

KNOWLEDGE-FIRST FUNCTIONALISM

Mona Simion

COGITO Epistemology Group, University of Glasgow

Abstract

This paper has two aims. The first is critical: I identify a set of normative desiderata for accounts of justified belief and I argue that prominent knowledge first views have difficulties meeting them. Second, I argue that my preferred account, knowledge first functionalism, is preferable to its extant competitors on normative grounds. This account takes epistemically justified belief to be belief generated by properly functioning cognitive processes that have generating knowledge as their epistemic function.

1. Introduction

Several philosophers take it that, in an interesting way to be further specified, when it comes to epistemological affairs, knowledge comes first. It is the goal of inquiry. It is a mental state in its own right. It specifies the condition for epistemically proper assertion and practical reasoning. Other epistemic and/or doxastic states are to be analysed in terms of knowledge. It is the norm of belief.

This paper assumes that the knowledge first picture—in the broad shape outlined above—is correct. Against this backdrop, in what follows, I will try to identify a normatively satisfactory account of epistemic justification that does justice to these claims.

Justification is widely taken to be normative. The following is an attractive way of capturing this thought:

The Deontic Thesis (DT). One's φ -ing is prima facie practically, morally, epistemically, etc. justified if and only if one prima facie practically, morally, epistemically, etc. permissibly φ s.ⁱ

If DT captures the way in which justification is normative, then, plausibly enough, the following captures the sense in which the epistemic justification of belief is normative:

The Deontic Thesis for Belief (DTB). One's belief that p is epistemically justified if and only if one epistemically permissibly believes that p .ⁱⁱ

That said, given a substantive account of permissible belief, we can, of course, use DTB to derive a substantive account of justified belief and vice versa.

Crucially, also, given DTB, desiderata for satisfactory accounts thereof will be reciprocally inherited: our account of justification had better be a normatively satisfactory account. Similarly, our account of the norm of belief had better accommodate our intuitions about justification.

This paper is concerned with the former. That is, in what follows, I will examine norm-based accounts of justified belief in the knowledge-first camp and compare their normative credentials to those of the account I prefer. Importantly, this will not exhaust the field: several knowledge-first epistemologists have proposed virtue-based rather than norm-based views (e.g. (Kelp 2016, 2018), (Miracchi 2015)) that follow a similar normative pattern to the view defended in this paper. As such, they will likely enjoy similar normative benefits to the view proposed here, insofar as their proponents want to incorporate DTB in their accounts.ⁱⁱⁱ The focus of his paper, then, is restricted to proposals that incorporate DTB, either explicitly or implicitly, in their knowledge-first accounts of justified belief.

Here is the game plan: #2 puts forth three minimal normative desiderata that any account of justification that incorporates DTB needs to satisfy. In Section #3 and #4 I argue that several popular norm-based knowledge first views fall short of meeting at least one of the normative desiderata identified. In #5, I defend my view, Knowledge First Functionalism, and argue that it does better than its rivals. According to the view defended here, a belief is epistemically justified if and only if it is formed via a properly functioning cognitive process that has the etiological function of generating knowledge.

2. Normative Desiderata

Theories about stuff X, be they scientific, philosophical or about how to make the best cheesecake, can be better or worse: they can, that is, display more or less theoretical virtues.^{iv} Here are a few such virtues: evidential adequacy, causal adequacy, explanatory depth, internal consistency, internal coherence, prior plausibility, generalizability, beauty, simplicity, unification, durability, fruitfulness, applicability... Now, doing well on all of these counts is notably hard, in particular since some theoretical virtues seem to come in conflict at times; importantly, though, it is not the case that anything goes either: evidential accuracy, prior plausibility and generalizability are arguably (some of) the most stringent requirements for acceptable theories of X:

PRIOR PLAUSIBILITY: Theory T is externally coherent (it fits widely accepted claims in T's domain of inquiry).

GENERALISABILITY: Theory T can be successfully applied to settings other than that in which it was originally tested.

EVIDENTIAL ADEQUACY: Theory T fits the evidence well.

Now, recall that, given DTB, desiderata for satisfactory accounts of justification, respectively the norm of belief, will be reciprocally inherited; here are, then, three minimal normative desiderata that get inherited via DTB right to left: First, one's account of justification had better enjoy minimal prior plausibility; i.e., sit nicely within general normativity theory:

NORMATIVE PRIOR PLAUSIBILITY: Theory T is externally coherent (it fits widely accepted claims in the theory of normativity).^v

Alternatively, if the proposed account of justification takes epistemic normativity to work in special ways in one respect or another, a good case should be made for this claim.

Second, note that epistemologists do not enjoy exclusivity on justification. Justification can be e.g. practical, moral or social alike; also, both states and actions can be justified. Ideally, then, we want our account of justification to be generalizable:

NORMATIVE GENERALISABILITY: Theory T can be successfully applied to normative settings other than that in which it was originally tested.

That is, we want T to be generalizable to other normative domains – prudential, moral etc.—and to other targets of application—actions, other types of states etc. Again, alternatively, if this fails, we should be given a good reason to believe epistemic justification is somehow special.^{vi}

Last, the account proposed should exhibit minimal intuitive adequacy in normative respects. In a nutshell, this desideratum asks for the proposed account to be able to account for intuitively distinct normative statuses:

NORMATIVE EVIDENTIAL ADEQUACY: Theory T fits normative intuitive data well.

I take these three desiderata to be fairly minimal, and thus fairly uncontroversially constraining any account of justification that accepts DTB. In the light of this framework, in what follows I will (1) check the normative credentials of several popular, DTB-friendly knowledge first accounts of justification in the literature, against the desiderata identified (2) offer a novel, knowledge first functionalist account of justified belief and (3) argue that the account put forth does better than the competition in meeting the minimal normative desiderata identified.

3. Simple Knowledge First

Several philosophers believe that knowledge is the norm of belief:

KNB: A belief is epistemically permissible if and only if knowledgeable.

DTB in conjunction with KNB gives:

JB=K: One's belief that *p* is epistemically justified if and only if knowledgeable.

Several prominent philosophers defend JB=K, e.g. Brian Ball (2013) Clayton Littlejohn (2013), Adrian Haddock (2010), Jonathan Sutton (2005, 2007) and, most notably, Tim Williamson (2007, 2009, 2011).

JB=K faces one important objection: on JB=K, the concept of justification seems to fail to play some of the most important roles we have historically wanted it to play. In particular, it looks as though the view is too strong: it fails to grant justification to both deceived^{vii} and Gettierized victims^{viii}—after all, they don't know—and, by factivity of knowledge, it fails to allow for justified false beliefs.

Traditionally, JB=K theorists appeal to normative pluralism to escape objections along these lines. The move is one whereby they distinguish between blameworthiness and norm violation; the two are distinct normative concepts, and can be independently instantiated. According to the JB=K theorists, we are not very good at distinguishing between intuitions pertaining to these two different normative concepts; the ‘warm and fuzzy feeling’,^{ix} as it were, that we get when we consider the unfortunate epistemic situation of, for instance, the deceived victim (or the Gettierized victim, or the epistemically conscientious false believer), we confuse for approval sourced in compliance with the norm of belief, when, in fact, its source is mere blameless norm violation.^x

Now, of course, since blamelessness is, itself, a normative notion, there are conditions—i.e. normative constraints—that one needs to meet in order to qualify as a blameless norm violator. Here are but a few identified in the relevant literature: lack of control over one’s actions, ignorance ((Kelp and Simion 2017), (Littlejohn Forthcoming), (Zimmerman 1997)), being generally disposed to conform with the norm, acting as one who is so disposed (Williamson Forthcoming). The question, then, becomes: what is the status of these further constraints in the normative landscape?

According to Tim Williamson, these normative constraints are, crucially, derivative of the primary norm at stake—in our case, the knowledge norm of belief (Forthcoming, 7-8). There is an important difference to be made between the normative status conferred by the primary norm governing φ -ing—mapping on to justified φ -ing—and the normative status conferred by mere compliance with the derivative norms—mapping on to blameless norm violating φ -ing. In this vein, according to Williamson, the deceived, the Gettierized victim and the conscientious false believer, while doing what someone disposed to comply with the norm of belief would do—and thereby complying with the relevant blamelessness conferring derivative norm—they fail to comply with the primary, justification conferring norm.

Insofar as it is overwhelmingly plausible that, in general, norms can be blamelessly broken, the JB=K line drops out nicely from general normativity theory, and thereby meets **NORMATIVE PRIOR PLAUSIBILITY**.^{xi}

There are two different normative-theoretic problems with this move, though. The first pertains to the account’s potential to capture all the needed normative distinctions these cases point to (**NORMATIVE EVIDENTIAL ADEQUACY**). In a nutshell, here is the worry: there intuitively seems to be an epistemically interesting normative distinction between massively deceived and Gettierized victims, or everyday conscientious false believers on one hand, and merely well-meaning believers, blamelessly employing epistemically dubious methods and processes in forming beliefs on the other, which are not easily explained by any of the extant accounts of blameless norm violation in the literature. Consider, for instance, Ben, Alvin Goldman’s (1988) benighted cognizer: this fellow forms beliefs about the outcome of the battle based on astrological readings, due to the fact that this is how he has been culturally determined to proceed, by stipulation, through no fault of his own. Goldman’s benighted guy is in breach of any plausible norm of belief, however blamelessly so, both intuitively and by any account of blamelessness in the literature. There seems, however, to be an important difference between Ben and, say, the Gettierized victim, and the difference seems to be a distinctively normative one: the Gettierized is just better off, epistemically, than Ben. To see this, note that if we were to get the opportunity, it would be fine for us to tell Ben to change his epistemic ways, in a way in which we would definitely not want to

do in the case of Gettierized victims. To the contrary, we want to encourage the Gettierized victim to persist in her epistemic practices.^{xiii} The JB=K blamelessness strategy, however, stays silent on this intuitive distinctively normative difference. The view is too coarse grained to be able to capture it: according to JB=K champions, Ben and the Gettierized fellow are both in mere blameless breach of the norm of belief. Furthermore, on all accounts of blamelessness in the literature, they come out as being in the same boat, i.e. blameless in the same way: they are both unaware of being in breach of the knowledge norm, being so is beyond their control, they are both the kind of people who do what a good believer would do in their circumstances etc. As such, for now, it is hard to see how the JB=K picture can explain the intuitive normative difference between the two.

Second, and most importantly, the failure of the JB=K view to allow for justified false belief is problematic for failing to meet **NORMATIVE GENERALISABILITY**. To see this, note that, it looks as though, in other normative domains, justified action is distinct from successful action. Giving money to charity is a morally justified action even though, on a particular occasion, due to a strike of bad luck, it fails to reach its intended target, i.e. achieve success. Morally justified action is distinct from morally successful action. Similarly, I am justified to take Euston road to King's Cross, given that I've done so successfully hundreds of times in the past, even if, on this particular occasion, a jokester daemon moves the station with the result that I fail to get there (no success). Practically justified action comes apart from practically successful action.^{xiiii} It is not clear, then, why it should be the case that in the case of epistemic justification, things are different; that is, why justified and successful (i.e. knowledgeable) belief should amount to one and the same thing.

However, on the JB=K knowledge first picture, successful belief and justified belief coincide: they both amount to knowledgeable belief. This turns the concept of epistemic justification in an odd ball in the normative landscape; as such, the JB=K theorist owes us an explanation as to why this should be so.

In sum: the simple variety of knowledge first justification picture we have been looking at—JB=K—scores well on **NORMATIVE PRIOR PLAUSIBILITY**: in general, norms can be broken blamelessly. However, the view does less well in accounting for all the intuitive normative distinctions needed (i.e., **NORMATIVE EVIDENTIAL ADEQUACY**), and fails to allow epistemic justification to fit in the same normative boat with justification of other sorts – thus scoring badly on **NORMATIVE GENERALISABILITY**. What I will do next is look at how more complex incarnations of the knowledge first view fare with regard to these desiderata.

4. Complex Knowledge First

In response to worries along the previous lines, several more complex knowledge first views have been put forth in the literature. The recipe is, again, normative pluralism. This time around, though, contra JB=K, the permissibility at stake in DTB is taken to map on to milder, non-factive norms.

To see how this goes, note that DTB makes no mention of what norm the permissibility at stake—i.e. the one mapping on to justification—pertains to. Of course, if there is one and only one norm of belief, the answer is easy. On a normative pluralist picture, however, what we need is a more restricted version of DTB, that makes it clear to what norm the permissibility

at stake relates to. For instance, the Williamson view can be seen as defending the following restricted variety of DTB:

DTB*: One's belief that p is epistemically justified if and only if one's belief that p is permissible by the *primary* epistemic norm of belief.

We have a normative pluralist picture with several norms that derive from a primary one, but, crucially, the latter is the one which, when met, confers the status of justified to the relevant belief. All the derivative norms only map on to blameless belief.

We have seen that there are several normative issues with this view, however, sourced in its being too strong. That, though, need not mean that one needs to abandon either KNB or DTB. To the contrary, one can just go for normative pluralism in conjunction with a different, milder restriction on DTB. On this type of view, while knowledge is *a* norm of belief—regulating what a successful belief is—justification will not map on to permissibility by the knowledge norm, but by a weaker (derivative) norm governing belief.

Roughly, according to the champions of this family of views, what justification maps on to is what is happening internally in cases of knowledge. According to people like Alexander Bird (2004), Jonathan Ichikawa (2014) and Steven Reynolds (2013), justification is, in a sense to be specified further, would-be-knowledge; that is, it maps on to some internal features of the believer, which, in conjunction with friendlier external conditions, constitute knowledge.

When it comes to fleshing out what the relevant internal state may be, the accounts are quite different: for Bird, what matters is that the believer at stake have the same mental states at this world as a knower does at a different possible world. Roughly, then, the view takes it that justification “is a certain kind of approximation to knowledge, [...] where the failure to know (if any) is explained by factors external to the subject's mental states.” (Bird 2007, 86) According to Jonathan Ichikawa, “a subject's belief is justified just in case her intrinsic state is consistent with her having knowledge” (2014, 189). Reynolds goes one step further in the direction of internalism, imposing a (non-stringent) accessibilist condition on justification: on his account, justification is the appearance of knowledge, where appearance is taken to be a fairly non-sophisticated second-order state, in order to accommodate non-sophisticated cognizers (2013, 369).

Let's start with **NORMATIVE GENERALISABILITY**. Recall that we have seen that, when it comes to justification in general, successful φ -ing comes apart from justified φ -ing. At this point, it should be fairly easy to see that, at least at first glance, internalist-friendly knowledge first views vindicate this thought, and thus score better than JB=K on this count. According to its champions, knowledge is the goal of belief—as such, successful belief will be knowledgeable belief—while the permissibility at stake in DTB maps on to a weaker state: would-be-knowledge.

When it comes to **NORMATIVE PRIOR PLAUSIBILITY**, some of the views in this family have better answers than others. Here is how: the question that naturally arises is: why should we care about would-be-knowledge? Why should we think it has the normative significance ascribed to it by these views? This is a central concern for the internalist-friendly knowledge firster to address. The champions of the view are themselves fairly concerned with answering this worry. In a nutshell, here are the three proposals they put forth:

Reynolds takes the normativity of assertion to aid explain the normative status of possible knowledge. According to him, knowledge is the norm of assertion, and, one needs to keep ones beliefs in constant check for assertability. Awareness of knowledge, then, is instrumental to permissible assertion, which explains its distinctive normative status (Reynolds 2013, 367). There are two main problems with this move: First, assertion is a social phenomenon. Surely, though, we already had justified and unjustified beliefs before living in a society. As such, it seems fairly implausible that the normative significance of justified belief derives from the normative significance of permissible assertion.

Second, crucially, note that for Reynolds' argument to work, what is needed is a fairly sophisticated second order state of awareness, rather than the very basic one Reynolds advertises when putting forth the view. That is, quite a bit of reflective work seems needed for selecting assertable beliefs; it seems implausible that someone who does not have the relevant concepts—i.e., of assertion, permissibility, belief etc. —is able to assess a belief for assertability. If that is so, though, non-sophisticated awareness of knowledge will be of no use. The alternative, of course, would be to make the view stronger than that, but then it would lose much of its externalist appeal, in virtue of lack of friendliness towards non-sophisticated cognizers.

Ichikawa identifies a different normative source: he takes matters internal to the believer to map on to blamelessness. According to him, then, justified belief is a kind of blameless belief, which explains its normative significance (Ichikawa 2014, 193). Accordingly, then, Ichikawa's view shares with JB=K its **NORMATIVE PRIOR PLAUSIBILITY**: blamelessness is a perfectly respectable normative phenomenon. On this view, then what we get is two normative dimensions: good belief—i.e., knowledgeable belief, mapping on to permissibility by the knowledge norm—and justified belief—mapping on to permissibility by a blamelessness conferring norm that is restricted to regulating the internal features of the believer.

Bird disagrees; according to him, there is a clear normative difference between blamelessness and justification; justified belief is not merely blameless, but praiseworthy (2007, 108). Accordingly, then, on Bird's view, we get three important normative distinctions: successful belief (knowledge), praiseworthy belief (justified belief) and blameless belief. In aiming at knowledge, one can fail to reach one's aim while doing nothing wrong (blamelessly); that is, for instance, the case of someone who is brainwashed into believing a falsehood. In contrast, one can fail to believe knowledgeably while, at the same time, doing something right, i.e. praiseworthy. This latter normative dimension, according to Bird, maps on to justified belief, and regards one's proper 'ordering' of one's mental life (2007, 108).

Importantly, note that, while **NORMATIVE PRIOR PLAUSIBILITY** dismisses Reynolds's answer, it does nothing to adjudicate between Ichikawa's and Bird's. After all, both blamelessness and praiseworthiness are perfectly respectable normative concepts, so whichever of the two justification maps on to will do just fine.

Alas, though, the move will get the complex knowledge firster in trouble elsewhere. And here is why: first, in virtue of borrowing the normative framework from the JB=K view, Ichikawa's account inherits both its ups and its trouble. Recall Ben, the believer in astrology: no account of blamelessness in the literature is able to explain why there seems to be an important, epistemically normative difference between him and, say, the Gettierized victim. Of course, Ichikawa's account has a clear answer to why the latter is justified in her beliefs, while Ben is not: Ben fails to have would-be-knowledge,

in the relevant sense. Note, though, that we are after is the normative significance of would-be-knowledge. If all there is to it is blameless norm violation, it follows that Ben's blameless beliefs are normatively on a par with those of Gettierized agents. This, of course, does not get **NORMATIVE EVIDENTIAL ADEQUACY** right, in virtue of missing important normative distinctions.

Bird's proposal does better on this front, in virtue of being more fine-grained. According to Bird, justified believers are better off than mere blameless believers: they are not only not deserving of a negative normative evaluation, but they are even worthy of praise.

Unfortunately, Bird's account of praiseworthiness is too externalist for **NORMATIVE PRIOR PLAUSIBILITY**. To see this, note, first, that one can either think that praiseworthiness is a purely internal matter, or not. The first option looks normatively more plausible: it makes sense to have some positive evaluation for agents who do internally perfectly fine in one respect or another, but fail to reach their goal due to external factors. Bird himself motivates his view along such lines:

[Some] some failures can be laid at the door of the believer, because the source of failure is one or more of the believer's mental states, and some failures can be ascribed to mischance, in that the failure is due to some mentally extraneous factor. The role of the concept of justification is to mark the difference between these different sources of failure (Bird 2007, 96)

According to Bird, then, the crucial role played by the concept of justification—i.e. epistemic praiseworthiness—is to mark the difference between failure that is due to our environment and failure that is due to us. Now, here is the problem: Bird accepts content externalism. As such, what mental states one is in will not merely supervene on internal features of the subject. Bird's justification preserves externalist flavor in that it is not entirely dependent on matters internal to the believer. But then, it is not clear how his account of justification—and, in turn, epistemic praiseworthiness—fits the motivations put forth in its support. After all, given that my mental states are not entirely an internal affair, they are still dependent on environmental luck. Failure to be in the right mental state is not, as Bird puts it, something that can be 'laid at the door of the believer' exclusively, any more than knowledge is. They both depend on cooperation of the environment. In this, the view remains silent on the normative significance of would-be-knowledge after all.

5. Knowledge First Functionalism

The worry I closed the last section with was one of ad hoc-ness: if we are to draw the line between justified and unjustified belief somewhere in the environment, rather than along the agent's skin, we need a good (normative) reason to draw it there. The view I will be proposing in what follows purports to do just that.

Before doing so, however, I will try to figure out what might have gone wrong with the accounts we have been looking at on a more general normative level. To this effect, it will be useful to have a look at a widely employed distinction in normativity theory:^{xiv} the distinction between evaluative and prescriptive norms. Prescriptive norms are primarily about what one ought to do.^{xv} Paradigmatic examples of prescriptive norms include moral

norms such as ‘Don’t steal’ or ‘Don’t lie’ but also traffic norms such as ‘Drive 50km/h within city bounds’ and rules of games such as ‘only move the bishop diagonally’ in chess.

In contrast, evaluative norms are ought-to-bes rather than ought-todos: they regulate what it takes for a token of a particular type to be good or bad with regard to its type. Take, for instance, the norm that a good hospital is a clean hospital, that a good knife is sharp or that good driving is safe driving. Evaluative norms use ‘good’ in Geach’s (1956) attributive sense, where “good” functions as a predicate modifier, rather than as a predicate in its own right. When the evaluative norm states that good knives are sharp, it merely states that knives *qua knives* are good only if they are sharp. It does not entail that good knives are good *simpliciter*, or good for some purpose or another.

Evaluative and prescriptive norms can come apart. It is entirely possible for an evaluative norm to be violated without a corresponding prescriptive norm being violated. Consider for instance the evaluative norm: ‘hospitals ought to be clean’, and one of its prescriptive counterparts: ‘hospital management ought to hire cleaning personnel’. The management team in question may have complied with the prescriptive norm, while the evaluative norm may still be violated – if the personnel in question was negligent, or ill equipped etc. And, conversely, it is possible to violate a prescriptive norm without violating any evaluative norm. Suppose, for instance, no cleaning personnel is hired, but the doctors and nurses decide to work extra hours to clean the hospital – and successfully do so.

While evaluative norms thus differ from prescriptive norms, the two may still be related. In particular, prescriptive norms often enough derive from evaluative norms. They serve to ensure that evaluative norms are likely enough complied with. For instance, prescriptive norms of driving, such as the norm ‘Drive no more than 50km/h within city bounds’, serve to ensure that the evaluative norm of driving—according to which good driving is safe driving—is likely enough complied with. In this way, evaluative norms often come first and prescriptive norms are in their service.

With the distinction between prescriptive and evaluative norms in play, let’s return to the case of belief. In particular, let’s ask whether the accounts of justification/the norm of belief we have been looking at should be considered as mapping on to evaluative or prescriptive norms. One answer to this question is suggested in the following passage from Williamson:

“If justification is the fundamental epistemic norm of belief, and a belief ought to constitute knowledge, then justification should be understood in terms of knowledge too” (2014, 5).

It would seem that what Williamson has in mind is an evaluative norm, stating that what it takes for a belief to be a good belief is for it to be knowledge. After all, what Williamson says here is that a belief ought to be knowledge and we saw that ought-to-bes are the stuff that evaluative norms are made of.^{xvi} How about would-be-knowledge? Is this condition plausibly an evaluative or a prescriptive norm? I want to suggest the former. After all, again, would-be-knowledge accounts purport to tell us what a belief ought-to-be in order to be justified: it ought to be the internal duplicate of a knower’s belief, or it needs to bear the appearance of knowledge etc.

Now, we have seen that, generally speaking, evaluative and prescriptive norms can come apart in the sense that it is at least in principle possible to violate a certain evaluative norm without violating any prescriptive norm and vice versa. Contra these accounts, the key hypothesis that I want to

defend next is the following:

Normative Hypothesis: The norm at stake in DTB is a prescriptive norm.

The idea is to allow that, in the case of belief too, the evaluative and prescriptive norms can come apart. In particular, it is possible to violate the evaluative norm of belief (i.e., in our case, knowledge) whilst complying with the prescriptive norm. This will allow us to adopt a knowledge evaluative norm of belief, whilst leaving room for a prescriptive norm that is weaker than knowledge.

The view I want to propose in what follows is one that follows this normative pattern. I dub the view *Knowledge First Functionalism*. The account is functionalist in that it follows (Burge 2010), (Graham 2012) and (Millikan 1984) in taking the epistemic normativity of belief to drop out of the epistemic function of our cognitive processes. It is knowledge-first epistemological in that, unlike traditional, truth-first functionalism, it unpacks the function at stake in terms of knowledge. Here is a more precise formulation of the view:

Knowledge First Functionalism (KFF): A belief is justified if and only if it is generated by a properly functioning cognitive process that has the etiological function of generating knowledge.

On the etiological theory of functions,^{xvii} functions turn on histories that explain why the item exists or operates the way it does. Take my heart; plausibly, tokens of the type pumped blood in my ancestors. This was beneficial for my ancestors' survival, which explains why tokens of the type continue to exist. As a result, my heart acquired the etiological function (henceforth also e-function)^{xviii} of pumping blood.

While etiology does require some history of beneficial effects, it does not require an awful lot of it;^{xix} what it all amounts to, eventually, is explaining the existence/continuous existence of a trait through a through a longer or shorter history of positive feedback:

Functions arise from consequence etiologies, etiologies that explain why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output (Graham 2014, 35).

Functions can be of different sorts: there are biological functions, aesthetic functions, social functions, etc. In contrast to the Graham/Millikan view, my account takes functions to be typed by the corresponding benefit.^{xx} As such, if a trait produces a benefit of type B in a system, the function thereby acquired will be a function of type B. The heart's function to pump blood is a biological function in virtue of the fact that the produced benefit is also biological – i.e., survival. The function of art is an aesthetic function in virtue of the fact that the produced benefit is an aesthetic benefit. Now, of course, aesthetic benefit might, and often will, also result in biological benefit. This, however, in no way renders the function at stake a biological function. What is important to keep in mind is that the benefit that is *essential* to aesthetic function acquisition is the aesthetic one. The fact that biological benefit is also associated with the latter is a mere contingent matter of fact. Here is, then, the full etiological account to be employed by this paper:

E-Function: A token of type T has the e-function of type B of producing effect E in system S iff (1) tokens of T produced E in the past, (2) producing E resulted in benefit of type B in S/S's ancestors and (3) producing E's having B-benefitted S's ancestors contributes to the explanation of why T exists in S.

Note that etiological functions are successes. They explain the continuous existence of the trait that bears them because that is so. The etiological economic function of knife producing economic systems is not just that of producing knives; it's to produce good, sharp knives. To see this, note that the positive feedback loop that is presupposed by etiological function acquisition—the trait produces the effect, the effect benefits the system and thereby contributes to the explanation of the continuous existence of the trait—presupposes a history of success. What contributes to the explanation of the continuous existence of knife producing economic systems is their producing good, sharp knives. Blunt ones would plausibly not have done the trick.

Just as the economic function of knife producing systems is to produce good, sharp knives, the epistemic function of our belief forming systems is to produce good beliefs. Mere belief, then, is a failure on the part of our cognitive system to fulfil its epistemic function just as blunt knives are failures on the part of knife producers to fulfil their economic function.

On a knowledge-first picture, good belief is knowledgeable belief. Note, also, that knowledge meets E-Function: (1) our cognitive processes have produced it in the past (in our friendly epistemic environment, knowledge is readily available), (2) this benefitted our ancestors (the widely accepted value of knowledge thesis), which (3) contributes to the explanation of why our cognitive systems continue to exist. So generating knowledge is *a* function of our cognitive processes. Is it the main function? Note that what determines the proper level of generality for main function individuation is the T-value of the relevant T-function. The main biological function of the heart, for instance, maps on to its most valuable biological contribution: its main function is not 'pumping blood and making a ticking sound', but neither is it merely 'pumping', for instance. Plausibly, that is because if the heart pumps orange juice in our circulatory system, that's not very valuable for our survival. I submit that knowledge is more valuable than any lesser epistemic standing. That much is very widely accepted in the literature.^{xxi} It is easy to see that, if I am right about main function individuation, the distinctive value of knowledge thesis, together with E-Function, deliver the result that the main epistemic function of our cognitive processes is generating knowledge.

Here is the view, fully spelled out:

Knowledge First Functionalism (KFF): S's belief B is justified if and only if it is generated by a properly functioning cognitive process that (1) is a token of a type that has generated knowledge in the past, (2) it's having generated knowledge resulted in epistemic benefit in S/S's ancestors and (3) the fact that generating knowledge has epistemically benefitted S/S's ancestors contributes to the explanation of why the process exists in S.

The etiological account is an account of functions as purposes: by being selected for it, our hearts have acquired the purpose of pumping blood in our organisms (Graham 2012, 449). Reaching that purpose—i.e., successfully pumping blood—will amount to function fulfilment. But purposes will also come with associated norms prescribing the right way to proceed in order to reliably reach the corresponding purpose. Because its function contributes to the explanation of its very existence, the trait in question *ought* to perform in a

way that is associated with likely function fulfilment. Now, according to the etiological theory of functions, this is but the way in which the trait functioned back in the day when it acquired its function. Your heart will pump blood in normal conditions, i.e., conditions similar to those in which it was selected, when functioning normally, that is, when functioning in the way in which it was functioning when it was selected for its beneficial effects. Plausibly, in normal conditions, a normally functioning heart will fulfil its function of pumping blood in your system by beating at a particular rate. According to the etiological theory, then, normal functioning is proper functioning: a heart functions in the way it should (i.e., by the norm) when it functions in the way it did back in the day when it acquired its function: when it beats at a particular rate.

Note, then, that there are two ways a functional device might go right, and two ways it may go wrong. The unhappy cases are: breach of the norm, i.e., malfunction (in the case of the heart, not beating) and failure to reach the corresponding purpose, i.e. failure to fulfil its function (not pumping blood) (Graham 2012, 449). The happy scenarios are, of course, proper functioning (beating) and function fulfilment (pumping blood).

Crucially, failure/success in the former respect need not imply failure/success in the latter: proper functioning need not imply function fulfilment. To see this, think of a situation where a surgeon takes the heart out of your chest, places it in a vat full of nutrients for a short while and plugs it to a pipe circuit filled with orange juice. Your heart, of course, will fail to fulfil its function of pumping blood under these circumstances; it will, as a matter of fact, be pumping orange juice. But this does not make it into a malfunctioning heart; to see this, compare it to a heart that has stopped pumping blood because it has been stabbed by a dagger. When functioning normally—whether in normal conditions or not—your heart will function properly, i.e. it will meet the norm constitutively associated with the purpose of pumping blood. It will work the way it is supposed to work, where the right way of working is partly constituted by fulfilling its function in normal conditions.

To return to belief: on this view, epistemic justification supervenes on the proper performance of cognitive systems that have generating knowledge as their epistemic function. The standards for proper functioning are thus constitutively associated with promoting knowledgeable beliefs.

Of course, proper functioning need not imply function fulfilment; our cognitive processes can function normally, but still fail to produce knowledgeable beliefs. It is easy to see that, in virtue of this feature, KFF enjoys excellent **NORMATIVE GENERALISABILITY**: just like in many other normative domains, epistemic justification and epistemic success can come apart: our belief acquisition processes can function properly and still fail to fulfil their function, due to unfriendly environmental conditions.

In virtue of its normative richness, KFF also scores well on **NORMATIVE EVIDENTIAL ADEQUACY**. Gettierized, deceived victims, and conscientious false believers alike, employ properly functioning cognitive processes that have the function of generating knowledge, and thus come out justified. In contrast, mere blameless believers such as Ben do not—albeit, through no fault of their own.

Last, there are two ways in which KFF does exceptionally well on **NORMATIVE PRIOR PLAUSIBILITY**. First, of course, in virtue of being a variety of proper functionalism, on KFF, epistemic normativity drops right out of functional normativity, and, thereby, it has a nice explanation of how normativity exists in nature. This is not all there is to it, though. Note that KFF employs a historical account of functions, and justification is analysed in terms

of knowledge. The cognitive processes at stake need to first produce the success—i.e. knowledge—a few times, in order to acquisition the relevant function via producing the relevant type-specific benefit. Since justification supervenes on proper functioning, and there is no proper functioning unless the trait in question already has a function, justification will come second temporally. This picture receives strong naturalistic support: after all, this is how evolution works: it all starts by something good, a good trait or mutation which is randomly produced, and then reproduced in future generations due to its benefits. Note, also, that this story gets further support if one accepts the Williamsonian ‘knowledge is a state of mind’ picture. After all, if knowledge is a state of mind in its own right, it makes sense for it to be generated independently, rather than to be generated from a different state of mind.

6. Conclusion

Knowledge comes first in our epistemic affairs; keeping this value-theoretic point fixed, I have asked what the most normatively fit way to conceive of epistemic justification might be. I have argued that the geography of our normative landscape has difficulties in accommodating popular knowledge first accounts of justification, and I have defended an alternative view. On this account, beliefs are justified if and only if generated by properly functioning cognitive processes that have generating knowledge as their epistemic function. The account is proper functionalist in that it takes the epistemic normativity of belief to drop out of the epistemic function of our cognitive processes. It is knowledge-first epistemological in that, unlike traditional proper functionalism, it unpacks the function at stake in terms of knowledge.

References

- Bird, A. (2007) Justified Judging. *Philosophy and Phenomenological Research* 74: 81-110.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Brown, J. (2018). *Fallibilism: Evidence and Knowledge*. Oxford University Press.
- Buller, D. (1999). *Functions*. Albany: SUNY Press.
- Cohen, S. and Lehrer, K. (1983). Justification, Truth, and Knowledge. *Synthese* 55.
- Chisholm, R. M. (1966). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- Geach, P. (1956). Good and evil. *Analysis*, 17: 33-42.
- Gerken, M. (2011). Warrant and action. *Synthese*, 178, 529-547.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23: 121-23.
- Godfrey-Smith, P. (1993). Functions: Consensus Without Unity. *Pacific Philosophical Quarterly* 74: 196-208.
- Goldman, A. (1988). Strong and weak justification. *Philosophical Perspectives*, 2: 51-69.
- Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73:771-91.
- Graham, P. J. (2010). Testimonial Entitlement and the Function of Comprehension. *Social Epistemology*, D. Pritchard, A. Millar, A. Haddock, eds.,

- OUP:148-174.
- Graham, P. J. (2012). Epistemic entitlement. *Nous*, 46(3): 449-482.
- Graham, P. (2014). Functions, warrant, history. Fairweather, A. and Flanagan, O., eds., *Naturalizing Epistemic Virtue*. Cambridge University Press: 15-35.
- Greco, J. (2010). *Achieving Knowledge*. Oxford: Oxford University Press.
- Greenough, P. (2012). Discrimination and Self-Knowledge. in D. Smithies and D. Stoljar (eds) *Introspection and Consciousness*, Oxford: Oxford University Press.
- Ichikawa, J.J. (2014) Justification is Potential Knowledge. *Canadian Journal of Philosophy* 44/2: 184-206.
- Kelp, C. (2018). *Good Thinking: A Knowledge First Virtue Epistemology*. New York: Routledge.
- Kelp, C. (2018). Assertion: A Function First Account. *Nous* 52: 411-42.
- Kelp, C. and Simion, M. (2017). Commodious Knowledge *Synthese*, vol 194/5: 1487-1502.
- Kelp, C. (2016). Justified Belief: Knowledge First-Style. *Philosophy and Phenomenological Research* 93, 79-100.
- Littlejohn, C. (Forthcoming). A Plea for Epistemic Excuses. In F. Dorsch and J. Dutant (eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- McHugh, C. (2012) The Truth Norm of Belief. *Pacific Philosophical Quarterly*, 93, (1), 8-30.
- Millar, A. (2010) Knowledge and Recognition. In Duncan Pritchard, Alan
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: The MIT Press.
- Millikan, R. (2000). *On Clear and Confused Ideas*. Cambridge University Press, Cambridge.
- Miracchi, L. (Forthcoming). Competent Perspectives and the New Evil Demon Problem. *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, Oxford University Press, eds. Fabian Dorsch and Julien Dutant.
- Miracchi, L. (2015). Competence to Know. *Philosophical Studies*, 172 (1): 29-56.
- Neander, K. (1991). The teleological notion of function. *Australasian Journal of Philosophy* 69: 454-468.
- Reynolds, S. (2013) Justification as the Appearance of Knowledge. *Philosophical Studies* 163: 367-83.
- Simion, M. (Manuscript). *Epistemic Norms and Epistemic Functions*.
- Simion, M. and Kelp, C. (Forthcoming). The Constitutive Norm View of Assertion. In Goldberg, S. (ed), *Oxford Handbook of Assertion*. Oxford University Press.
- Simion, M. (2018). Saying and Believing: The Norm Commonality Assumption. *Philosophical Studies*.
- Simion, M., Kelp C. and Ghijssen, H. (2016). Norms of Belief. *Philosophical Issues, A Supplement to Nous*. vol. 26/1: 375-392.
- Simion, M. and Kelp, C. (2016a). The Tertiary Value Problem and the Superiority of Knowledge. *American Philosophical Quarterly*, vol. 53/4: 397-411.
- Simion, M. (2016b). Perception, History and Benefit. *Episteme*, 13 (1), 61-76.
- Sosa, E. (1993). Proper Functionalism and Virtue Epistemology. *Nous*, 27: 51-65.
- Sutton, J. (2005) Stick to What You Know. *Nous* 39: 359-96.
- Sutton, J. (2007) *Without Justification*. Cambridge, MA: MIT Press.
- Thompson, J.J. (2008). *Normativity*. Carus Publishing.
- Walsh, D. (2008). Teleology. In Ruse, M. (ed) *The Oxford Handbook of Philosophy of Biology*. Oxford University Press: 113-138
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.

Williamson, T. (forthcoming). Justifications, excuses, and sceptical scenarios. In F. Dorsch & J. Dutant (Eds.), *The new evil demon*. Oxford: Oxford University Press.

Zimmerman, M.J. (1979). Moral Responsibility and Ignorance. *Ethics* 107, no. 3: 410-426.

Wright, L. (1973). *Functions*. *The Philosophical Review* 44: 409-422.

ⁱ For the purposes of this paper, I will bracket the phenomenon of defeat. For ⁱⁱ I will use ‘justification’ as shorthand for ‘epistemic justification.’

ⁱⁱⁱ See, however (Simion 2019) for why we should prefer proper functionalism over virtue epistemology on a more general normative level, pertaining to the resources that these accounts have for accommodating ignored evidence, propositional warrant and defeat.

^{iv} See Keas (2017) for a nice taxonomy.

^v That is not to say that epistemic normativity should behave similarly to all, or even most other types of normativity. Rather, it should not be the case that, on the account proposed, epistemic norms contradict overwhelmingly plausible claims about general, un-typed normative behavior.

^{vi} This is not to say that epistemic justification should behave in identical ways to all other types of justification, but rather that it will nicely drop out of a general account of (un-typed) justification. Domain-specific differences between species of the general type are to be expected, of course.

^{vii} See (Cohen and Lehrer 1983); throughout this paper, when I refer to deceived victims, the case I have in mind is the version where an adult normal perceiver is abducted and deceived. Two reasons for this: I find the intuition of positive epistemic normative status in this case much stronger, and (2) I want to avoid complications pertaining to concepts acquisition.

^{viii} See (Gettier 1963) and (Goldman 1976).

^{ix} Clayton Littlejohn (pc).

^x (Littlejohn Forthcoming), (Williamson Forthcoming).

^{xi} For worries, see (Gerken 2011), (Brown 2018).

^{xii} See also (Kelp 2016, 2018), (Miracchi Forthcoming) and (Simion et al. 2016).

^{xiii} For normative distinctions along these lines, see also (Simion et al. 2016), (Greenough 2012), (McHugh 2012), and (Jarvis-Thompson 2008).

^{xiv} See (McHugh 2012) for discussion on the evaluative vs. prescriptive status of a truth norm of belief.

^{xv} This is not to say that they need be action guiding in a strongly internalist sense. On the contrary, I take it that most prescriptive norms are not thus operationalizable.

^{xvi} This also gets further confirmation in several other places in Williamson (2000, 2014). For instance, according to Williamson, “[k]nowledge sets the standard of appropriateness for belief. [...] Mere believing is a kind of botched knowing. In short, belief aims at knowledge” (2000, p.47).

^{xvii} Defended by people like David J. Buller (1999), Ruth Millikan (1984), Karen Neander(1991), Peter Godfrey-Smith (1994) and Larry Wright (1973).

^{xviii} For application of the etiological account of functions to the normativity of assertion, see e.g. (Graham 2010), (Kelp 2018) and (Simion Forthcoming, 2018).

^{xix} This addresses early worries against natural selection based proper functionalist accounts, i.e. Sosa’s appeal to Swampman (1993).

^{xx} See (Simion 2016b) for reasons against reducing epistemic functions to biological functions.

^{xxi} See (Kelp and Simion 2017), (Simion 2016a).