Proper Functionalism

by Kenneth Boyce and Alvin Plantinga

PENULTIMATE DRAFT

Introduction

Warrant is that, whatever precisely it is, which makes the difference between knowledge and true belief.[1] Proper functionalism with respect to warrant is the view that (1) it is a necessary condition for a belief's being warranted that it arise by way of cognitive proper function and (2) any adequate analysis of warrant centrally involves that notion. There can be proper functionalist theories of other epistemic concepts besides warrant; unless said otherwise, however, when we use the term 'proper functionalism' it will be proper functionalism about warrant to which we are referring. In this introductory essay, we will describe proper functionalism about warrant and point to some of the reasons for it. We will also respond to some objections leveled against this view and say a few words about how it is related to other issues in epistemology.

I. Motivations for Proper Functionalism

Some Preliminary Remarks

The notion of proper function appealed to in proper functionalist analyses of warrant is that found in the biological sciences (as when a biologist says that the average human heart functions properly when it beats 70 times per minute).[2] Of particular import, the concept of proper function presupposes the notion of a *design plan* – something that specifies the way in which a thing is *supposed to work.*[3] We need not initially take the notions of *design plan* and the *way in which a thing is supposed to work* to entail

---

[1] See Plantinga (1993a), p. 3 for the original characterization of warrant along these lines. We are grateful to Trenton Merricks for pointing out that the definition found there contains a misprint. The definition of 'Warrant' found in Merricks (1995), p. 841 and reproduced here reflects Plantinga's original intentions.

[2] See Plantinga (1993b), pp. 5-6.

[3] More precisely, we might think of a design plan as a set of triples specifying circumstance, response, and purpose or function. See Plantinga (1993b), pp. 22-24 for a fuller discussion.

conscious design or purpose.[4] In central and paradigm cases, however, design plans do indeed involve the thing's having been designed by one or more conscious designers who are aiming at an end of some sort.

Roughly (far too roughly to be of any serious use), proper functionalism maintains that a belief's being warranted is a matter of its being produced by a properly functioning cognitive faculty, one whose aim (as specified by the design plan) is the formation of true beliefs. That is, a belief's being warranted is a matter of its being produced by way of truth-aimed cognitive proper function. But why think anything like that? Why hold such a view? An easy way to see some of the motivations is to look at how it circumvents many of the problems that bedevil alternative theories. It is our conviction that any attempt to provide an analysis of warrant that does not acknowledge the central role of cognitive proper function will inevitably flounder.

 The Failure of Internalist Theories

According to internalist theories, what primarily determines whether a belief has warrant (or warrant-like status) are factors or states that are in some sense internal to the believer. (An externalist, by contrast, denies this). More specifically, it is a matter of factors or states to which the believer has some sort of special epistemic access. Perhaps they are factors that a subject can tell are present by reflection alone; or perhaps they are factors that the subject can be certain are present; or perhaps she cannot be culpably mistaken concerning their presence.[5]

These days, in light of the Gettier problem, nearly all parties agree that no strictly internalist conditions are jointly sufficient for warrant. That is because standard Gettier cases show that no matter how well things are going for a believer from an internalist point of view, her beliefs may still be true "just by accident," and therefore not items of knowledge. However, many theories of warrant remain substantially internalist in character, maintaining that whether a belief is warranted for an individual is mostly a matter of that individual's satisfying various internalist conditions, in conjunction with

---

[4] Although, for an argument that there can be no naturalistic account of these notions, see Plantinga (1993b), ch. 11.

[5] Here we offer a rough and ready characterization of internalism, one according to which internalism (or at least any paradigmatic case thereof) is to be identified with access internalism. As far as our current purposes are concerned, this characterization might as well be taken stipulatively.

some (admittedly hard to specify) externalist condition needed to mollify Gettier. Here we will argue that no such substantially internalist theory of warrant can succeed as long as it does not recognize the need for a proper function condition.

In standard Gettier cases, though everything is going well for a believer from an internalist perspective, her beliefs lack warrant because something about her cognitively external environment is amiss. It is tempting, therefore, to think that a substantially internalist analysis of warrant can pacify Gettier merely by adding some (admittedly hard to specify) condition to the effect that things in the believer's cognitively external environment are not awry. This temptation evaporates, however, when we attend to Gettier cases (or at least cases that are *Gettier-like*) in which, though everything is going right from an internalist perspective, the problem lies not in the believer's environment, but in the believer herself.

Consider, for example, a case that Chisholm attributes to Meinong: An aging forest ranger lives in a cottage in the mountains. There is a set of wind chimes hanging from the bough just outside the kitchen window; when these wind chimes sound, the ranger forms the belief that the wind is blowing. As he ages, his hearing (unbeknownst to him) deteriorates; he can no longer hear the chimes. He is also sometimes subject to small auditory hallucinations in which he is appeared to in that wind-chimes way; and occasionally these hallucinations occur when the wind is blowing.

We can stipulate that, on these occasions, the ranger meets whatever plausible internalist conditions for warrant one might like to propose.[6] But clearly, on such occasions, the aging forest ranger's beliefs have little by way of warrant for him. And this is so in spite of the fact that there is nothing amiss in his cognitively external environment. Rather, the problem is due to cognitive malfunction on the part of the ranger himself. Examples of this sort can be multiplied as well as made more specific to various internalist proposals.[7] We take such counterexamples to illustrate that any

---

[6] We take it to be a necessary condition for the internalist conditions being plausible that they do not straightforwardly preclude the possibility of auditory perceptual knowledge.

[7] There are many cases of the sort found in Plantinga (1993a). There are, for a few examples, Chisholm's dutiful epistemic agent who, whenever he is appeared to redly, always believes that nothing is appearing redly to him (pp. 44-45), Pollock's cognizer who by virtue of malfunction has the wrong epistemic norms (pp. 169-170), and the Coherent but Inflexible Climber (p. 82).

substantially internalist analysis of warrant that omits a proper function condition will find itself subject to Gettier-like counterexamples.

The Failure of Reliabilism

It's not just substantially internalist accounts of warrant that flounder on account of omitting a proper function condition. Externalist theories that ignore the role of cognitive proper function also do so to their peril. Here we will support that charge by focusing on reliabilist theories of warrant.

Reliabilism comes in at least two styles. The first sees warrant in terms of origin and province: a belief is warranted for one if it is produced and sustained by a reliable belief-producing mechanism. The second sees warrant as a matter of *probability*; a person is said to know a (true) proposition A if he believes it, and if the right probability relations hold between A and its significant others. Of course, these two different styles of reliabilism can be (and have been) developed with considerable sophistication. And we can't hope to address all the nuances of these developments here. [8] But the thing to see is that on either of these styles of reliabilism, a true belief can have those properties that (according to the theory) confer warrant-like status *by accident*, in the way that precludes them from having warrant.

Consider reliabilism of the second style. For definiteness, let the reliabilist view in question be the view that a belief, B, has warrant just in case it is causally sustained in virtue of S's obtaining, where S is a particular state of affairs, and the objective conditional probability (appropriately relativized to background conditions) of B given that S obtains is sufficiently high.[9] Now consider the case of the Epistemically Serendipitous Brain Lesion. Suppose that Sam suffers from a serious abnormality – a brain lesion, let's say. This lesion wreaks havoc with Sam's noetic structure, causing him to believe a variety of propositions, most of which are wildly false. It also causes him to believe, however, that he is suffering from a brain lesion. Further, Sam has no evidence at all that he is abnormal in this way, thinking of his unusual beliefs as resulting from an engagingly original turn of mind. Since, in this case, Sam's belief that he has a brain

---

[8] See Plantinga (1993a) ch. 9 for much more involved discussion of various reliabilist views.
[9] This simplified reliabilist account is *loosely* based on the much more sophisticated reliabilist account offered in Dretske (1981) and discussed in Plantinga (1993a), pp. 192-197.

lesion is causally sustained in virtue of its being the case that he has a brain lesion and since the objective conditional probability that Sam has a brain lesion given that *Sam's having a brain lesion* obtains is sufficiently high (unity!), it follows from the view we are considering that Sam's belief is warranted. But, obviously, Sam's belief that he has a brain lesion is not warranted.

Consider a paradigmatic reliabilist theory of the first style. According to such a paradigmatic reliabilist theory (minus a few bells and whistles) a belief is warranted for an individual just in case it is produced by a token of a sufficiently reliable type of belief-forming process.[10] Now, adapt The Case of The Epistemically Serendipitous Lesion described above. There is a rare but specific sort of brain lesion (we may suppose) that is always associated with a number of cognitive processes, most of which cause its victim to hold absurdly false beliefs. One of the associated processes, however, causes the victim to believe that he has a brain lesion. Suppose that Sam suffers from this disorder and accordingly believes that he has a brain lesion. Add that he has no evidence at all for this belief (also add, if you like, that he has much evidence against it; but then add also that the malfunction by the lesion makes it impossible for him to take appropriate account of this evidence). Then the relevant type of belief-forming process (while it may be hard to specify in detail) will certainly be highly reliable; but the resulting belief – that he has a brain lesion – will have little by way of warrant for Sam.

What these examples illustrate is that no mere causal or reliability condition is sufficient for warrant.[11] It is possible for such conditions to be met, for the resulting belief to be true, and yet, owing to a lack of cognitive proper function, for the agent to have acquired a true belief merely by accident. So reliabilist theories of warrant, like substantially internalist theories, ignore the need for a proper function condition to their detriment.

The Upshot: Proper Function and Non-Accidentality

We have seen how views of warrant that fail to acknowledge the need for proper function – whether substantially internalist or paradigmatically externalist – face Gettier problems.

---

[10] For a classic and paradigmatic exposition of such a reliabilist view, see Goldman (1979).

[11] See Plantinga (1993a) chapter 9 for a much more detailed argument for this claim.

That is, these views are subject to counterexamples in which a true belief fulfills all of the specified conditions for warrant and, yet, manages to fulfill them "merely by accident". The lesson to be learned, we take it, is that a necessary condition for a belief's being warranted is that it be formed by way of cognitive proper function.

We also take there to be a stronger lesson to be learned. What examples of the above sort illustrate is that, for any merely internalist, merely causal, or merely reliabilist conditions (or any combination thereof), it is possible for a belief to satisfy those conditions "merely by accident" (in the sense at issue in Gettier cases). Furthermore, in many of these sorts of counterexamples, it seems that the only way to *explain why* it is that the relevant beliefs satisfy those conditions merely by accident is to appeal to the fact that they do so in a way that is accidental *with respect to the belief-forming mechanisms specified by the believer's cognitive design plan.* And so it seems that the notion of non-accidentality that pertains to our concept of knowledge is one that cannot be analyzed merely in terms of the meeting of certain internalist, causal or reliability conditions (or combinations thereof). Rather, it is a concept that must be analyzed in terms that make reference to a cognitive design plan. And so the above examples seem to show not only that cognitive proper function is necessary for warrant, but that the notion of cognitive proper function is centrally involved in the concept of non-accidentality that any adequate analysis of warrant must capture.

II. Formulating a Proper Functionalist Analysis of Warrant

We concluded from our consideration of the examples offered in the previous section that cognitive proper function is *necessary* for warrant. Though necessary, however, it is easy to see that it is not sufficient (not even nearly sufficient) for warrant. First, not all aspects of the design of our cognitive faculties need to be aimed at the production of *true* belief; some might be such as to conduce to survival, or relief from suffering, or the possibility of loyalty and so on. What confers warrant is one's cognitive faculties working properly, or working according to the design plan, *insofar as that segment of the design plan is aimed at producing true beliefs*.

Still more is required. For suppose a well-meaning but incompetent angel – one of Hume's infant deities, say – sets out to design a variety of rational persons, persons

capable of thought, belief and knowledge. As it turns out, the design is a real failure; the resulting beings hold beliefs, but most of them are absurdly false. If one of these beings were to happen upon a true belief, it would merely be by accident. The beliefs of these beings, therefore, lack warrant.

What must we add? That the design plan is a *good* one – more exactly, that the design governing the belief in question is a good one; still more exactly, that the objective probability of a belief's being true, given that it is produced by cognitive faculties functioning in accord with the relevant module of the design plan, is high. Even more exactly, the module of the design plan governing its production must be such that it is objectively highly probable that a belief produced by cognitive faculties functioning in accord with that module (in a congenial cognitive environment) will be true or verisimilitudinous.

Even the above conditions are not enough. It is at least broadly logically possible that beings with a good design plan might find themselves in environments radically different from the environments for which they were designed, environments in which the ways of belief formation specified by their design plan are wildly unreliable (though they would be reliable in the environment for which they were designed). If so, the beliefs of these beings, though fulfilling all of the previously mentioned conditions, would, nonetheless, lack warrant. So we need to add that the beliefs in question must be produced in the right sort of cognitive environment,[12] the sort of cognitive environment for which their faculties were designed.

As we have seen, the condition that a belief come about by way of cognitive proper function is not sufficient (not even nearly sufficient) for warrant. But we maintain that something in the neighborhood of the conditions specified above being jointly fulfilled *is* sufficient (or at least nearly sufficient) for warrant. Accordingly, we offer the following as a first approximation to an adequate analysis of warrant:

A belief B has warrant for S if and only if the relevant segments (the segments involved in the production of B) are functioning properly in a cognitive

---

[12] Here, as well as elsewhere, we are ignoring complications that arise from cases in which there is a difference in how a belief is produced and how it is sustained.

environment sufficiently similar to that for which S's faculties are designed; and the modules of the design plan are (1) aimed at truth, and (2) such that there is a high objective probability that a belief formed in accordance with those modules (in that sort of cognitive environment) is true; and the more firmly S believes B the more warrant B has for S.[13]

The above, as was just said, is merely a first approximation. Many amendments and refinements are required (some of which we will mention later on).[14] However, for current purposes, this first approximation will do.

III. Objections to Proper Functionalism

Sadly enough, in spite of the luminous obviousness of proper functionalism, not everyone is on board. Many objections have been proposed.[15] We can't hope to address them all in the space allotted here; we will confine ourselves to addressing what we consider to be two of the most important kinds of objections raised against the view.

Swampman

First, objections according to which there are counterexamples to the claim that cognitive proper function is necessary for warrant, counterexamples in which there are beings who have warranted beliefs but who nevertheless, owing either to poor design or lack of design altogether, fail to exhibit cognitive proper function. A paradigmatic example of such a purported counterexample has been put forward by Ernest Sosa.[16]

Sosa adapts, for his own purposes, Donald Davidson's "Swampman" scenario. In that scenario, by sheer accident, a lightning strike results in the formation of a molecule-for-molecule duplicate of Donald Davidson out of the materials of a nearby swamp. As Sosa envisions the scenario, this "Swampman" exits the swamp having various perceptual experiences, and forming beliefs as a result, in the same way that Davidson

---

[13] This is the first approximation given in Plantinga (1993b), p. 19.

[14] See Plantinga (1993b) ch. 2 for some proposed refinements. See also our discussion below concerning what needs to be added to this analysis in order to avoid Gettier problems.

[15] For a whole host of them, from a number of leading epistemologists, see Kvanvig (1996).

[16] See Sosa (1993), p. 54. For similar purported counterexamples to proper functionalism, see Taylor (1991), pp. 187-188, Feldman (1993), pp. 47-48 and BonJour (2002), p. 225.

himself would have. According to Sosa, "it … seems logically possible for … Swampman to have warranted beliefs not long after creation if not right away." But it also seems that Swampman (having been created just by accident) has no design plan and therefore is not properly functioning.

It's not clear, however, that the Swampman scenario that Sosa envisions *is* (broadly) logically possible (is it possible for a person to be created just by accident in this way?). So it is at least not clear that we are offered a genuine counterexample to proper functionalism here.[17] But suppose that it is possible. It is also at least conceivable that an entity acquires a design plan by accident. Or, in any case, it is not immediately clear that a being has a design plan only if it acquired one as the result of intentional design. It seems that biological systems have design plans (there are ways, for example, that the heart is "supposed to" function) even though, if naturalism and the theory of evolution are both true, these systems were not intentionally designed. So perhaps we may think of Sosa's Swampman case as a case in which a creature acquires a design plan by accident, or at least not by way of intentional design.

But suppose we grant Sosa's case is logically possible and that in it Swampman is not properly functioning. Having granted this, it is no longer clear that, in Sosa's scenario, Swampman has genuinely warranted beliefs. Certainly, if Wilma, an ordinary human being, forms the belief that there is a tree in front of her by way of veridical perception of a tree, in the way that human beings ordinarily form such beliefs, then (provided her cognitive environment is Gettier-free), we are inclined to say that Wilma's belief is warranted. We are inclined to say this because we take it for granted that such a belief-forming process is an epistemically appropriate one for ordinary human beings like Wilma to have. We take it for granted that she forms her belief in the way that human beings are "supposed to" form such beliefs.

It is far from obvious, however, that we are entitled to take it for granted that Swampman's belief-forming processes are epistemically appropriate in the way that Wilma's are. Since, by what has been granted, Swampman's beliefs are not formed by way of cognitive proper function (or by way of cognitive malfunction, for that matter), it

---

[17] For more on this issue, see the exchange between Sosa and Plantinga in Sosa (1993), pp. 55-57 and Plantinga (1993c), pp. 76-78.

is not clear that there are any ways in which it can be said that Swampman is "supposed to" form the beliefs that he does. It seems plausible that if Swampman forms true beliefs, even by way of reliable belief-forming processes, that he does so merely by accident (since he does not form his beliefs appropriately; there is no way of forming beliefs that is appropriate for him). And so, upon reflection, it seems plausible to regard the cases in which Swampman forms true beliefs as Gettier cases (or at least as relevantly analogous to Gettier cases). And for that reason, it is plausible to deny that Swampman's beliefs have warrant.[18]

Gettier Problems

Here's a Gettier-style counterexample to the first approximation of a proper functionalist analysis offered in Section II. It is an amalgamation of examples provided by Richard Feldman, Peter Klein and Robert Shope (the narration is from Plantinga's first person perspective):

> I own a Chevrolet van, drive to Notre Dame on a football Saturday, and unthinkingly park in one of the many places reserved for the football coach. Naturally his minions tow my van away and, as befits such *lèse majesté*, 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 belief that I own such a van is true, but 'just by accident' (more accurately, it is only by accident that I happen to form a true belief); hence it does not constitute knowledge. All of the non-environmental conditions for warrant, furthermore, are met. It also looks as if the environmental condition is met: after all, isn't the cognitive environment here on Earth and in South Bend just the one for which our faculties were designed? What is important about the example is this: it is clear that if the coach's minions had been a bit less zealous and had *not* destroyed my van, the conditions for warrant outlined above would

---

[18] Bergmann (2006), pp. 147-149, in defense of a proper functionalist theory of justification, makes similar points concerning Sosa's Swampman scenario.

have obtained and I would have known that I own a Chevrolet van. In the actual situation, however, the one in which the van is destroyed, my belief is produced by the very same processes functioning the very same way in (apparently) the same cognitive environment. Hence, on the original 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 one isn't.

Upon reflection, it is easy to see that there are many cases of this sort. There is Bertrand Russell's pre-Gettier case, for example, in which one glances up at a clock that has stopped twelve hours ago and forms a true belief about what time it is. There is also, for another example, the case in which one glances across one's lawn and believes truly that it is Paul standing in his driveway, unaware that Paul's identical twin Peter (who one does not even know exists) is currently visiting. In all of these cases, the global environment in which one finds oneself is the one we enjoy right here on Earth, the one for which we were designed by God or evolution. Yet, in each of these cases, one's beliefs are robbed of warrant because of some local environmental anomaly. And so, the original account requires amendment.

Call the sort of global environment alluded to above, the one which we enjoy on Earth and the one for which our cognitive faculties were designed, the "maxi-environment." This environment would include such features as the presence of light and air, the presence of visible objects, of other objects detectable by our kind of cognitive system, of some objects not so detectable, of the regularities of nature, of the existence of other people, and so on. There is also a much less global kind of cognitive environment. For any belief B and (more relevantly) the exercise E of my cognitive powers issuing in B, there is a much more specific and detailed state of affairs we might call its "cognitive mini-environment". We can think of a cognitive mini-environment of a given exercise of cognitive powers E as a *state of affairs* (or proposition) – one that includes all the relevant epistemic circumstances obtaining when that belief is formed (though diminished with respect to whether E produces a true belief). Let 'MBE' denote the mini-environment with respect to a particular exercise of cognitive powers, E, issuing in

the formation of a particular belief, B. To be on the safe side, let MBE be as full as you please, as large a fragment of the actual world as you like.

Now, what's relevant here is that 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 favorable. 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 with respect to a particular exercise of cognitive powers producing the belief that p. So what the original account of warrant requires is an addition to the effect that the relevant mini-environment not be misleading. What must be added to the other conditions of warrant, then, is a *Resolution Condition*:

(RC) A belief B produced by an exercise of cognitive powers has warrant sufficient for knowledge only if MBE (the mini-environment with respect to B and E) is favorable for E.

What does 'appropriateness' or 'favorability' or 'non-misleadingness', for a cognitive mini-environment consist in? Can we say anything more definite?

In the above cases (the Notre Dame van case, the case of Peter's visiting Paul, the case of the stopped clock), the problem is that the relevant exercises of cognitive powers lack sufficient resolution in the cognitive mini-environment. One cannot, while sitting in the Notre Dame stadium, distinguish its being the case that one still owns the van one drove to campus this morning from its being the case that one's van was unexpectedly destroyed. Similarly, one can't, from across the street, distinguish Paul from his identical twin Peter. And, one can't tell, just at a glance, that the clock is stopped. In each case in which one has a belief, B, produced by a given exercise of cognitive powers, E, there are some states of affairs that E is competent to detect in the relevant cognitive mini-environment, MBE, and some which it is not. So consider the conjunction of circumstances C contained in MBE such that C is detectable by E; call this state of affairs DMBE.

With these notions on hand, Plantinga earlier proffered the following, tentative suggestion of what it is for a mini-environment to be favorable:

> (FC) MBE is favorable just if there is no state of affairs S included in MBE but not in DMBE such that the objective probability of B with respect to the conjunction of DMBE and S falls below r,

where r is some real number representing a reasonably high probability. [19] In the Twin case, for example, a state of affairs S such that B is not probable enough with respect to the conjunction of DMBE and S would be Peter's being in the vicinity as well as Paul, and being indistinguishable from him across the street. In the Notre Dame van case, S is the state of affairs of one's van having been destroyed by the coach's over-zealous minions. Etc.

However, this proposal won't quite do. While the original proper functionalist account countenanced as warranted true beliefs that clearly weren't items of knowledge, the current proposal excludes beliefs that do seem to count as items of knowledge. This is illustrated by the following counterexample from Thad Botham:

> Consider a version of the identical twin case where we stipulate that Paul is an only child, thereby removing Peter from the scene. In addition, the moment before you believe the proposition 'There's Paul', his uncle – unbeknownst to you – lies to a friend, telling her that Paul has an identical twin brother who's visiting Paul at that very moment and that he just spoke with each of them on the telephone. Paul's uncle asserts this falsehood while in London, thousands of miles away. In this case, you know the proposition 'There's Paul'. However, the objective probability of your belief is not reasonably high given DMBE together with the state of affairs S – viz., its being the case that Paul's uncle just told a friend that Paul has an identical twin who is presently in the vicinity of Paul.

---

[19] Plantinga (2000), p. 160.

Thus, by (FC) your MBE is unfavorable with respect to E, which together with (RC) entails that you do not have knowledge. But you do know that it is Paul.[20]

Perhaps there is an answer here. In each of the above cases, there are certain states of affairs which are, so to speak, "taken for granted" by E when it issues in B, but which are not included in DMBE. For instance, in the Notre Dame van case, it is taken for granted by E that matters haven't changed with respect to Plantinga's owning the Chevy van – not merely with respect to his owning *a* Chevy van, but with respect to his owning *that* Chevy van. In the case of the stopped clock, it is taken for granted that the clock on the wall is functioning properly. In the twin case, it is taken for granted that there is no one visually indistinguishable from Paul in the vicinity. Etc.

It is difficult to say just what this property of being taken for granted consists in. It is not a matter of one's having any particular *belief* that the relevant state of affairs obtains (in some of these cases, that belief may well be absent). Nor is it obviously a matter of B's being, in some way, *based upon* what is taken for granted. Nevertheless, we seem to have an intuitive grip on the notion. And perhaps that is enough. Provided that it *is* enough, we might try replacing (FC) with the following:

(FC*) MBE is favorable just if for every state of affairs, S, such that S is taken for granted by E in the issuing of B, there is no state of affairs S* such that S* is included in MBE but not in DMBE and such that S* precludes S.

This may work. Induction suggests, however, that still another counterexample will rear its ugly head. Even if one were to appear, however, proper functionalism would be no worse off in this respect (i.e. that of not having a fully adequate solution to the Gettier problem) than any other extant theory of knowledge. There would also be some reason to think it is better off in this respect than many of its rivals. As we argued in the first section, the notion of cognitive proper function must be centrally involved in any adequate analysis of the concept of non-accidentality required to circumvent Gettier

---

[20] See Botham (2003), pp. 435-436. We have slightly altered the quote by omitting various subscripts that Botham uses.

counterexamples. And so it appears that any successful solution to the Gettier problem will be one in which the notion of proper function plays a central role.

IV. Further Directions

So far we have been defending a particular kind of proper functionalist account of warrant. However, there is no reason that the notion of cognitive proper function can't play a significant role in other projects in epistemology. We have already noted, for example, that substantially internalist theories and reliabilist theories of warrant flounder insofar as they ignore the importance of cognitive proper function. But perhaps these theories might be significantly strengthened (while retaining their spirit) by supplementing their proposed analyses with a proper function condition. Investigating such prospects would make for a worthwhile project, one that we commend to others.

Another worthwhile project, and one that is currently a flourishing research program in epistemology, is that of virtue epistemology. According to John Greco, "the central idea of virtue epistemology is that, Gettier problems aside, knowledge is true belief which results from one's cognitive virtues", where "a cognitive virtue … is an ability [or "cognitive faculty"] to arrive at truths in a particular field, and to avoid believing falsehoods in that field, under the relevant conditions."[21] Greco classifies proper functionalism, as well as Ernest Sosa's Goldman-inspired brand of reliabilism, as a paradigm instance of virtue epistemology.[22] Sosa himself concurs with this assessment, maintaining that his own view, some of Alvin Goldman's more recent views, and proper functionalism are all properly seen as "varieties of a single more fundamental option in epistemology, one which puts the explicative emphasis on truth-conducive intellectual virtues or faculties, and is properly termed 'virtue epistemology'".[23]

We agree that proper functionalism is correctly regarded as a version of virtue epistemology. Or rather, as we would like to put it, the various views that fall within the scope of virtue epistemology should really be thought of as varieties of proper functionalism! That is because virtue epistemology of any stripe presupposes a notion of cognitive proper function. Virtue epistemology centrally involves an appeal to cognitive

---

[21] Greco (1993), p. 414.
[22] Ibid., p. 413.
[23] Sosa (1993), p. 64.

virtues, understood as cognitive faculties or cognitive abilities. It is this appeal that is supposed to help evade counterexamples to reliabilism, such as *The Case of the Epistemically Serendipitous Brain Lesion.* Sosa, for instance, responds to the latter by holding that a belief is warranted only if it is produced by a faculty, and a brain lesion isn't a faculty. This seems right. But the notion of a faculty itself involves the notion of proper function. A faculty or power – perception, or memory, or reason in the narrow sense, or digestion, or one's ability to walk, is precisely the sort of thing that can function properly or improperly. Indeed, this is just the difference between the brain lesion and a faculty: the concepts of proper and improper function don't apply to the brain lesion. It isn't functioning either properly or improperly in producing the belief it does: it isn't that sort of thing. So, the relevant notion of a cognitive faculty required by the virtue epistemologist presupposes the notion of cognitive proper function. Insofar as it does, we can see different views within virtue epistemology as varieties of proper functionalism.

So far we have described how proper functionalism is interestingly related to programs in epistemology that are primarily concerned with the concept of *warrant*. It is also possible, however, to develop proper functionalist theories of other significant epistemic properties besides warrant. Michael Bergmann, for example, has offered an interesting and promising proper functionalist theory of epistemic justification. We will close this section with a brief summary of some of the details and merits of Bergmann's theory.

In order to be clear about what the advantages of a proper functionalist theory of epistemic justification might be, it would be helpful to have a clear notion of just what property the term 'epistemic justification' is meant to pick out. Unfortunately, in spite of the extensive literature on the concept of epistemic justification, it is difficult to say, exactly, what epistemic justification is, or even whether there is a single property being referred to in these debates.[24] We can say at least this much; the concept of justification is a transparently normative concept. And perhaps we can understand this concept (where this concept shows up in many of these discussions) as the kind of rational normativity involved in warrant. In particular, what seems relevant here is rational normativity *downstream* from experience (i.e. forming appropriate beliefs, in response to

---

[24] See Alston (2005), ch. 1.

experience, irrespective of the way in which the subject's experience is connected with the world).[25] From the standpoint of proper functionalism, this kind of normativity is to be analyzed in terms of truth-aimed cognitive proper function. And Bergmann proposes just such an analysis.

Before we describe Bergmann's analysis in more detail, however, it will be helpful to give a quick overview of some of the difficulties faced by alternative theories of justification.[26] Not only is the concept of justification a normative one; it is also (at least insofar as it is a concept of the kind of normativity involved in warrant) a concept that involves the notion of truth-aptness.[27] That is, there is some connection between a belief's being produced in accordance with the norms required for justification and its being likely to be true. What norms of belief formation are truth-apt also depends, however, to a large extent, on the kinds of environments in which believers find themselves. It is appropriate, for example, for human beings to believe that ordinary objects continue to exist, even when they are not being perceived. But it is not difficult to imagine a cognitive environment in which this norm fails to be truth-apt. Or, for another kind of example, Goodman's New Riddle of Induction shows that our ways of making inductive projections are, at best, merely contingently reliable; in different kinds of cognitive environments, it would have been the case that projections involving predicates like 'grue' and 'bleen' would have been truth-apt whereas projections involving 'green' and 'blue' would not have been. And so which epistemic norms are the truth-apt ones is a contingent matter, one that depends on the nature of the relevant cognitive environment. But since the concept of justification at issue is (at least something in the neighborhood of) the concept of rational normativity *downstream* from experience, it also seems that whether or not a belief is justified is independent of environmental factors upstream.

---

[25] Note that this is *not* the deontological concept of justification discussed in Plantinga (1993a), ch 1. It is more akin to what is there called "broad justification" (see p. 11). It should also be noted that this way of picking out the relevant concept of justification differs from the one that Bergmann himself provides (for Bergmann's own way of picking it out, see Bergmann 2006, pp. 1-9). Nevertheless, we think that this characterization of justification gets at the same sort of thing that Bergmann is interested in.

[26] What follows is, for the most part, a summary of many of the points found in Bergmann 2006, ch. 5.

[27] On the seemingly puzzling connection between justification and truth, see Cohen (1984).

Understandably then, it is difficult to find a theory of justification that successfully accommodates all of the above features (normativity, truth-aptness, environmental independence). The difficulty of accommodating all of these features is made especially vivid by "The New Evil Demon Problem", originally put forward by Keith Lehrer and Stewart Cohen, as a problem for reliabilist theories of justification.[28] Consider a population of beings, just like ourselves, who form their beliefs downstream from experience in just the ways that we do, but who (unlike us) are victims of a Cartesian demon who renders their belief-forming processes unreliable. From many reliabilist theories of justification, it follows that these beings have far less by way of justified belief than we ourselves do (since most of their beliefs are not formed in a truth-reliable manner).

Intuitively, however, it seems that their beliefs are justified to the same extent that our beliefs are. Their living in a demon world is an epistemic misfortune, to be sure, one that robs most of their beliefs of warrant. Nevertheless, surely the beliefs they have are still the appropriate ones for beings like them to hold in their circumstances; surely their beliefs are at least justified! One might think that an easy fix here is simply to jettison reliabilist theories of justification and maintain that which norms of belief formation are necessary and sufficient for justification is a non-contingent matter that does not depend on the epistemic congeniality of the environment. However, adopting such a theory seems to sever justification from truth-aptness, since, as we have already pointed out, which epistemic norms are truth-apt varies across different kinds of cognitive environments.

Bergmann's analysis of the concept of justification, by contrast, readily accommodates the notions of normativity, truth-aptness, and environmental independence involved in that concept. (Roughly) according to the analysis of justification that Bergmann provides, a belief is justified just in case it fulfills all of the proper functionalist conditions mentioned in the analysis proposed in Section II, absent the environmental condition.[29] Bergmann argues that this analysis of justification accommodates the intuition that beings like ourselves living in a demon world have

---

[28] Lehrer and Cohen (1983), pp. 192-193.
[29] Bergmann also adds a no-defeater condition (see Bergmann 2006, p. 133).

justified beliefs to the same extent that we do. For, provided that these beings have a cognitive design plan comparable to ours and they are properly functioning, many of their beliefs are justified, even though their ways of forming beliefs are, for the most part, unreliable. This analysis also, as Bergmann points out, accommodates the intuition that justification is connected with truth-aptness. For, insofar as the beings living in a demon world fulfill Bergmann's conditions for justification, the manner in which they form their beliefs would be truth-apt if they were placed in the environment for which their cognitive faculties were designed.[30]

Insofar as these advantages adhere to Bergmann's proper functionalist theory of justification, they also adhere to proper functionalism about warrant (since the concept of justification, as we have characterized it, is a concept closely related to the rational normativity required for warrant). And so the advantages that Bergmann claims for his own theory also serve to illustrate another way in which proper functionalism about warrant accounts for features of that concept that alternative theories leave out. Thus, the advantages of Bergmann's proposal offer additional support for our claim that any adequate analysis of warrant will centrally involve the concept of truth-aimed cognitive proper function.[31]

Works Cited

Alston, William P. (2005) *Beyond 'Justification': Dimensions of Epistemic Evaluation* (Cornell University Press).

Bergmann, Michael (2006) *Justification Without Awareness: A Defense of Epistemic Externalism* (Oxford University Press).

BonJour, Laurence (2002) "Internalism and Externalism" in *The Oxford Handbook of Epistemology* edited by Paul K. Moser (Oxford University Press), pp. 234-263.

Botham, Thad M. (2003) "Plantinga and Favorable Mini-Environments" *Synthese*, 135, pp. 431-441.

Cohen, Stewart (1984) "Justification and Truth" *Philosophical Studies*, 46, pp. 279-295.

---

[30] Bergmann (2006), pp. 141-143.

Dretske, Fred I. (1981) *Knowledge and the Flow of Information* (The MIT Press).

Feldman, Richard (1993) "Proper Functionalism" *Nous*, 27, pp. 34-50.

Goldman, Alvin J. (1979) "What is Justified Belief" *Justification and Knowledge* edited by George S. Pappas (Dordrecht, Holland: D. Reidel Publishing Company), pp. 1-23.

Greco, John (1993) "Virtues and Vices of Virtue Epistemology" *Canadian Journal of Philosophy*, 23, pp. 413-432.

Kvanvig, Jonathan (1996) *Warrant in Contemporary Epistemology* (New York: Rowman and Littlefield).

Lehrer, Keith, and Stewart Cohen (1983) "Justification, Truth, and Coherence" *Synthese*, 55, pp. 191-207.

Merricks, Trenton (1995) "Warrant Entails Truth" *Philosophy and Phenomenological Research*, 55, pp. 841-855.

Plantinga, Alvin (1993a) *Warrant: The Current Debate* (Oxford University Press).

– (1993b) *Warrant and Proper Function* (Oxford University Press).

– (1993c) "Why We Need Proper Function" *Nous*, 27, pp. 66-82.

– (2000) *Warranted Christian Belief* (Oxford University Press).

Sosa, Ernest (1993) "Proper Functionalism and Virtue Epistemology" *Nous*, 27, pp. 51-65.

Taylor, James E. (1991) "Plantinga's Proper Functioning Analysis of Epistemic Warrant" *Philosophical Studies*, 64, pp. 185-202.