

Justified Judging

Alexander Bird

(forthcoming in Philosophy and Phenomenological Research)

1. Introduction

When is a belief or judgment justified? One might be forgiven for thinking the search for single answer to this question to be hopeless. The concept of justification is required to fulfil several tasks: to evaluate beliefs epistemically, to fill in the gap between truth and knowledge, to describe the virtuous organization of one's beliefs, to describe the relationship between evidence and theory (and thus relate to confirmation and probabilification). While some of these may be held to overlap, the prospects for fulfilling all may well seem poor. Furthermore the internalist requires that justification be an introspectible property of beliefs and a fundamental epistemic concept, while the externalist is often happy to ignore the concept altogether or at best regard it as an embarrassing add-on to their epistemology. In the light of this one might reasonably give up on justification altogether or adopt pluralist approach, denying that justification is any single property of beliefs of judgments.

In this paper I shall propose a unified account of justification, providing necessary and sufficient conditions for a judgment or belief to be justified. It is an externalist account. Yet it provides a natural, non-redundant role for the concept and at the same time it satisfies at least one key internalist intuition. The account relates justification to the concept of knowledge but can also explain the role of justification both in evidential support and in the proper organization of one's beliefs. That a single account can provide such a unified account may be a surprise for the reasons given, and indeed clearly not all the diverse claims made about justification can be expected to come out true. Nonetheless, it is revealing that many can.

The account is intended to supplement Timothy Williamson's project in epistemology, the key feature of which is that knowledge takes centre stage. Williamson's project is avowedly externalist. Knowledge is a factive mental state—the most general factive mental state (Williamson 1995). Knowledge is prime, a unitary mental state, not decomposable as a hybrid of

an internal condition and an external condition. Williamson's project thus stands in contrast to a long tradition in epistemology that has sought to provide an account of knowledge in terms of belief and justification, of which the simplest is the Justified True Belief account.¹ That account could be taken, and often was, to provide a internal component (justified belief) and an external component (truth). As regards the internal component, foundationalism and coherentism are the standard accounts of the structure of justification. In so doing they also provide accounts of what justification is and why it is that a belief is justified—a belief is justified *because* it stands at the end of a chain of appropriate inferences starting from a foundationally justified belief or *because* it of its place in a coherent system. In this tradition justification plays the central and basic epistemic role.² I shall not discuss the problems of the tradition beyond the following remarks. First, to the extent that coherentism and certain versions of foundationalism aim to provide internalist accounts of justification, problems continue to plague both.³ Secondly, the project of repairing the JTB account of knowledge in response to Gettier's counterexamples (by finding some 'X' for 'knowledge = (true) belief + X') is about as degenerative a research programme as one could wish for.

Thirdly, the attempted post-Gettier accounts of knowledge, typically externalist, have found it difficult to accommodate the concept of justification. In particular they run the risks of being too strong (i.e. entailing truth) and of being too far removed from the sorts of intuition that motivate internalism to be convincing. For example a reliabilist position may initially assert: knowledge = reliably formed true belief, where reliability does not entail truth. And we can give a corresponding reliabilist account of justified belief as reliably formed belief. But this is subject to its own Gettier counter-examples. So we may make it a requirement that reliability does entail truth. But then justification would entail truth, and indeed justified belief would be knowledge, which, intuitively, is too strong. And in both cases the conception of reliability is heavily external, so that two subjects may have formed the same belief in the same way, yet one be

¹ The JTB account has often been taken to be an intended *analysis* of knowledge. I shall take it simply as an intended necessary equivalence.

² Indeed for some the problems of scepticism have suggested abandoning knowledge altogether. Wright argues that knowledge is not the concern of epistemology since we can happily concede that we do not know so long as we are nonetheless justified (Wright 1991, 88).

³ If this paper is right, then the question that foundationalism and coherentism (whether internalist or externalist) attempt to answer, what is the structure and source of justification? reduces to the question, what is the structure and source of knowledge? It *may* be proper to give a coherentist or foundationalist answer to that question, though the account will inevitably be an externalist one, if the Williamsonian perspective I adopt is correct. See below . . .

justified and the other not. One may be a friend of externalism in general yet sympathize with the internalist in finding this counterintuitive.

These problems may be avoided by giving an account of justification within the framework Williamson has provided. Justification is not the central epistemic concept, knowledge is. And so my account of justification is an account in terms of knowledge. Thus the direction of explanation is the reverse of that in the traditional project. Jonathan Sutton reaches the same conclusion from consideration of the Gettier cases:

... any genuine concept of justification at work in our so labelling this category of justified belief is parasitic on the concept of knowledge. We only understand what it is to be justified in the appropriate sense because we understand what it is to know, and can extend the notion of justification to non-knowledge only because they are would-be knowers. We grasp the circumstances—ordinary rather than extraordinary—in which the justified would know. Justification in the relevant sense is perhaps a disjunctive concept—it is knowledge or would-be knowledge. (Sutton forthcoming)

This captures the thrust of my argument. In what follows I shall explicate what counts as ‘would-be knowing’—it is a certain kind of approximation to knowledge. Sutton develops this into an extreme view, that a belief is justified only if it amounts to knowledge. He too starts from a position of sympathy with Williamson’s arguments. I shall show how the latter can be endorsed while taking the more natural view that some cases of justification fall short of knowledge. (I address Sutton’s view in §5.)

Knowledge and justification are externalist concepts according to my proposal. Nonetheless this account of justification does endorse a central internalist intuition, that justification supervenes on mental states. (How these may be reconciled I explain below.) Furthermore, the account of justification also satisfies the desideratum that knowledge entails justification, i.e. necessarily, S knows that $p \rightarrow S$ is justified in believing that p —but not vice-versa. Williamson’s primeness view means that the mental state of knowing is not a sub-species of the state of believing. Nonetheless that is consistent with the claim that knowing entails (justified) believing, that in all worlds where S knows that p , S (justifiably) believes that p . I shall take this to be the case in what follows.⁴

⁴ Williamson is agnostic on the point. If however there are cases where knowing is not accompanied by belief (e.g.

2. *The account of justification*

Knowledge is epistemically central. Justified belief is a certain kind of approximation to knowledge. It is an approximation that is independent of one's mental states. If one attempts to know but fails for some reason that is located outside one's mental states then one's belief is justified. But if one's failure to know is due to some feature internal to one's other mental states, then the judgment is not justified. Here the notions of 'reason' and 'due to' are constitutive not causal. Thus the falsity of a belief is one factor that, in the constitutive sense, explains why it does not amount to knowledge. The unreliability of the belief forming mechanism would be another such factor. Our focus is on those features of the world that are lacking or awry when believing fails to yield knowing.

The account is as follows:

(JuJu) If in world w_1 S has mental states M and then forms a judgment, that judgment is justified if and only if there is some world w_2 where, with the same mental states M, S forms a corresponding judgment and that judgment yields knowledge.

(JuJu) respects various intuitions concerning justification, including some internalist ones, as I shall argue in a later section. According to (JuJu) knowledge entails justification. But not because knowledge is some special subset of justified belief, but rather because knowledge is the standard against which beliefs are compared to see whether they count as justified. Furthermore we can see why Gettier's cases may arise. Let us call a state of affairs *mentally extraneous* if it does not supervene on a subject's mental states. If the only mentally extraneous reason for a failure to know were the falsity of a belief then a true justified belief would always yield a case of knowledge. But the fact that there can be mentally extraneous reasons for failure to know that are compatible with the truth of a belief entails that there can be cases of true justified belief that are not knowledge. A normally reliable and mentally extraneous belief-forming process may suddenly cease to be reliable but the belief it forms is accidentally true nonetheless. Since there

Radford's cases, or cases of low-level perceptual knowledge), then the question of justification does not arise. We would have to modify the claim 'knowledge entails justification' to 'knowledge that p entails that if the subject believes that p , then that belief is justified'.

is another possible world in which one has the same antecedent mental states and in which the same process remains reliable, thereby generating knowledge, it follows from (JuJu) that in the actual world the belief in question is justified—it is a true justified belief that isn't knowledge.

(JuJu) tells us that justified judging is a certain kind of approximation to knowledge. It is an approximation where the failure to know (if any) is explained by factors external to the subject's mental states. More succinctly, justified judging is possible knowing (where one world is accessible from another if and only if they are identical with regard to a subject's antecedent mental states and judgment forming processes).

This account not only provides a useful role for the concept of justification, it helps us see why there should be beliefs that are justified but fall short of knowledge. Our cognitive processes typically depend on the cooperation of the environment for their success. But we should not expect that cooperation to be infallible. On occasion the environment will not be propitious. Thus we should *expect* there to be states of judging that are just like knowing, except that that environmental—mentally extraneous—factors undermined the processes that would have led to knowledge.

Note that (JuJu) is diachronic. It tells us about the justification of a belief, given earlier mental states and a belief-forming process. It does not state a synchronic condition on justification, the most straightforward of which would be: S is justified in believing that *p* iff there is a possible world where with the same mental states S knows that *p*. This will not do since knowing is itself a mental state. That condition on justification would require that to be justified in believing that *p* S must know that *p*. A diachronic condition seems appropriate. We think that a belief is justified because of the way in which it came to be formed.

Typically the relevant antecedent mental states, M, that we need to consider are just those the subject possesses at the start of the judgment forming process. But sometimes it would be appropriate to extend M to include more or indeed all of the subject's mental history, for the reason that what we once knew but have (perhaps culpably) forgotten is relevant to what we are not justified in judging. Let *p* be misleading evidence such that anyone who knows *p* cannot justifiably believe *h*. But *h* is in fact true, and it is possible to know *h*, so long as one doesn't

know p . S does know p and is aware of its relevance to h . However, at t_1 S culpably forgets p and comes to acquire all the good evidence for h . At t_2 , reflecting on this evidence, S judges that h is true. The normative element of justification tells us, it can be plausibly argued, that S is not justified in judging h to be true. But if we consider just S's mental states at t_2 then there is an S* like S as regards her mental states at t_2 who never knew p at all and so who does know h as a result of reflecting on the evidence she has. (JuJu) would thus wrongly tell us that S is justified. To avoid that conclusion we must regard M as including not just S's mental states at t_2 but also all S's relevant prior mental states. Thus any comparison S* must a be someone who had previously known p and so is barred from knowing h . Indeed, I think we could in every case extend M to include a person's total mental history. After all, when asking whether S is justified in making a judgment we are asking, on this view, whether S herself could have gained knowledge thereby—which implicitly means S just as she is, including her mental history, up to the point of judgment. Another way of looking at this is in counterfactual terms: had mentally extraneous matters alone been otherwise and propitious for knowing, S would have come to know by judging.⁵ Given Lewis' constraints on nearness of possible worlds, this requires us to consider worlds where S's mental history is as it actually is.⁶

As it stands (JuJu) could be accepted in conjunction with a variety of differing accounts of what knowledge is—as it stand it is consistent with a reliabilist analysis of knowledge and also with contextualism. Nor does the development of the account in the next section does not depend upon any particular account. Nonetheless, I do also believe that the proposal gains most plausibility by being embedded in the approach to epistemology briefly referred to in the introduction, that found in Williamson's *Knowledge and its Limits*. In Kuhnian terms this proposal may be thought of as a puzzle-solution within the Williamsonian paradigm: given the 'knowledge first' approach, what should one then say about justification? I concede that the paradigm is not yet widely adopted, and in particular not the central claim that knowledge is a genuine mental state. That said, it is also true that the success of this proposal *within* the terms of that paradigm provides further testament to its power.

⁵ This counterfactual formulation can be shown to be equivalent to (JuJu) when M is taken to include all of S's mental history.

⁶ Lewis 1986, 43-48.

3. *Ways of judging*

(JuJu) says that a judgment is justified if a *corresponding* judgment in some (possibly other) world is knowledge. I do not say ‘the *same* judgment’ and here is why. S calculates 659×953 with a calculator which hitherto has been perfectly reliable. The calculator gives the answer 625027. In fact the answer is 628027, but two of the LEDs in the calculator are no longer functioning. Since S is not sufficiently good at mental arithmetic to spot this error, he judges that $659 \times 953 = 625027$. Because S’s judgment is false, he does not know. Is he justified? It seems clear that he is. But note that there is *no* world where S knows that $659 \times 953 = 625027$, since the latter is a necessary falsehood. The natural thought is that S is justified because had the LEDs continued to function as they had previously done, then S would have judged correctly that $659 \times 953 = 628027$; indeed, given this continued reliability of the calculator, S would have known that $659 \times 953 = 628027$. In this case S’s actual judgment is that $659 \times 953 = 625027$ and the corresponding judgment in another possible world, which amounts to knowledge, is that $659 \times 953 = 628027$.

What counts as a corresponding judgment seems intuitively clear and my account could rely on its being so. Nonetheless, it will be useful to say some more about when a judgment in some other possible world corresponds to an actual judgment. There are two components to correspondence: content and process of formation. Clearly a corresponding judgment should be at most a minor variant in *content* on the actual judgment. If the calculator had badly malfunctioned, and given $659 \times 953 = 48$, then S would not be justified in forming that judgment. What counts as a corresponding belief is restricted not only by the content of the judgment but also by the *process* whereby the judgment was formed. In the original case of S we ascribed justification because we compare the actual circumstance to one where S forms a belief in the same way and comes to know. In general, two judgments are brought about in the same way if they are the produced by the same mental dispositions and capacities. When in the actual world S uses a faulty calculator while in another world S uses a reliable calculator, they are operating with the same mental dispositions, such as the disposition to believe the read-out of the calculator.

Note that the issues of minor variation in content and judgment forming process are related, for what counts as a minor variation in content will depend on the subject’s dispositions and capacities. Consider another version of our case, where S is a research scientist, competent at

mental arithmetic. This time the faulty calculator gives $659 \times 953 = 628020$ and S judges that $659 \times 953 = 628020$. Now S ought to have seen that this cannot be the right answer, since two numbers ending in 9 and 3 will never multiply to give a number ending in 0, but always in 7. Furthermore let us assume that self-deception causes S to judge that $659 \times 953 = 628020$ rather than spot the error, since 628020 accords better with the predictions of his pet theory. So S is not justified in judging that $659 \times 953 = 628020$. Considering the content alone '659x953=628020' would seem just as minor a variant (or less of one) on '659x953=628027' as '659x953=625027'. However, when we consider that a competent user of mathematics, such as S, can normally detect an error in the last digit of a long multiplication but not an error in a middle digit, then we see that in this context the former is not a minor variant. What counts as corresponding judgment is one that is a minor variant relative to the process of judgment formation.

That an account of justification should refer to the process by which a judgment is formed is no surprise, as I remarked above when motivating the diachronic approach. Justification marks the fact that a subject formed a judgment in a legitimate way. This account tells us that a legitimate way is one that could have led to knowledge. Nonetheless, talk of the 'way' or 'process' by which a subject forms a judgment is vague and might remind us of the difficulties that reliabilists have in providing criteria for the appropriate characterization of methods or processes of belief formation. For example, when a subject believes something in a newspaper report, is the method 'believing what one reads in print', 'believing something in a report in the *Times*', or 'believing a report written by a reliable, well-respected, and disinterested journalist' etc.? Depending on the characterization of the method, its reliability may be evaluated differently.

For two reasons I can largely finesse such problems. First, because I am interested only in *sameness and difference* of process in subjects that are otherwise in identical mental states. And secondly, because I am interested only in sameness and difference as regards the deployment of a subject's mental capacities and dispositions. Thus in the original calculator case, it is easy to see that when in w_1 S uses a faulty calculator while in w_2 S uses a reliable calculator, they are operating with the same mental dispositions, without having to be able to say precisely what the method is that S uses.

((JuJu) could be combined with a reliabilist account of knowledge, e.g. knowledge is (true) belief produced by a reliable process, in which case one would also require an account of ‘process’.⁷ But that is a problem for the account of knowledge, not for this account of justification.)

The operation of the same dispositions (i.e. the same process leading to a corresponding judgment) *may* be accompanied by the acquisition of different mental states along the way. In the calculator case S in w_1 has a visual experience of the calculator display showing ‘625027’ while S in w_2 has a visual experience of the display showing ‘628027’. This does raise a different issue concerning the individuation of processes of forming a belief. S is talking to a friend who is normally entirely reliable. At t_0 , the friend tells S that p , and so at t_1 , very shortly after t_0 , S believes that p . Although normally reliable, the friend has on this occasion made an uncharacteristic error—it is false that p . At t_2 , immediately after t_1 , S infers that q from the belief that p (by an unproblematic and obvious inference). S’s belief that p is justified, and (JuJu) tells us so, since there is a world where the friend does not make the slip of the tongue and S comes to know the corresponding proposition that p^* . However, is S’s belief that q also justified? Consider S’s antecedent mental states at t_1 , which include the false belief that p and hence does not include knowledge that p . There is no world where with those antecedent states, which don’t include knowledge, one can come to know q by inference. So it looks as if (JuJu) says that S is not justified in believing q . But this seems counterintuitive. However, if we consider S’s antecedent mental states at t_0 , there may well be a world where S comes to believe and to know q or a counterpart of q . Thus in a world where it is true that p and the friend knows this, then S will come to know q . Or in a world where the friend correctly informs S that p^* (rather than p) and S infers q^* . If we consider the process of coming to believe that q/q^* on the basis of hearing that p/p^* as a single process (albeit one with parts), similar to the calculator case, then the relevant antecedent mental states are those at t_0 and we have seen that there may well be another possible world where the same belief forming process starting from those mental states and leads to knowledge.

The question then is, when do we have a sequence of mental states a *single* belief-forming process (with parts) (i.e. the manifestation of a single disposition) and when is the sequence merely a sequence involving several processes. (JuJu) does not itself help provide an answer but

⁷ Again I note that I believe that (JuJu) is most appropriately situated within Williamson’s epistemology according to

depends on an answer being provided from elsewhere. That we do make such a distinction is clear, even if the distinction is itself vague. A baseball player may score a run by hitting a home run or by advancing through the bases one by one; the former is a single process, the latter is not. There are clear cases where a sequence of beliefs is a single process. For example the rules of conversational implicature or the employment of tacit assumptions may be such that a speaker will utter '*p*' with the primary intention of telling us that *q*. We may believe that *p* along the way, but that is often part of a single process leading to the almost immediate belief that *q*. On the other hand if the belief that *p* comes as part of an extended inquiry whose eventual conclusion is that *q*, then we have multiple processes. It may be up to cognitive science to tell us which compound processes are genuinely single and which are multiple. One relevant feature may be whether the process is hard-wired as a single process whose sub-processes are not easily dissociated—or whether the whole process involves the operation of component processes that normally operate independently.

That an account of justification should stand in need of such a distinction should not come as a surprise to the externalist about justification. Consider S who leads a normal life and comes to know many things. Now consider S* whose life was just like S's until *t* (a long time ago) when at a young age he was envatted; S* has been given the same sensory stimuli as S ever since and so comes to have the same beliefs. (Two things: Perhaps the match between S's and S*'s stimuli is accidental; there is no reliability in S*'s case. Secondly, we'll ignore issues concerning the reference of S*'s thought-components.) Where S knows that *p* S* merely believes that *p*. The externalist about justification may be inclined to deny that S* is even justified in believing that *p*. S*'s beliefs are so far from being connected with reality that they have almost no positive epistemic status at all. However, if we regard S's life history of sensory stimuli since *t* as a *single* process whose product is the belief that *q*, then we have to say that S* formed the same belief in the same way. *If* that were a legitimate perspective then it would be difficult to deny that S* is justified. Certainly (JuJu) gives that conclusion and some species of reliabilism about justification may well do so also, since the process is one that would normally yield true beliefs. Internalists are much more inclined than externalists to hold that S*'s beliefs are justified. One diagnosis of this is that the internalist does adopt the perspective whereby S*'s life history is compared to the life history of S, implicitly thinking of them as identical ways of forming a

which a reliabilist analysis of knowledge is not possible.

belief. However, in my view, that should *not* be taken to be a legitimate perspective. A life history is not a way of forming a belief but is a concatenation of many processes, a concatenation that has no intrinsic unity. Although I reject the internalist's take on this case, I will later show how (JuJu) does respect a different but related internalist intuition.

4. The 'transmission' of justification

A consequence of the preceding discussion is that we need not expect justification to be transmitted in every case by (legitimate) inference from a justified belief. This stands in direct contrast, for example, to a foundationalist view whereby justification originates with certain foundational beliefs and is spread around by inference (rather like a communicable disease). In general I reject any view of justification as the epistemically fundamental property whose transmission may be followed by knowledge should additional conditions are met. That there is something wrong with the latter should already be clear. For one can have a justified but logically false belief from which one eventually draws by inference a patent contradiction. In such cases the justification does not get transmitted to the conclusion. It is true that in such cases the original proposition loses its justification—but there is no obvious and immediate explanation of that fact in terms of the picture of justification as a quantity that gets transmitted by inference. I replace this picture (knowledge follows (in propitious cases) from the transmission of justification) by the reverse—it is knowledge that can be transmitted by suitable (e.g. deductive) inference while justification that piggy-backs on (the possibility of) knowledge. Knowledge is spread by such inferences because they preserve a reliable connection between true belief and corresponding fact. It is because of the possibility of spreading knowledge by inference that justification too can typically be transmitted in this way (even in cases where knowledge is only possible and not actual). Consider a straightforward case. Let S have justified belief that p . Let S infer from the proposition that p the proposition that q , by a simple deduction. We need to show that S's belief that q is also justified. Since S is justified in believing that p , by (JuJu) there is a possible world w (which might or might not be the actual world) in which S has some corresponding knowledge. In this case let us assume that it is knowledge that p . Let it be that in some world w^* near to w (which might be w itself) in which S also knows that p , S infers that q by the aforementioned simple deductive inference. Since knowledge is transmitted by such inferences, in w^* S thereby comes to know that q . Since w^* is a possible world in which S comes to know that q by the same process by which S comes to believe that q in the actual world, then

by (JuJu) S's belief that q in the actual world is justified. In other cases the corresponding knowledge may not be knowledge that p but knowledge that p^* . For many such cases there will be a q^* corresponding to q (i.e. q^* is deducible from p^* just as q is from p). Since S knows that p^* in w^* and deduces q^* , S knows q^* in w^* , and hence, once more, by (JuJu) S's belief that q , in the actual world, is justified. However, there need not always be a corresponding q^* . Say that p is a complex logical falsehood justifiably believed on the basis of expert testimony. Let q be a manifest contradiction deduced from p . Since p is justified there is a corresponding proposition p^* known to S in some world w . But there will not be a q^* standing to p^* and q stands to p . Since p^* is known it is true, and so any q^* deducible from it will be true also—and so cannot be a manifest contradiction. In short (JuJu) has the consequence that justification cannot be transmitted to a manifest contradiction because there is no possible case of a corresponding transmission of knowledge.

The above suggests that my account will deny justification where some other accounts may (mistakenly) ascribe it. I now want to consider a possible objection that *prima facie* suggests that my account is too liberal with ascriptions of justification. A judgment is justified if and only if, with the same mental states, including belief-forming dispositions, one could in some possible world have formed a corresponding belief in the same way and thereby achieved knowledge. This means that in some circumstances one can use thoroughly unreliable methods of belief formation and yet be justified. That conclusion may lead to the following objection. Imagine that S forms beliefs by examining traces of tea leaves. If there is some other possible world in which tea leaf reading is a reliable method of belief formation, then it may be that in such a world S's counterpart forms a belief in the same way as S and gets to know. In which case we must say that S is justified in this world. This may seem to some to be counter-intuitive. How can in this world tea leaf reading lend justification to a belief? In response, let us first exercise our intuitions in the other direction. The challenge asserts that there really is a world in which S* gets to know things by reading tea-leaves. Note that S is mentally just like the knower S*. If one does everything that is done by someone who gets to know, then surely one is at least justified in one's belief. We may then try and explain why the case seems counterintuitive. We may think that the world in which tea leaf reading is reliable is so remote that it would be quite difficult, perhaps impossible, to find an S* who is mentally like S in our world. At the very least, S and S* would have to be very unsophisticated indeed. For in the actual world there is ample evidence that tea leaf reading and other such methods of belief formation are unreliable. If S knows that tea leaf

reading is unreliable, then, *ex hypothesi*, S* knows this too. But if S* knows this, then it is true (in S*'s world) that tea leaf reading is unreliable. In which case there is no such world as the case suggests, and on my account S is not justified in having beliefs based on tea leaf reading. Even if S is aware of the evidence that this is a world in which tea leaf reading is unreliable, then, since S* must be aware of that evidence too, that may be enough either to make S*'s world a world in which tea leaf reading is unreliable or one where, if it is reliable, S* cannot use it to know (since S* has beliefs that strongly point to its unreliability). Thus if S is to have justified beliefs from tea leaf reading S will have to be ignorant of a lot of facts we take for granted. As regards my account, the verdict on this case is that it is under-described. Filled out in some ways, whereby S is unsophisticated and ignorant of many things we know, S might be justified, but filled out in other ways, whereby S knows much of what we do, S is unjustified. This explains why our intuitions may go either way.

This looks to be a plausible result. Knowing things may broaden the range of new things one can get to know and justifiably believe. But it can also *constrain* what one can justifiably believe. Conversely, ignorance may *increase* the range of things one may justifiably believe. This of course flies in the face of foundationalism, which suggests that increased justification *always* requires a greater rather than smaller evidence set. But my conclusion is by no means paradoxical. In the extreme if I know that p , I cannot justifiably believe that $\neg p$, but if you don't know that p you may well be justified in believing that $\neg p$. More generally, we do not condemn individuals in the past or in other kinds of society as having a raft of unjustified beliefs, even if members of our society could not justifiably have those beliefs. For example, S might justifiably believe in witchcraft if brought up a society where otherwise reliable informants tell S that witches exist and present S with only the favourable evidence. The additional evidence we possess may enable us to justifiably believe sophisticated scientific propositions S cannot justifiably believe, but at the same time such evidence may also prevent us from justifiably believing what S does believe with justification.

5. Norms for judgment and belief

With my account of justification in place I shall now give that account some motivation. The general picture is that justified beliefs and judgments are those that approximate to knowledge in a certain way. Justification is a kind of near miss. Knowledge is the epistemic bull's-eye.

Justification occurs when an unexpected gust of wind nudges the arrow off-centre. Unjustified belief occurs when one's aim was poor or one fumbled the release. The key element of this view is that belief aims or ought to aim at knowledge; knowledge is thus the primary norm for belief.⁸

It is sometimes said that *truth* is a norm for belief. This is unarguable if this means that truth is a standard by which beliefs can be assessed. It will be one among a variety of standards of greater or lesser importance—we can also assess beliefs according to whether they are offensive, plausible, provable etc. But the truth norm for belief is sometimes held to be more important than any other norm. It is the norm, so it is said, that belief properly aims at, and in that sense constitutive of belief. A false belief, whether or not it is justified, fails to be what it ought. It must be admitted that even among those who concur that 'belief aims at truth' there is disagreement about what exactly this slogan means.⁹ My proposal is this. Cognitive faculties have essential functions, as do bodily organs and the like.¹⁰ The function of the liver is to filter toxic impurities out the blood. Likewise the function of the faculty of belief is to produce truth/knowledge (depending which you think the aim is). A little more precisely, the function of the faculty is to supply true contents (or knowledge) for the purposes of reasoning. Thus when engaged in practical deliberation whose outcome is an action, the function of belief is to supply inputs to that process, and furthermore its function is to supply true/known inputs, so that the output, the action, will be successful. This is analogous to the function of the blood, which is to supply oxygen to the muscles and other organs, so that they may carry out their functions properly. On this view it is the function of the faculty that provides its constitutive aim (this function/aim makes it the faculty that it is).

Now I have sketched what I hold is meant by 'the aim of belief is ...' we must address the question of whether it is truth or knowledge that is the aim. Like bodily functions, the cognitive faculties can fall short of their constitutive aim. The damaged liver may fail to cleanse the blood of impurities; we may form beliefs that are false. The question amounts to this, is the truth of a

⁸ The plausibility of my account depends in part on its being a component of a package that overall provides a more satisfactory explanation of various phenomena than its competitors. Nonetheless, it is worth noting that the account does not strictly depend upon the remarks in this section concerning the aim of belief.

⁹ Alternative proposals are given by, among others, Williams (1973), Velleman (2000), and Wedgwood (2002). Owens (2003) denies that belief has an aim.

¹⁰ For accounts of biological function see Wright (1973), Cummins (1975), Millikan (1989), Neander (1991), and Walsh (1996).

belief enough for it to have achieved its constitutive aim, for it to have fulfilled its function fully? Or is mere truth not alone sufficient for the function to be fulfilled? Certainly falsity is a good reason to give up a belief—its presence marks a lack of achievement on the part of the believer. So truth is *part* of the aim of belief. But, I claim, it is not enough. If truth were sufficient for the fulfilment of its constitutive aim then one could have no complaint with someone whose beliefs are true. But clearly this is not right. A belief that is merely accidentally true and lacks justification may rightly be criticized on this score (even a self-deceiver may be accidentally right). So, I suggest, justification is required as well. Are justification and truth jointly sufficient for the success of belief? It would be odd if the norm for belief were the conjunction of two independent factors. As Gettier-style examples show, there is nothing special about a belief that is both justified and true. The best explanation of all this data is, I propose, that *knowledge* is the aim of belief.¹¹ Since knowledge entails both justification and truth, it is clear both these are necessary. It is worth noting that Williamson argues that knowledge has an explanatory capacity (in explaining action) that true belief and even justified true belief do not (Williamson 2000, 63). If his argument is sound, it suggests also that it is better for an agent to have knowledge than true belief. In which case it is plausible, given the sketched account of ‘belief aims at ...’, that the aim of belief is knowledge rather than truth. (It is also suggestive that when deliberating about a course of action we say that we do so because we want to *know* what to do.)

In summary the proposal is this. Belief is a cognitive attitude that aims at knowledge. Anything short of knowledge is failure. But some failures are worse than others. And in particular 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.

That knowledge is the appropriate norm for judgment fits with Timothy Williamson’s proposed norm for assertion, which also is knowledge. If one is entitled to judge that *p* one should be entitled to assert that *p* to oneself. Since the latter entitlement entails knowledge, so then does the former. At the same time, what one is entitled to assert to others one is entitled to judge. Again, if the latter entitlement entails knowledge, then so does the former. Hence knowledge is the

¹¹ Williamson also asserts that knowledge is the aim of belief (Williamson 2000; 1, 48); c.f. below.

norm for judgment if and only if it is the norm for assertion. This last argument, *mutatis mutandis*, shows that the norms for judgment and assertion must be identical, whatever they are.

Jonathan Sutton argues that taking knowledge to be the norm for assertion means that justification cannot fall short of knowledge. If a belief were justified but not knowledge then S would be justified in asserting *p* even though S does not know *p*. But he takes Williamson's arguments to show that one is not justified in asserting *p* unless one knows *p*: "If beliefs so transmitted [by assertion] meet the primary standards governing good belief for both speaker and hearer—that is they are justified in an evaluative sense—and meet standards of permissible belief ... it would be mysterious if the assertions transmitting the beliefs failed to meet the standards governing good assertion."

This is true, but to be an argument for justified belief = knowledge it must trade on a equivocation concerning 'standards governing good belief/assertion'. First, there is a standard concerning the *aim* of belief and assertion and its fulfilment. The standard in both cases is knowledge. But there is also a standard for justified belief and justified assertion. That standard is less stringent. A case of practical action provides an appropriate analogy. S is attempting to save a drowning man by throwing him a lifebuoy. It is a good throw but a sudden gust of wind blows the lifebuoy out of reach. S's aim is not achieved; there is a stringent standard by which his action must be judged a failure. But that action is clearly justified. Furthermore, I think that it is in order to say in advance of the throw: 'S ought to throw the lifebuoy within reach of the drowning man'. If that is correct then S did not do what S ought to have done. Nonetheless, what S did was entirely justified. It makes sense to have different standards, to evaluate the action in terms of its goals, to evaluate an action in terms of ethically satisfactory outcomes, and to evaluate them as they reflect on the moral character of the agent. These will be different standards since some failures to reach a satisfactory outcome (the man drowns) do not reflect on the character of the agent. There is another possible world where an agent, S*, whose physical motions are just like those of S but where the drowning man is saved (there is no gust of wind). Justification is the standard by which we judge S and S* alike. But that is consistent with there being a standard of success which S*'s action met but S's did not. Furthermore, the latter is primary since (trivially) all actions aim at success. Sutton is thus mistaken in thinking that justification is the primary standard for belief—knowledge is the primary standard and justification is a derivative, secondary standard. We have seen why we have two standards. The

primary standard is the success the action aimed at—knowledge in the case of judgment and thus is typically external; the secondary standard is used to evaluate the internal component of the action, whether it was appropriate to the primary standard. For if it is, it is worth identifying, praising, and reinforcing. Corresponding to justified belief there is a justified assertion. (It may be more perspicuous to say ‘S was justified in believing *p*’ and ‘S was justified in asserting *p*’.) This is consistent with that belief and that assertion failing to be what they ought to be—a belief and an assertion backed up by knowledge. (And, consequently, Williamson’s arguments do *not* show (and are not intended to show) that an assertion is justified only if it is known.)

So far I have not distinguished between belief and judgment. The main reason for not doing so is that it would be odd if one could be justified in judging that *p* but not justified in believing that *p*, or vice-versa. But there do seem to be differences. On one view belief relates to action while judgment is closer to assertion. On another judgments are themselves actions, typically subject to the will, while sometimes believing seems to be something that simply happens to us. Related to both of these is the thought that judgment is a matter of intellectual assent to a proposition while belief carries with it degrees of conviction that are more visceral and which may conflict with the intellectual assent of judgment. Whether these differences affect the account of justification, and if so how, remains to be explored elsewhere. One source of complaint might be that the idea of a norm at which belief aims is inappropriate if what we believe is not subject to choice and control. To this last thought there are two responses. First, norms of a more or less teleological kind can be applied to objects and processes that are involuntary, such as the functioning of bodily organs. And a Darwinian story may easily be told about the function of the faculty of belief. Cognitive faculties that produce accidentally true beliefs are not functioning properly since they may well have produced a false beliefs. Moreover, there is a selective advantage in knowing over accidentally true believing, since knowledge is more robust in the face of misleading countervailing evidence.¹² Consequently a cognitive system that produces knowledge is likely to supersede a system that continuously gets lucky in producing true beliefs. Secondly, judgments and beliefs often are subject to the will in the sense that they are the intended outcomes of processes of reasoning we choose to engage in. It is true that many beliefs are not like this, basic perceptual beliefs for example, or the beliefs of the brainwashed or insane. Those who wish to pursue the thought that for such beliefs norms are inappropriate (the first reply notwithstanding),

¹² C.f. Williamson 2000; 75-80.

may wish to limit the account of justification (this or any other) to judgments and those beliefs that accompany judgments.

6. *An internalist intuition*

My approach is thoroughly externalist since it depends on an externalist concept of knowledge. Nonetheless it respects an intuition which, I believe, is part of what motivates the internalist. That an externalist account of justification can do so is an advantage.

Articulating epistemological internalism. Ralph Wedgwood claims the following:

... the following intuition seems compelling. Consider two possible worlds, w_1 and w_2 . In both worlds, you have exactly the same experiences, apparent memories, and intuitions, and in both worlds you go through exactly the same processes of reasoning, and form exactly the same beliefs. In this case, it seems, exactly the same beliefs are rational in both worlds, and exactly the same beliefs are irrational in both worlds. Now suppose that in w_1 you are bedevilled by an evil demon who ensures that many of your experiences are misleading, with the result that many of the beliefs that you hold in w_1 are false. In w_2 , on the other hand, almost all your experiences are veridical, with the result that almost all the beliefs that you hold in w_2 are true. Intuitively, this makes no difference at all. Exactly the same beliefs are rational and irrational in both worlds.¹³

Wedgwood then goes on to give an account of the thought processes by which beliefs that are rational are formed. That is the internalist approach. But the intuition he states can be used to support an externalist account of justified belief.

When Wedgwood articulates his understanding of this intuition he extends the supervenience basis for rational belief beyond just experiences, apparent memories, intuitions, and belief forming processes, for these clearly nowhere near exhaust the range of doxastic and epistemic

¹³ Wedgwood 2002; 349.

states let alone other mental states a subject may be in. Wedgwood says that rational belief supervenes on the above and all other ‘internal facts’. Internal facts include a subject’s non-factive mental states (and causal relations among such states).

This is where I part company with Wedgwood and internalists more generally. For I propose that the supervenience basis should include all mental states, factive as well as non-factive. S’s mental state M_p with content p is factive iff S has M_p entails p is true. Knowing is factive, since S knows p entails p . Williamson argues that knowing is a mental state (Williamson 1995). If this is right then knowing is a mental state that is neither include among nor supervenient on Wedgwood’s ‘internal states’. We should include factive mental states in our supervenience basis for the following reasons. First, what one knows is relevant to what one rationally believes, especially on Williamson’s equation of evidence with knowledge (Williamson 1997). If the statement of the intuition is not to be question-begging, then we should require that the subject be in all the same antecedent mental states in the two worlds. Secondly, Williamson argues that knowing is *prime*, that is, it cannot be factored into constituent parts: it is not the case that knowing is a composite of an internal mental component (such as believing) and an external component (which would including truth). If that is right, then there is no guarantee that there is a non-factive mental state that a subject is in whenever that subject is in a state of knowing. And so when we explain why a judgment is rational or justified by reference to a factive mental state (what someone saw, remembered, knew), there is not obviously any guarantee that there is a non-factive mental state that could do the same job. Thirdly, the restriction to non-factive mental states seems unnatural. Wedgwood’s intuition seems just right and straightforward as relates to all mental states. But when it is qualified by reference to the semi-technical idea of a non-factive mental state, it begins to look less like a ubiquitous intuition and more like a substantial thesis in need of a defence. Fourthly, it is not immediately clear that the restriction to non-factive mental states will in any case satisfy the internalist. For if there are factive mental states, there might, *prima facie*, be highly external non-factive mental states. Might there not be some states with content p that have entailments other than p but related to p (e.g. states that entail $\neg p$, or possibly p , or probably p)? For example ‘refuting p ’ looks like a mental state or act that entails $\neg p$. Refuting is not factive, but it is external in the same way that factive states are. If Wedgwood is to exclude these states also, then it seems likely he cannot do this simply by excluding certain easily identified kinds. Rather, an amount of contentious internalist theory will be required to say which mental states are to be employed in specifying the internalist ‘intuition’ (in which case its

status as independent evidence in favour of internalism is weakened). This reinforces the third point, that the internalist ‘intuition’ construed so as to refute externalism is no longer so intuitive, for the intuition so construed cannot even be formulated without substantive and debatable presuppositions.

As argued above, considering the justification of a belief, it is best to consider the process of belief-formation diachronically. If S sees normally while S* is a victim of an illusion, what makes S* and S equally justified is the fact that they started with the same mental states and then formed a judgment in the same way. The internalist intuition may therefore be expressed thus

(InIn) If in worlds w_1 and w_2 S has the same mental states and forms a judgment in the same way, then the judgment formed in w_1 is justified if and only if the judgment formed in w_2 is justified.

(JuJu) entails (InIn) and thus we may say that (JuJu) respects an important internalist intuition.¹⁴ I do not wish to suggest that (InIn) is the only way that the guiding internalist intuition can be expressed. Rather it is to show that there is one reasonably compelling way in which the intuition can be expressed that preserves its intuitive force while also being consistent with (and indeed explained by) my externalist account of justification. Since intuitions such as (InIn) are typically employed to motivate support for internalism against externalism, the fact that this externalist account can respect and explain it is an advantage of the account.

7. *The lottery paradox*

In this section I shall consider the relationship between justification and purely probabilistic thinking, as exemplified by inferences in a lottery case. It is sometimes supposed that the justificatory relationship between evidence and a belief is simply a matter of the former raising the probability of the latter’s truth. In this section I ask whether the relationship can be *simply* probability raising. I argue that if merely probabilistic reasoning cannot deliver knowledge it cannot deliver justification either. This, I argue, ought to be a feature of Williamson’s position,

¹⁴ The case of forgotten misleading evidence suggests that (InIn) is not correct. The natural accommodation is to replace ‘mental states’ in (InIn) by ‘mental history’. The corresponding reading of (JuJu) entails (InIn) thus modified.

and is certainly a direct consequence of my account. It is a view which has been argued for independently, e.g. by Dana Nelkin. I suggest that my account of justification provides a more attractive diagnosis of the parallel rejection of both knowledge and justification in lottery cases. I will return to sketch a positive account of evidential justification and the role of probability in the following section.

Concerning lottery cases, it is held by some philosophers that both of the following are true:¹⁵

- (a) a holder of a single ticket in a million ticket lottery does not know that he will lose;
- (b) the ticket holder is justified in believing that he will lose.

Above I motivated (JuJu) by reference to an approach that takes the aim of belief to be knowledge rather than truth:

- (c) belief aims at knowledge.

There appears to be a powerful tension here, verging on inconsistency, in the combination of (a), (b), and (c).¹⁶ If (c) is true then the justified belief in (b) has an aim which according to (a) it cannot achieve. Can the belief really be justified if its aim cannot possibly be achieved? The tension is exacerbated when we further assume that the lottery ticket holder is someone who has reflected on lotteries and knows that (a) is true. Can the lottery ticket holder be justified in his belief when he *knows* that it is impossible for the belief to achieve its aim? It seems not. If one knows that the very point and aim of a belief is impossible one ought to desist from believing in

¹⁵ E.g. Lewis 1999; 421. I believe that the same may be true of Williamson. See for example Williamson 2000; 200-201; in this case, however, Williamson employs not a lottery but draws of balls from a bag with replacement. Arguably such a case could have a different structure from a lottery, one which would allow for justification on the grounds I articulate below, concerning the possibility of a causal link. Williamson does write “I reason probabilistically, and form a justified belief ...” which does suggest that the reasoning is no different from a lottery case. Williamson does discuss lotteries in connection with assertion (2000; 238-269) and (1996). Williamson uses the fact that, intuitively, one should not assert “Your ticket has lost” in an ordinary lottery case to motivate a rejection of rules for assertion that fall short of his proposed knowledge rule (cf. above). However, if Williamson did not think that such cases yield justified belief, then the intuition that the assertion is inappropriate would not rule out (as constitutive of assertion) the rule that one should assert *p* only if one is justified in believing that *p*. At (2000; 259) Williamson makes it clear that he thinks that a lottery case can generate reasonable belief on merely probabilistic grounds.

¹⁶ In my view Williamson is committed to all three (see footnote 13) and so is guilty of this near inconsistency. In discussion Williamson has suggested that ‘full-on belief’ (the highest degree of belief) is not justified, while lesser degrees of belief may be. Thus (b) is false for full-on belief or judgment and (c) would be false for less than full-on beliefs. It does seem far less acceptable to say that the ticket holder is justified in *judging* that the ticket will lose than to say that he is justified in *believing* it will lose, and this may relate to the fact that we are happy with degrees of belief but not degrees of judgment. With this in mind, we may restrict our attention to (a), (b), and (c) understood with respect to full-on belief/judgment.

that case (and perhaps chose instead a belief that can be known, e.g. that the ticket is highly unlikely to win). If one nevertheless persists in one's belief it cannot be justified.

Thus it is a consequence of the position that I adopt here that one should either accept both knowledge and justification in lottery cases or deny both. Nothing I have said so far is strictly inconsistent with the first of these alternatives. Nonetheless, the view that there is no knowledge in lottery cases is widely held and so it is worth considering whether the implication that, at least for full-on belief and judgment, (b) is false. The independent reason for denying (b) is that holding it leads to the well-known paradox, whereby a subject is justified in believing every one of a set of propositions (for each ticket, the proposition asserts that it will lose) but is not justified in believing their conjunction, since this is known to be a contradiction.

The extensive literature on the lottery paradox suggests a variety of responses which it will not be appropriate to review here. Instead I note that Dana Nelkin also asserts that one's response to the justification version of the lottery paradox should match one's response to the knowledge version, while also arguing forcefully that the paradoxes require denying that there is either knowledge or justification.¹⁷ This linking would be of no more than aesthetic value unless there were also a connection in the diagnosis of the failure to know and the corresponding lack of justification. Nelkin's diagnosis in the case of knowledge is that the belief in question comes about as a result of a P-inference:

(P) p has a statistical probability of n [where n is a very close to 1] therefore p .¹⁸

This inference cannot on its own be knowledge generating (even when the conclusion is true) since it provides no appropriate link between the truth of the resulting belief and the fact that makes it true. Nor in the lottery case does any such link exist externally.¹⁹ An 'appropriate link' would be the sort of causal, explanatory, or reliable natural connection that externalists have focussed on post Gettier.

¹⁷ Nelkin 2000. To be precise, Nelkin talks of the rationality version of the lottery paradox rather than the justification version.

¹⁸ Nelkin 2000; 388 (Nelkin says 'where n is a very high number').

¹⁹ The fact that the P-inference as applied to a lottery case does not generate knowledge, the causal or reliable link being absent, does not mean that the same inference when supplemented by further premises concerning such a link, or made in a case where such a link does exist externally, cannot generate knowledge. There is no reason to suppose, pace Hawthorne (2004; 135), that Nelkin's diagnosis generalizes to rule out a lot of knowledge claims expressed

The problem, as Nelkin notes, with extending this to the justification/rationality version of the paradox is that justified beliefs (unlike knowledge) do not have to be true. A false but justified belief cannot be appropriately connected to the truth. Nelkin's solution is to internalise the appropriate connection condition on knowledge. For a belief to be rational it must be possible for the subject to suppose that there is some appropriate connection between one's belief and its object. ". . . [I]t must be possible," says Nelkin," for one to recognize one's commitment to the idea that one believes something *because* it is true."²⁰ Since in the case of subject who makes and understands a P-inference there can be no such connection the subject cannot properly represent there being such a connection.

As an explanation of why P-inferences do not yield justification this may do the trick but does so at the cost either of breaking the relationship between knowledge and justification or of making knowledge available only to sophisticated, reflective creatures. Nelkin faces the following dilemma. Either one takes the (fully) externalist view regarding knowledge, that an appropriate connection between truth and belief can be sufficient for knowledge. Or one takes a more internalist view, requiring that the subject additionally must be able to represent herself as standing in such a connection. If the externalist view is correct, then neither internal reflection nor the possibility of such reflection is required for knowledge. Given Nelkin's view that the latter is necessary for justification, then a subject who knows that *p*, but who, although appropriately connected to the truth, is unable to represent that fact to herself, is not justified in believing that *p*. Consequently, knowledge does not entail justification. On the other hand, Nelkin can retain the traditionally accepted relationship between knowledge and justification if she applies her internalist condition on justification to knowledge as well.²¹ That is, a subject could know some fact only if the subject were able to represent herself as standing in the appropriate connection to the (supposed) fact believed. That would exclude from the possibility of knowledge creatures (such as children and animals) whose reflective capacities are not well developed.

'categorically' rather than in probabilistic terms (c.f. Williamson 2000, 252).

²⁰ Nelkin 2000; 397.

²¹ I believe that Nelkin's own view is that justification as she conceives it is also a necessary condition for knowledge, along with the appropriate external connections (personal communication).

I suggest that my account of justification provides a better story concerning the link between the knowledge and justification versions of the lottery paradox without jeopardising the entailment of justification by knowledge. The latter is preserved as we have seen above: since a justified judgment is one that could have been knowledge; a judgment that actually is knowledge is *a fortiori* justified. I agree with Nelkin as regards the reason why (a) is true: a P-inference cannot deliver knowledge because of the lack of an appropriate connection between the truth of its conclusion and the corresponding fact. The account of justification I have given yields a straightforward explanation of why if (a) is true then (b) is false; that is, an explanation of why, given that the ticket holder cannot know the ticket will lose, he is not justified in judging it will lose. For the ticket holder to be justified there must be some world where by forming his judgment (that the ticket will not win) in the same way, he comes to know that he will not win. Here the way he forms his judgment is a matter making a P-inference. But there is no world in which making a P-inference will give him knowledge. Hence he is not justified in judging that his ticket will not win.

8. *Support, warrant, and blamelessness*

Jim Pryor usefully distinguishes three epistemic notions (or families of notions) that have been discussed in connection with justification. In this section I shall ask how these relate to my account. In the first place we have the idea that evidence may support a belief. As explained by Pryor:

If you believe that Bing is a male canary, and that 9 out of 10 male canaries can sing, and you have no other relevant evidence, then regardless of what you're *inclined* to believe on the basis of that evidence, there's an intuitive sense in which your evidence *supports* the conclusion that Bing can sing. Similarly, if you have a visual experience as of something yellow, and no countervailing evidence, what your evidence supports is the belief that there is something yellow in front of you.²²

Pryor calls this *justificatory support*. The second notion is that of *warrant* in Plantinga's sense—whatever it is that is needed to raise true belief to knowledge. The third notion is that of

²² Pryor 2001; 112.

epistemic blamelessness—the idea is that certain failures to know or to have true belief are not as bad as others, since those failures may be due to factors beyond the limits of one’s epistemic responsibility.

I suspect that there might well be more than one relationship between evidence and belief that would fall under the heading of justificatory support, as Pryor has presented it. For example, the evidence might simply raise the probability of the belief’s being true. But, as we have seen in the lottery case, that relationship can hold without giving justification to the belief. Furthermore, it is often the case that one thinks that a belief has a high probability on the grounds that the belief has a lot of evidential support (rather than vice-versa). These considerations each lead towards the view that there is a sense of evidential or justificatory support that is not entirely derived from probabilistic considerations (and hence that a merely Bayesian account of justificatory/confirmatory support is incomplete). A piece of evidence may add to one’s justification for p without itself alone making belief in p justified. This suggests that e gives p justificatory support if e is a non-redundant component of some set of possible facts E^* such that a belief in p based on knowledge of E^* is justified. This won’t quite do, since there might also be a set of facts E^\dagger such that e is a non-redundant part of E^\dagger and E^\dagger justifies $\neg p$. So one would also want evidence that suggests that E^* is more likely than E^\dagger . This gives a possible starting point for constructing an account of justificatory support founded on the ordinary concept of justification.

The notion of warrant is also necessarily more restricted than that of justification, since warrant is *supposed* to exclude Gettier cases whereas justification does not.²³ However, it is unclear that there really is any useful notion of warrant that does this. If warrant entailed the truth of p , warrant would not be distinct from knowledge. So the warrant for p is consistent with the falsity of p ; correspondingly there is a danger of allowing in Gettier cases of true warranted belief that are not cases of knowledge.²⁴ The only way to exclude cases arising from Gettier-style constructions is to deny that warrant is transmissible even over simple known entailments. Unless we commit ourselves to implausibly strong denials of closure, we must accept that knowledge is indeed transmissible over simple known entailments. Since knowledge but not

²³ C.f. Plantinga 1993. Plantinga allows for degrees of warrant. I am therefore here discussing that quantity of warrant that is supposed to be sufficient to raise true belief to knowledge.

²⁴ This point is emphatically discussed in Zagzebski 1994.

warrant is transmissible, it is difficult to see how warrant could in any sense be a component of knowledge. Moreover the Gettier-style constructions are not the only way of generating a counter-example to the equation knowledge = justified true belief. The important thing is that the belief in question should be true only accidentally. If warrant is consistent with falsehood then it should be consistent with accidental truth. If accidental truth were somehow excluded then the only cases of warrant would be knowledge and unavoidably false belief. In my view there is no good reason to think that there should be anything that fills the gap between true belief and knowledge. I explained that Gettier cases arose because justified beliefs are those that fail to be knowledge because of mentally extraneous factors. Such factors are typically consistent with accidentally true belief. To avoid Gettier cases for warrant without knowledge would require excluding any such extraneous factors consistent with the truth of the belief. But there is no reason to think that such factors form a unified set.

Since nothing can satisfy this notion of warrant, justification cannot be expected to. What is wrong with it is the thought that knowledge can be built out of truth plus something else. Nonetheless there is something right about this notion of warrant, viz. that we can explain justification by reference to its relationship to knowledge, and that its exactly what my account does.

We may expect epistemic blamelessness to have a wider extension than justification. While a killing in self-defence may be justified, a killing by someone who has been brainwashed or who is insane is not justified but may be blameless. Similarly someone who has from an early age been taught a set of faulty inference procedures or who suffers from a mental illness may have beliefs that are not justified; nevertheless, we may still not want to blame the epistemic shortcomings of the beliefs on the unfortunate individual concerned. My account of justification respects this difference. In the cases mentioned it is not possible to be in the same mental states as the brainwashed or insane persons and form beliefs in the way that they do and thereby acquire knowledge. Hence on my account their beliefs are not justified. At the same time we may see why intuitions concerning epistemic deontology have their bite. My account supplies knowledge as a norm for belief and this norm generates an epistemic obligation, roughly that one order one's mental life so as to achieve the goal of belief. So if, as in a case of unjustified belief, the explanation for one's failure to achieve the goal is internal to one's mental sates, there is a prima facie case that one has culpably failed to meet the obligation to order one's mental life properly.

But that case may be defeated in certain instances where the individual is not *capable* of ordering their mental life. Brainwashing and mental illness supply obvious cases. Epistemic culpability extends only as far as one's control over one's mental life. But the limits of that control may not be clear, just as they are not always clear in questions concerning moral responsibility for one's behaviour.

If blamelessness is not the right approach, then praiseworthiness might well be. This I think is the proper function of the concept of justification, to provide a certain kind of positive evaluation. The praiseworthiness of an action or belief is related to the praiseworthiness of the agent or subject. Praise is used to reinforce their dispositions and ways of doing things. Thus we may praise an unsuccessful action if it is nonetheless carried out in the correct way or with the right motivation. This, I have emphasized, is the evaluative role of the concept of justification. Bad luck may prevent our actions from having the optimal outcome or our judgments from amounting to knowledge. That it is bad *luck* means there is room for an important kind of positive evaluation, one that notes that were it not for that bad luck, what one did would have led to the optimal outcome, which in the epistemic case is knowledge.²⁵

9. *Degrees of justification*

One problem with the view that knowledge is a subspecies of justified true belief has attracted relatively little comment. It is the complaint that since justification comes in degrees, it would seem to follow from that account that knowledge should come in corresponding degrees. But it does not. Although we talk of certain facts being well-known that refers not to their strength of justification but to the number of knowers. For sure, we can create degrees of knowledge. We can say that the existence of electrons is more strongly known than the existence of quarks, but it comes less naturally than saying that belief in electrons is more strongly justified than belief in quarks. So it seems that justification naturally comes in degrees but not knowledge.

This would be an oddity if knowledge were justified true belief. But on the view of justification I am presenting there is no tension of this kind and indeed we would expect justification to come in

²⁵ Zagzebski (1996, 284) notes that in Gettier cases it is bad luck that prevents the subject from knowing. (It is bad luck that one is in fake barn country, even if one is looking at a (locally) rare real barn; it is bad luck that Jones has just sold his Ford; it is bad luck that the sheep in the field is hiding behind a rock.)

degrees but not knowledge. In general terms belief aims at knowledge and a belief that fails to achieve knowledge may be relatively nearer or further from its goal. Degrees of justification relate to the distance from the goal of knowledge. However, the answer cannot simply be a matter of the distance in Lewisian terms between the actual world and the nearest world in which one knows. This is because that distance may be a result of mentally extraneous factors alone. Consider two actual beliefs that are justified but which do not yield knowledge. It will be only mentally extraneous factors that differentiate the actual world from worlds in which the subjects know. These mentally extraneous reasons for failure might be extreme in one case but less extreme in the other. If that were enough to make one belief more justified than the other, then the difference in degree of justification would depend on mentally extraneous factors alone. That would be in tension with a central thesis of this paper that whether or not a belief is justified supervenes on mental factors alone.²⁶ An account of degree of justification should likewise make degree of justification supervenient on the mental alone.²⁷

A typical use of the idea of degree of justification is this. One might start with a promising hypothesis. It has some small but positive degree of justification. We adduce evidence in its favour, and its degree of justification grows. We might improve or add to our belief forming methods, thereby also adding to the degree of justification (should the improved methods continue to support the hypothesis). Of course, we may be cautious believers and not actually believe the hypothesis until it has a high degree of justification. Nonetheless, the degree of justification attaching to the hypothesis is presumably equal to the degree of justification a belief in the hypothesis would have, under the same circumstances. These changes in evidence or belief forming process are changes to mental states and processes, or involve such changes. The increased justification comes from the contribution that these mental changes make towards bringing the hypothesis (or the corresponding belief) closer to knowledge. More generally, S*'s belief is more justified than S's belief if the former is closer to knowledge due to a difference in mental states and belief-forming processes between S* and S.

The latter implies that the maximum degree of justification is possessed by any piece of knowledge and therefore by any justified belief. And we don't typically use degrees of justification when a belief has reached the threshold for knowledge. Even so, one might think

²⁶ The mental factors include, of course, externalist mental states and processes.

that there is a naturally extension of the idea to apply to those beliefs that amount to knowledge. For example, a belief might increase in justification if we acquire ever more evidence, even once it constitutes knowledge. That thought is easily accommodated in the present framework. To use an analogy, an arrow that hits the bull's-eye might have been blown off target had there been a gust of wind, but the stronger the force with which the arrow is released, the stronger the wind required to cause a miss. Under the circumstances a quantity of evidence may be enough to generate knowledge, but had the circumstances been different that evidence may not have generated knowledge, e.g. if there were additional misleading countervailing evidence, or if some belief-forming process had been unreliable. If the initial quantity of evidence had been greater or had lead to the relevant belief via a number of different and independent processes, then the knowledge generated is more robust and less likely to be undermined by slight differences in evidence or belief forming process. So cases of knowledge may be graded according to how close they are to cases of not knowing. The further away they are (e.g. by being supported by more evidence) the better justified they are.

Thus beliefs of all kinds (including those that constitute knowledge) may be graded according to their degree of justification. But since knowledge itself is a matter of hitting the bull's eye, it is an all-or-nothing affair, and consequently it is not typically appropriate to attach degrees to the concept of knowing.

10. Conclusion

Above I mentioned that Jonathan Sutton takes an approach broadly similar to mine but reaches the conclusion that a subject is justified in believing p only if they know p . For example, he argues that the reason why we should not ascribe justification in the lottery case is that the subject does not know. I have proposed instead that there is no justification because the subject *could not* know in a lottery-type case (not that they just don't *actually* know). This is a better explanation in that it is a consequence of an account of justification that both ascribes to justification a distinctive role while preserving important intuitions about justification. If justified belief just is knowing, then one might wonder what the distinctive function of the justification concept is. While Sutton's account doesn't provide this, my account makes precise

²⁷ I am grateful to Carlton Gibson for emphasizing this point.

the very natural idea that saying that a belief is justified is a way of saying something positive about the subject even if they fail to know—it is to say that they are just like (in the internal aspects of belief-formation) a possible someone who does know. This satisfies the intuition that knowing entails justification but not vice-versa. It furthermore yields the intuition that motivates internalism, that justification supervenes on the internal aspects of belief-formation. Sutton’s account does not—for him justification is a strongly externalist as knowledge. Two subjects can form beliefs in the same way (from the internal perspective) but one know and the other not know—Sutton’s view has it that the one is justified but the other not.

Traditional approaches to justification have raised more questions than they have answered. In particular two fundamental problems have remained as far as ever from solution: What is it that gives a belief its justification? And what is the relationship between justification and knowledge? My view is not that these problems are too difficult or that we have not been trying hard enough to find a solution. It is, rather, that the traditional approaches have simply grasped the stick at the wrong end. Instead of constructing an account of justification, whether foundationalist or coherentist, and then building on top of that to provide an analysis of knowledge, we should start from the concept of knowledge and work our way down. Perhaps to devotees of the traditional approaches this will look like cheating, helicoptering onto the top of a mountain and walking down rather than slogging one’s way up. But this is philosophy not mountaineering, and the ease of the downward route should itself be evidence of its veracity.²⁸

References

- Cummins, R. 1975 “Functional Analysis” *Journal of Philosophy* 72, pp. 741-765.
Hawthorne, J. 2004 *Knowledge and Lotteries* Oxford: Oxford University Press.

²⁸ I would like to thank Dana Nelkin, Finn Spicer, Ralph Wedgwood, and Timothy Williamson for their helpful

- Lewis, D. 1986 “Counterfactual Dependence and Time’s Arrow”, in *Philosophical Papers Volume II* Oxford: Oxford University Press, pp.32-66.
- Lewis, D. 1999 “Elusive Knowledge”, in *Papers in Metaphysics and Epistemology* Cambridge: Cambridge University Press, pp. 418-445.
acknowledgments obscured for anonymity
- Neander, K. 1991 “Function as Selected Effects: A Conceptual Analyst’s Defense” *Philosophy of Science* **58**, pp.168-184.
- Nelkin, D. 2000 “The Lottery Paradox, Knowledge, and Rationality” *Philosophical Review* **109**, pp. 373-409.
- Plantinga, A. 1993 *Warrant and Proper Function* Oxford: Oxford University Press.
- Owens, D. 2003 “Does Belief Have an Aim?” *Philosophical Studies* **115**, pp.283-305.
- Pryor, J. 2001 “Highlights of Recent Epistemology” *British Journal for the Philosophy of Science* **52**, pp.95-124.
- Sutton, J. forthcoming 2005 “Stick to What You Know” *Noûs*.
- Walsh, D. 1996 “Fitness and Function” *British Journal for the Philosophy of Science* **47**, pp.553-574.
- Wedgwood, R. 2002a “Internalism Explained” *Philosophy and Phenomenological Research* **65**, pp. 349-369.
- Wedgwood, R. 2002b “The Aim of Belief” *Philosophical Perspectives* **16** *Language and Mind*, pp. 267-297.
- Williams, B. 1973 “Deciding to Believe” in *Problems of the Self* Cambridge: Cambridge University Press, pp. 136-151.
- Williamson, T. 1995 “Is Knowing a State of Mind?” *Mind* **104**, pp. 533-65.
- Williamson, T. 1996 “Knowing and Asserting” *Philosophical Review* **105**, pp. 489-523.
- Williamson, T. 1997 “Knowledge as Evidence” *Mind* **106**, pp. 717-741.
- Williamson, T. 2000 *Knowledge and its Limits* Oxford: Oxford University Press.
- Wright, C. 1991 “Scepticism and Dreaming: Imploding the Demon” *Mind* **100**, pp. 87-116.
- Wright, L. 1973 “Functions” *Philosophical Review* **82**, pp. 139-168.
- Velleman, J D. 2000 “On the Aim of Belief” in *The Possibility of Practical Reason* Oxford: Oxford University Press, pp. 244-281.

Zagzebski, L. 1994 "The Inescapability of Gettier Problems" *Philosophical Quarterly* 44, pp. 65-73.

Zagzebski, L. 1996 *Virtues of the Mind* Cambridge: Cambridge University Press.