it remains only to argue that it is correct.

5. Normal Belief is True and Safe from Error

Knowledge is normal belief. If this definition is to be remotely plausible, the conditions of normal belief formation had better guarantee knowledge. But this does seem to be the case. For the possession of beliefs is no great advantage to a subject in and of itself; rather, beliefs are valuable in this way only if they are useful, fulfilling some role within the overall cognitive economy of the subject possessing them. In other words, having beliefs, and therefore a belief producing mechanism, is beneficial only to those subjects that also have mechanisms of belief consumption. And I propose, more specifically, that this mechanism of belief consumption is desire: for it is when a belief is acted upon in such a way as to satisfy a desire that it benefits the one who holds it.¹¹

This assumes, of course, that belief interacts with desire so as to produce action. Some have thought that this is indeed the essence of belief; according to Stalnaker, for instance:

To believe that P is to be disposed to act in ways that would tend to satisfy one's desires, whatever they are, in a world in which P (together with one's other beliefs) were true. (1984: 15)

But one need not accept this reductive dispositional account of belief in order to concede the current point. In particular, one might take the more modest view that belief is a state which *involves* a disposition to act in ways that tend to satisfy the subject's desires in worlds in which the proposition believed is true. What matters for our purposes is just that it is advantageous to an

¹⁰ This process of belief formation, as I am thinking of it, is very narrowly individuated: it is simply the generation of a mental representation with a certain associated functional role – namely, that of belief. Beliefs themselves, of course, belong to the central cognitive system in Fodor's (1983) sense, and are available to all central processes; but belief formation does not have access to e.g. what other beliefs are in place. We might therefore think of it as supremely informationally encapsulated; more accurately, though, it is not a rational/informational device at all, but a brutally causal process.

¹¹ I assume here that the organism's desires are well suited to its needs. That this will be so in those cases in which action is reproductively beneficial seems to me highly plausible.

agent to have a belief producing mechanism only if it has desires which (sometimes) interact with the beliefs produced in such a way as to satisfy those desires in worlds in which the beliefs in question are true.

Beliefs interact with a wide range of desires to produce action. If a given belief could only interact with one particular desire in a causally efficacious manner, we would have no grounds for regarding that cause of action as a complex comprising component parts – no grounds, that is, for distinguishing the belief from the desire.¹² Beliefs are items which interact with our desires ‘whatever they are’: there is no limit to the range of desires to the satisfaction of which a given belief may be relevant.¹³ This is crucial since it gives us reason to believe that when our beliefs are normally produced they are true. A false belief might interact with a single desire to cause some action which was beneficial to an agent; but the best explanation of *generally* successful action requires, at a minimum, true belief. A subject’s belief producing mechanism will be beneficial in such a way as to be selectively reproduced, then, only when it produces true beliefs.¹⁴

In fact, normally produced beliefs will be not only true, but also safe from error; that is, they will be such that one would not easily be mistaken on the matter they concern. One may see this point by reflecting on the fact that an agent which forms the belief that p may simply be ignorant on the unrelated issue whether q;¹⁵ yet her resultant belief must be capable of combining with her desires, whatever they may be, to yield successful action whether q or not - that is, in worlds in which q, and in worlds in which not q. Since this will hold for a variety of propositions q, the agent’s belief will be true not only in the actual world, but in a range of nearby possible worlds as well, if her belief has been normally formed.

¹² In short, the cause might be what Millikan (2005: chapter 9) has called a ‘pushmi-pullyu’ representation, undifferentiated between belief and desire.

¹³ In this respect it is unlike the representations employed in modular systems.

¹⁴ There might be thought to be exceptions to this rule: for instance, most people believe that they are better than average drivers; perhaps it is advantageous to an agent to have false overestimates of its own abilities and worth. But this seems to me to be a parasitic effect: it is only possible for false beliefs about oneself to have this beneficial effect given that one treats them as true; accordingly, we can assume that the particular processes leading to the formation of such beliefs arose after devices of belief formation were firmly in place, and they can safely be regarded as abnormal.

¹⁵ As special cases we might consider q = the agent in question believes that r, and q = that agent desires that r. In the latter case, as the mechanism of belief formation is blind to the desires the agent has, it will (normally) produce a belief that is able to satisfy this desire, whether the agent possesses it or not.

I have argued that normally produced beliefs will be both true and safe from error. These are the two least controversial necessary conditions on knowledge.¹⁶ In the absence of any reason to deny that normally produced belief is knowledge, I suggest that we have here a single phenomenon.

6. Conclusion

I have argued that knowledge is epistemically permissible belief (section 2), and that epistemically permissible belief is doxastically normal belief (sections 3, 4, and 5); and I have therefore concluded that knowledge is doxastically normal belief. In this way I have arrived at a new, naturalistic analysis of knowledge which, arguments to the contrary notwithstanding (section 1), has every chance of being not only philosophically interesting, but also materially correct.

References

- Ayer, A. J. (1956), *The problem of knowledge* (London: Macmillan) x, 258 p.
- Fodor, Jerry A. (1983), *The modularity of mind: an essay on faculty psychology* (Cambridge, Mass: MIT Press) 145 p.
- Gettier, Edmund L. (1963), 'Is Justified True Belief Knowledge?' *Analysis*, 23 (6), 121-23.
- Goldman, Alvin I. (1976), 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy*, 73 (20), 771-91.
- Hawthorne, John (2007), 'A Priority and Externalism', in Sanford Goldberg (ed.), *Internalism and externalism in semantics and epistemology* (Oxford: Oxford University Press), 201-18.
- Millikan, Ruth Garrett (1984), *Language, thought, and other biological categories: new foundations for realism* (Cambridge, Mass: MIT Press) x, 355 p.
- (2005), *Language: a biological model* (Oxford: Clarendon Press) vii, 228 p.
- Stalnaker, Robert (1984), *Inquiry* (Cambridge, Mass: MIT Press) xii, 187 p.
- Sutton, Jonathan (2007), *Without justification* (Cambridge, Mass: MIT Press) vi, 180 p.
- Williamson, Timothy (2000), *Knowledge and its limits* (Oxford: Oxford University Press) xi, 340 p.
- (2007), *The philosophy of philosophy* (Oxford: Blackwell Publishing) xiv, 332 p.
- Zagzebski, Linda (1994), 'The Inescapability of Gettier Problems', *The Philosophical Quarterly*, 44 (174), 65-73.

¹⁶ It is not essential to my view that knowledge should be subject to a safety constraint; but knowledge is not merely true belief, and safety seems a good candidate for what more is needed at a minimum. If some alternative condition on knowledge turns out to be necessary, obviously I am committed to the claim that normal belief satisfies it: but this cannot be argued until that condition is articulated; safety is therefore considered for convenience.