



Kelp, C. (2011) Not without justification. *Dialectica*, 65(4), pp. 581-595.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/140943/>

Deposited on: 15 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Not Without Justification

CHRISTOPH KELP*

ABSTRACT

In this paper I take issue with Jonathan Sutton's attempt at defending the thesis that knowledge is justified belief. I argue, first, that the arguments he adduces in support of it fail. Second, I provide independent reason to believe that knowledge and justified belief come apart.

1 Introduction

In a number of recent pieces, Jonathan Sutton (2005, 2007) has defended a novel approach to epistemology. The central thesis Sutton aims to establish and that I will take issue with in this paper is:

The identity thesis. Justified belief is knowledge.

At first glance, the identity thesis is rather unattractive to say the least. To begin with, it may be worth pointing out just how strong a claim Sutton aims to establish here. The thought is not that simply that there is *a* concept of justified belief that is coextensive with the concept of knowledge. Rather, the thought is that *any* concept of justified belief that is of core epistemological interest is coextensive with that of knowledge. Here is Sutton:

I aim to show, first, that we do not and maybe cannot have *a* serviceable notion of justification that is distinct from knowledge, and, second, that we do not need one—we can get by better in epistemology without one. (Sutton 2007, 2-3)

In light of the overall argument, it is no objection to my view that some particular epistemological issue that I have perhaps

*Institute of Philosophy, University of Leuven, Kardinaal Mercierplein 2, BE-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

not considered fully or at all might require a knowledge-independent notion of justified belief, since the argument, I contend, establishes that there is *no* such notion to be had. (Sutton 2007, 40)

Notice, next, that the identity thesis also is highly revisionist: traditionally, justified belief is taken to be necessary but insufficient for knowledge. What's more, Sutton also intends to reverse the traditional direction of analysis: instead of defining the concept of knowledge in terms of an independently understood concept of justified belief, Sutton holds that the concept of justified belief is to be defined in terms of an independently understood concept of knowledge.

Perhaps most importantly in relation to unattractiveness, however, is the fact that the identity thesis is intuitively highly implausible. After all, consider the following, intuitively highly plausible claims:

- Perceptual beliefs about dry, middle-sized goods that are in plain view formed by normal adult human beings in good lighting conditions are justified. For instance, suppose *A* is driving alongside a road in the countryside and sees a barn in the field on the right. *A*'s belief that he is currently looking at a barn is justified.
- Competent scientists beliefs in the best-confirmed scientific theories at the time are justified. For instance, Priestly was justified in believing that all flammable materials contain phlogiston.
- Some justified beliefs are false. For instance, suppose a spec of dust on the printing plate of the newspaper I follow and know to be highly reliable causes the print to look like '29' instead of '28'. My false belief that the upcoming election will be on the 29th, say, is nonetheless justified.

The intuitive plausibility of these claims constitutes evidence for their truth. If any one of these claims comes out true, then there must also be a concept of justified belief that is not coextensive with the concept of knowledge.¹

¹To see why the first claim supports the possibility of justified belief (in the relevant sense) that falls short of knowledge, suppose that *A* is looking at the only real barn in a field otherwise full of barn façades. The intuition that his belief is justified remains unscathed. However, intuitively, his belief does not count as knowledge—the case now is a run-of-the-mill Gettier case.

Thus, the identity thesis is not only a very strong thesis, but it also turns tradition on its head and is highly counterintuitive. As a result, in aiming to defend it, Sutton shoulders a heavy burden of proof. In the following section I will turn to Sutton's attempt at discharging it.

2 Sutton's arguments

Sutton adduces four arguments in support of the identity thesis. He is very clear that one of these arguments—the “modesty argument”, which is intended to establish that one cannot justifiably believe propositions one knows one doesn't know—by itself won't establish the identity thesis. For that reason I will, for the purposes of this paper at least, simply grant Sutton that it works in the way envisaged. As opposed to that, the other three arguments are intended to directly and independently establish the identity thesis. In what follows I will look at each of them and ask whether it does its job.

2.1 The assertion argument

Sutton's first argument for the identity thesis proceeds along the following lines. Sutton agrees with Williamson and others that knowledge is the norm of assertion. He then goes on to note:

One of the main goals of making assertions, if not *the* main goal, is to transmit beliefs from one thinker to another. If the beliefs so transmitted [are justified], it would be mysterious if the assertions transmitting the beliefs failed to meet the standards governing good assertion. On the contrary, the assertions in question have to meet the standards governing good assertion impeccably since they transmit impeccable beliefs. It is not, however, the knowledge rule that is at fault; the arguments of Williamson and others for that rule are good ones. It is our initial supposition that was at fault. There are no justified true beliefs falling short of knowledge. (Sutton 2007, 46)

Here is my explication of Sutton's argument:

AA1. The goal of assertion consists in the transmission of belief.

AA2. If the goal of assertion consists in the transmission of belief, then any belief that satisfies the epistemic norm of belief also satisfies the epistemic norm of assertion.²

²Sutton claims that any justified belief, if transmitted across assertion, satisfies “the standards for good assertion”. This claim appears to be too strong. For in-

AA3. Justification is the epistemic norm of belief.

AA4. Hence, any justified belief satisfies the epistemic norm of assertion. [AA1 – AA3]

AA5. Knowledge is the epistemic norm of assertion.

AA6. Hence, any justified belief also qualifies as knowledge.³ [AA5, AA6]

Should we be moved by the assertion argument? As I am about to argue, the answer here is “no”. To begin with, notice that at least some of the premises at issue in the assertion argument are themselves substantive philosophical theses and as such hardly self-evident. As a result, while the assertion argument establishes that we cannot have all of its premises and the negation of its conclusion, it is far from clear that it is the negation of the conclusion that must go. On the contrary, in view of the great *prima facie* implausibility of the conclusion, one might initially be inclined to think that one of the premises must be false. In order to get the assertion argument off the ground, Sutton must provide convincing support for at least those premises that are philosophically substantive. And this will have to include some good argument that the relevant theses are preferable to at least the most obvious alternatives on the philosophical market. However, as I will show in due course, Sutton fails to deliver here. If my arguments are successful, it follows that there is no reason to think that we must buy the conclusion of the assertion argument rather than give up one of its premises.

Let’s start with AA6, the knowledge norm of assertion. The obvious alternative to AA6 is what Sutton calls the “J rule” according to which justified belief is the epistemic norm of assertion. If the J rule rather than the knowledge rule governs assertion, the assertion

stance, there may be politeness standards governing assertion that an assertion, even one that transmits the relevant justified belief, may fail to satisfy. For that reason I made the restriction to epistemic standards on behalf of Sutton here.

³It may be worth noting that, even if sound, the assertion argument in its present form does not establish the identity thesis. After all, the identity thesis claims that justified belief entails knowledge and that knowledge entails justified belief. Even if successful, the assertion argument in its present form establishes only the first entailment, i.e. the entailment from justified belief to knowledge. Notice, however, that the other entailment, from knowledge to justified belief, is widely regarded as uncontroversial among epistemologists. So, even if the assertion argument in its present form does not establish the identity thesis, if successful, it does establish the controversial half of the argument for the identity thesis.

argument delivers the unexceptional result that any justified belief is justified rather than the surprising result that any justified belief must also qualify as knowledge. Thus, if the J rule rather than the knowledge rule governs assertion, the assertion argument will not go through.

In order to make the assertion argument cogent, Sutton needs to provide reason to think that the J rule is not a viable alternative to the knowledge rule. In order to achieve this, Sutton asks what the concept of justified belief would have to look like in order to explain the phenomena that, according to him, motivate the knowledge rule. He distinguishes between two kinds of non-knowledgeable belief that traditionalists about justified belief countenance: intuitively justified beliefs that are known to fall short of knowledge (“known unknown beliefs” as Sutton calls them) and intuitively justified beliefs that are not known to fall short of knowledge (“unknown unknown beliefs” as Sutton calls them). He then defines the concept of “*U*-justification” as follows: A belief is *U*-justified if and only if either it is a known belief or an unknown unknown belief. (Sutton 2007, 63) Now, Sutton concedes to the traditionalist that if the concept of justification at issue in the J rule is *U*-justification, then J rule explains the relevant phenomena just as well as the knowledge rule does. However, he also maintains that the knowledge rule retains an edge over the J rule so understood:

The J rule is, in fact, more complex than the knowledge rule ... *U*-justification is more complex than knowledge since the concept of *U*-justification is parasitic upon that of knowledge. A *U*-justified belief is a belief that constitutes knowledge or a belief that would have done so but for bad luck in the believer’s external environment. At the very least, the burden of proof is shifted back to the proponent of the J rule. He must show that *U*-justification is in fact at least as primitive a concept as the concept of knowledge rather than one that we understand as a disjunction of knowledge and would-be knowledge.⁴ (Sutton 2007, 63)

⁴Sutton equates “unknown unknown beliefs” and “would-be knowledge”, i.e. beliefs that would have qualified as knowledge “but for bad luck in the believer’s external environment”. Notice also that Sutton takes his distinction between “known unknown beliefs” and “unknown unknown beliefs” to partition the set of beliefs that, according to the traditionalist, can be justified whilst not qualifying as knowledge. Given that it does, however, the equation of “unknown unknown beliefs” and “would-be knowledge” is false. To see this, suppose one does not have sufficient evidence to come to know whether one’s evidence for *P* is strong

I must confess that I don't find Sutton's argument very convincing. I take it that no one in their right mind would deny that the concept of *U*-justification is more complex than the concept of knowledge and no one would attempt to show that it isn't. At the same time, no traditionalist I know has defined the concept of justification in the way Sutton defines the concept of *U*-justification. Of course, Sutton can define the concept of *U*-justification in whatever way he likes. He can also use the concept so defined to specify what the traditionalists' concept of justification would have to look like if it were to successfully explain certain phenomena. However, he cannot hope to rest an argument against traditionalists on the observation that the concept of *U*-justification that he defined and that no traditionalist holds is more complex than the concept of knowledge. What he needs to do instead is to show that there is no concept of justification that approximates the extension of the concept of *U*-justification closely enough to explain the phenomena. Moreover, Sutton cannot even shift the burden of proof back on the shoulders of traditionalists in this way. In order to shift the burden of proof to the traditionalist, he needs to show at least that no existing traditionalist concept of justification does the trick. However, he doesn't do this.

What's even worse, there is little hope that Sutton will be able to fill the gap in his argument. To see this, notice first that Sutton (2007, ch. 2.5.2) defines the concept of knowledge not reductively, i.e. by giving a small set of individually necessary and jointly sufficient conditions, but implicitly, i.e. through a set of platitudes, of common sense beliefs about knowledge. Now, if Sutton is entitled to define the concept of knowledge implicitly, then surely there is nothing wrong with defining the concept of justified belief implicitly as well. Notice, however, that we can now help ourselves to Sutton's (2007, ch. 2.1.3) "modesty argument"—which, recall, aims to show that there cannot be any "known unknown beliefs"—to argue that the implicitly defined concept of justified belief is coextensive with Sutton's concept

enough to give one knowledge that *P*. Suppose, in fact, it isn't. By the lights of the traditionalist, one's evidence may still be strong enough to justify belief in *P*. So, suppose it does. If one believes *P* on the basis of this evidence, one has an "unknown unknown belief" that isn't "would-be knowledge". After all, it is not bad luck in the environment that prevents one from knowing but the quality of one's evidence. Arguably, some of Sutton's arguments depend on both the partition of traditionally justified belief and the equation of "unknown unknown beliefs" and "would-be knowledge" and may therefore be problematic. However, I will not pursue this line against Sutton here.

of *U*-justification, in which case, by Sutton's own lights, the *J* rule will explain the relevant phenomena as well as the knowledge rule.

What these considerations show is that Sutton fails to provide convincing reason to prefer the knowledge rule to the *J* rule and that, given his further theoretical commitments and arguments, there is reason to think that any attempt to do so is bound to fail. Hence, the assertion argument, in Sutton's mouth at least, is bound to fail. Yet, one might wonder whether, independently of Sutton's commitments and arguments, there is sufficient reason to think that the knowledge rule is true. If there is, the assertion argument may still be successful.

(Un)fortunately, there is reason for pessimism here as well. The debate over the correct norm of assertion is ongoing. Apart from the knowledge rule and the *J* rule, there are a number of further competitors in the race. For instance, Kent Bach (2008) has argued for a belief rule, Matthew Weiner (2005) for a truth rule and Jim Stone (2007) for a context-sensitive rule of assertion. Defenders of the various norms have not only given alternative explanations of the phenomena that are to motivate the knowledge rule, but they have also pointed to further phenomena which they claim favour their account over the knowledge rule (see e.g. Lackey (2008), Weiner (2005)) or have given further theoretical reason why their account is preferable to the knowledge rule (see e.g. Douven (2006)). As a result, the prospects for an argument from the knowledge rule to the identity thesis are at present dim.

But perhaps the assertion argument is best understood as an argument that addresses sympathisers with the knowledge rule. In other words, what the argument shows is that if one is tempted by the knowledge rule, one must also buy the identity thesis. The assertion argument does not even establish this much. As I am about to argue, Sutton fails to make a convincing case not only for the knowledge rule of assertion but also for the thesis that the goal of assertion consists in the transmission of belief.

To begin with, recall that in order to get the assertion argument off the ground, Sutton must at the very least show that the obvious alternatives to the philosophically substantive premises are false. Regarding AA1, Sutton's thesis that the goal of assertion consists in the transmission of belief, the obvious alternative is that the goal of assertion consists in the transmission of *knowledge*.

Sutton considers this alternative but claims that it won't block the assertion argument. Here is how he puts the objection and how he aims to respond to it:

It is not one of the goals of assertion to transmit belief; in light of the knowledge rule, it is clearly to transmit knowledge. Indeed, one might go further and claim that [a hearer, *H*] cannot himself acquire a warranted belief from [a speaker, *S*] unless [*S*] expresses knowledge. No warranted false belief can transmit its warrant to another via testimony. If that is so, then the belief that [*H*] acquires will be as defective as [*S*'s] assertion, and there is harmony between the standards of belief and assertion.

These defenses do not remove our initial puzzle, however. If a belief can be warranted despite being false, it is bizarre that an assertion conveying such a belief is unwarranted. Why do the main goals of assertion fail to encompass the conveying of such beliefs? Why is an assertion not doing its job when it conveys such a belief? It would add to our bafflement if [*H*] could not acquire a warranted belief from [*S*'s] expression of a warranted false belief. Why would assertion not act as a conduit for warranted belief as well as knowledge proper? (Sutton 2007, 47)

Once again, Sutton's argument here is unconvincing. To begin with, if the goal of assertion consists in the transmission of knowledge, it is not at all bizarre how assertions of justified beliefs that fall short of knowledge can be unwarranted: such assertions do not only fail to satisfy the epistemic norm of assertion, but they also do not and cannot reach their goal. In this situation, it does not strike me as at all bizarre that such assertions are unwarranted. On the contrary, it would be bizarre if these assertions were warranted.

Relatedly, once the goal of assertion is understood in terms of the transmission of knowledge, there is no longer reason to think that there must be "harmony" between the standards of belief and assertion. After all, Sutton's thesis that the goal of assertion consists in the transmission of belief motivated the thought that there should be harmony here in the first place. Accordingly, there is no need to rescue this harmony by maintaining that justified false beliefs cannot be transmitted to a hearer by an assertion of a justified false belief and any bafflement that might result from making such a move can be avoided. Sutton thus fails to provide convincing support for the thesis that goal of assertion consists in the transmission of belief even

vis-à-vis its most obvious competitor according to which it consists in the transmission of knowledge. As a result the assertion argument fails on this count as well.

2.2 The lottery argument

The lottery argument draws crucially on the so-called lottery paradox. The lottery paradox shows that the following three claims, which are individually highly plausible, are jointly inconsistent:

The Lockean thesis. If it is very likely that P , then one has justification to believe P .

The conjunction rule. If one has justification to believe P and one has justification to believe Q , then one has justification to believe both P and Q .

The contradiction thesis. One does not have justification to believe P if one knows P to be a contradiction.⁵

Here is how Sutton ventures to argue from the lottery paradox to the identity thesis:

[The lottery] argument builds on the work of Dana Nelkin (2000), who presents two versions of the lottery paradox, one concerning knowledge and the other justification, and argues that they should receive a solution that locates the flaw in each paradoxical argument in the same place (in fact, counterpart premises) ... Nelkin's uniform solution to the paradoxes involves denying that the belief that one's ticket will not win is justified. I will argue that Nelkin's explanation of why it is unjustified is incomplete at best. The best explanation, I will argue, is that justification is knowledge. (Sutton 2007, 48)

Here is my explication of the argument:

⁵To see why these three claims are inconsistent, notice that no matter how high we set the standards for satisfaction of the predicate "very likely", there will be some fair lottery with exactly one winner and sufficiently many tickets that it is very likely that each ticket will lose. So suppose that a ticket will very likely lose if the chances that it will lose are greater than 0.99999 and let l be a lottery one knows to have 1000000 tickets and exactly one winner. By the Lockean thesis, for each ticket in l , one has justification to believe that it will lose. By the conjunction rule one has justification to believe that all tickets in l will lose and, since one knows that there is exactly one winner, that exactly one ticket will win. By the contradiction thesis, one does not have justification to believe this. Thus the three claims are jointly inconsistent.

LA1. There are two versions of the lottery paradox: one concerns justification (henceforth also “the justification argument”), the other knowledge (henceforth also “the knowledge argument”).⁶

LA2. The two versions of the paradox afford a uniform solution that faults the Lockean thesis.

LA3. The identity thesis constitutes the best explanation why, in the justification argument, the Lockean thesis fails.

LA4. Hence, the identity thesis holds. [LA1 – LA3]

In order to assess Sutton’s argument, let’s first take a closer look at LA3, Sutton’s claim that the identity thesis provides the best explanation of why the Lockean thesis for justification (henceforth also “LJ”) fails. True, the identity thesis provides a very straightforward explanation of the failure of LJ, to wit, that it is identical to the obviously false Lockean thesis for knowledge.⁷ In and of itself, that does not mean that it is the best available explanation. In order to show this, Sutton must show at least that his explanation of why LJ fails is better than the alternatives that can be found in the literature. It is somewhat surprising how little Sutton does here. He only discusses one alternative solution to the paradox that denies the LJ, i.e. the one given by Nelkin. However, the literature witnesses a wide range of solutions of this type: Jonathan Adler (2002), Jake Chandler (2010), Igor Douven (2002), Simon Evnine (1999), Keith Lehrer (e.g. 1990), Mark Kaplan (e.g. 1996), John Pollock (e.g. 1990), Sharon Ryan (1996) and Robert Stalnaker (1984) all provide solutions to the justification version of paradox that deny LJ. Since Sutton provides no reason to believe that his explanation of why LJ fails is better than these alternative explanations, he fails to provide sufficient support for LA3.

Let’s now focus on LA2, Sutton’s claim that the two versions of the paradox afford a uniform solution that faults the Lockean thesis. Why think that the two versions of the paradox afford a uniform solution? Sutton tells us that “it is a clear desideratum of a solution

⁶The knowledge argument can be obtained from the above justification argument by replacing “has justification to believe” by “is in a position to know” and adding the clause that *P* be true to the antecedent of the conditional in the Lockean thesis.

⁷Consider: While it is intuitively plausible that you have *justification to believe* that some ticket will lose a lottery it is very likely to lose, it is clear that you aren’t in a position to *know* that it will lose such a lottery. As a result, the Lockean thesis for knowledge must be false.

to one version that it can be applied to the other.” (Sutton 2007, 50) However, upon reflection, it is not so clear that this is in fact so. To see why not, notice, as Sutton himself acknowledges, that, unlike the justification argument, the knowledge argument “is not really a paradox at all.” (Sutton 2007, 50) After all, a moment’s reflection reveals that the Lockean thesis for knowledge does not hold.

Given that the knowledge argument isn’t really a paradox, however, it is far from clear that a uniform solution to the two arguments is indeed desirable. To see this, consider first the closure paradox, which shows that the following three individually highly plausible claims are inconsistent:

SK1. I know that I am in Manhattan.

SK2. Knowledge is closed under known logical entailment: If one knows that P entails Q , then, if one knows that P , then one knows that Q . Since I know that my being in Manhattan entails that I am not a brain in a vat on Alpha Centauri: If I know that I am in Manhattan, I know that I am not a brain in a vat on Alpha Centauri.

SK3. I don’t know that I am not a brain in a vat on Alpha Centauri.

In the spirit of Sutton and Nelkin, let’s distinguish two versions of the closure paradox, one concerning knowledge (henceforth also “the closure argument”), the other concerning strangeness (henceforth also “the strangeness argument”). Here is the strangeness argument:

ST1. It is strange that I am in Manhattan.

ST2. Strangeness is closed under known logical entailment: If one knows that P entails Q , then, if it is strange that P , then it is strange that Q . Since I know that my being in Manhattan entails that I am not a brain in a vat on Alpha Centauri: If it is strange that I am in Manhattan, it is strange that I am not a brain in a vat on Alpha Centauri.

ST3. It is not strange that I am not a brain in a vat on Alpha Centauri.

Unlike the closure argument, the strangeness argument is not really a paradox: a moment’s reflection makes it obvious that strangeness

is not closed under known logical entailment. At the same time, any temptation to think that it is nonetheless desirable to give a uniform solution to the two arguments is surely misguided. No one in their right mind would claim that since the culprit in the strangeness argument is the closure principle, it is desirable to say that, in the closure argument, it is also the closure principle that is at fault. It thus becomes clear that, where there are two superficially similar arguments one of which is a paradox, whilst the other one isn't, there is no need to give a uniform solution to them.⁸

It turns out that, once again, Sutton fails to provide convincing support for two crucial premises of the argument. He fails to establish that the identity thesis does indeed provide the best explanation of why LJ fails. What's even worse, however, we have seen that there is excellent reason to believe that the uniformity desideratum, on which the argument crucially rests, is false. As a result, the lottery argument fares no better than the assertion argument.

2.3 The posterior evaluation argument

The last argument that Sutton adduces in support of the identity thesis takes the following shape:

If a belief that P is one that would be justified were one to form it, and it is in one's interest to have a belief in whether or not P , and one is capable of forming such a belief, then, in some intuitive sense, one should believe that P ... This claim is clearly true if we read 'would be justified' as 'would constitute knowledge'. I will argue that it is not true on any more expansive conception of [justification]. (Sutton 2007, 57)

Once again, here is my explication of the argument:

- PE1. If one has an interest in whether P and one is capable of forming a justified belief that P , then there is some intuitive sense in which one ought to form a belief that P .
- PE2. There is no intuitive sense in which one ought to form non-knowledgeable but (allegedly) justified beliefs that one has an interest in and is capable of forming.

⁸It is plausible that the uniformity desideratum holds where two arguments are indeed two manifestations of the *same* paradox. Notice that it is also possible that two arguments are both manifestations of paradoxes and appear to be manifestations of the same paradox but in fact aren't. In that case, there is no need to provide a uniform solution to them either.

PE3. Hence, there aren't any non-knowledgeable justified beliefs.
[PE1, PE2]

Let's grant Sutton that PE1 is indeed true. What about PE2? Here is how Sutton ventures to mount a case for it. He asks us to consider the following two cases:

Case 1. At t_i , one has an interest in the question whether one has a colleague who owns a Ford. One also has the opportunity to acquire a belief that one has a colleague who owns a Ford (call this proposition "*FORD*"): one has excellent evidence in favour of *FORD*. One's colleague *A* has shown one papers of a Ford in his name, a video tape that shows him winning a Ford in a well-known TV show, a number of highly trustworthy people have testified that he owns a Ford, he arrives at work in a Ford etc. What's more, at t_i , one has no reason to distrust any of this evidence.

Crucially, (a) one does not form a belief in *FORD* and (b) one is not in a position to come to know *FORD* because one's belief would have been gettierised: the evidence produced by *A* is misleading, while, at the same time, another colleague, *B*, owns a Ford.

Case 2. At t_i , one has an interest in the question whether one has a colleague who owns a Ford. One also has the opportunity to acquire a belief in *FORD*: one has excellent evidence in favour of *FORD*. One's colleague *A* has shown one papers of a Ford in his name, a video tape that shows him winning a Ford in a well-known TV show, a number of highly trustworthy people have testified that he owns a Ford, he arrives at work in a Ford etc. What's more, at t_i , one has no reason to distrust any of this evidence.

Crucially, (a) one does not form a belief in *FORD* and (b) one is in a position to know *FORD*.

Next, Sutton invites us to suppose that, at some later time, t_j , one has come to know all the relevant facts about one's epistemic position at t_i and asks how one would, intuitively, evaluate one's situation at t_i with respect to *FORD*. He makes the following suggestion:

There is an intuitive (but nonepistemic) sense in which [one] should have formed the belief [in *FORD*] because it would have

been true [this is what happens in Case 1], and in a different case [that is, in Case 2] there is a different intuitive sense (one partially epistemic) in which one should have formed that belief because it would have constituted knowledge. What we lack, I claim, is the intermediate case. When one could have formed an allegedly justified, true belief [in *FORD*], but one failed to do so, there is no intuitive sense in which one should have done so beyond the nonepistemic sense in which one should have done so because it would have been true. (Sutton 2007, 50)

Sutton thus agrees that in both cases there is an intuition that one should have formed the belief one could have formed. Crucially, however, Sutton claims that in Case 1 this intuition is grounded in the fact that one could have formed a *true* belief rather than in the fact that one could have formed a *justified* one. If this is correct, then it looks as though there is evidence for PE2.

Now the problem is that in Case 1, the belief one could have formed would not only have been true but also justified. So, crucially, Sutton needs to back up his claim that it is the fact that one's belief would have been true rather than that it would have been justified that grounds the relevant intuition in Case 1. Sutton ventures to achieve this by arguing that the intuition that one should have formed a belief survives if we alter the case so that the proposition remains true but one has "no evidence whatsoever" for *FORD*. Here is Sutton:

However, we can say exactly the same about a similar situation in which one has no evidence whatsoever [in favour of *FORD*], although in fact one has such a colleague. There is an intuitive sense in which one should have formed the belief [in *FORD*] despite the fact that such a belief would have been unjustified by anyone's lights. Such a belief would have been true. (Sutton 2007, 58)

Somewhat surprisingly (or perhaps not so surprisingly), Sutton rests content with making this statement *in abstracto*, i.e. without actually describing the relevant situation. So let's fill in the details for him:

Case 3. At t_i , one is interested in the question whether one has a colleague who owns a Ford, but one has no evidence whatsoever bearing on this question. One does have the opportunity to acquire a belief in *FORD*: a famous hypnotist offers to instil

it in one (or, alternatively, perhaps one almost forms the belief on the basis of wishful thinking).

Crucially, (a) one does not form a belief in *FORD*, while (b) one's belief in *FORD* would have been true.

Even knowing that *FORD* is in fact true, I simply do not have the intuition that, at t_i , I ought to have accepted the hypnotist's offer (or ought to have formed the belief on the basis of wishful thinking) because that would have given me a true belief in *FORD*. What, intuitively, I should have done instead is suspend judgement on the question. Moreover, I am also fairly optimistic that I am not the only traditionalist with these intuitions. As a result, Sutton's argument that, in Case 1, the intuition that one ought to have formed a belief in *FORD* is grounded in the truth of *FORD* rather than in the fact that one could have justifiably believed it fails at least in the sense that it begs the question against those traditionalist who share my intuitions about Case 3.⁹

Another problem with Sutton's argument is that, even if we suppose that Sutton's intuitions here are unobjectionable, the argument simply does not establish that, in Case 1, there is no intuitive sense in which one ought to have believed P because one's belief would have been justified. If successful, the argument provides reason to think that there is *an* intuitive sense in which, at t_i , one should have believed P because it was true. It is consistent with this, however, that there is *a different* intuitive sense in which, at t_i , one should have believed that P because one's belief would have been justified. For the argument to go through, Sutton needs to argue, additionally, that there is no such intuitive sense. However, once again, Sutton does not do this.¹⁰

So, let's ask whether there is an intuitive sense in which that one should have formed a belief in *FORD* because it would have been justified. Following Sutton, I suggest that we answer this question by determining whether the intuition survives if we alter the case so

⁹It seems as though Sutton himself is aware that the intuitions that are to support his argument are not the most solid when he makes the following concession: "[T]he posterior evaluation argument . . . is arguably the least intuitively compelling of my arguments." (Sutton 2007, 60)

¹⁰Notice that Sutton's claim that there is an intuition that, in Case 3, one ought to form the belief in *FORD* won't help him here. After all, one's belief is, as Sutton also acknowledges, not justified. Hence, Case 3 is unsuited to show whether or not there is a sense in which one ought to form some belief because it is justified.

that one's evidence remains the same, while the target proposition is false. Consider:

Case 4. At t_i , one has an interest in the question whether one has a colleague who owns a Ford. One also has the opportunity to acquire a belief that one has a colleague who owns a Ford: one has excellent evidence in favour of *FORD*. Colleague *A* has shown one papers of a Ford in his name, a video tape that shows him winning a Ford in a well-known TV show, a number of highly trustworthy people have testified that he owns a Ford, he arrives at work in a Ford etc. What's more, at t_i , one has no reason to distrust any of this evidence.

Crucially, (a) one does not form a belief in *FORD*, and (b) one's belief in *FORD* would have been false: the evidence produced by *A* was misleading and no other colleague owns a Ford.

Is there an intuitive sense in which, at t_i , I should have formed a belief in *FORD*? It seems to me overwhelmingly plausible that there is: it was highly irrational of me not to when I had such excellent evidence in favour of *FORD*. Does this change if I evaluate the situation at a later time at which I have become aware of the fact that the apparently excellent evidence was in fact misleading? Not at all: it remains plausible that it was highly irrational of me not to form the belief and that, therefore, in a sense I should have formed it. If this is correct, PE2, is false.

It becomes clear that the posterior evaluation argument is bound to fail: The way in which Sutton ventures to defend PE2 begs the question against some traditionalists and, what's worse, there is reason to think that PE2 is in fact false. That means, however, that Sutton's last argument for the identity thesis also fails.

3 Why justified belief isn't knowledge

It becomes clear that Sutton fails to make a convincing case for the identity thesis. Even so, however, so far we have only seen that Sutton's case for the identity thesis is unconvincing. For all that I have shown so far, there is a case for the identity thesis out there. For that reason in this section I will provide some perfectly general reason to believe that the identity thesis is false: there is a concept of justified belief that is independent of the concept of knowledge in the sense that it is not coextensive with the concept of knowledge. In order to

achieve this, I will first argue that there are some distinctions that it makes sense for us to make but that we could not make unless we had an independent concept of justified belief.

There are propositions, even true ones, that, in certain cases at least, cognitive agents like us are not in a position to know. For simplicity's sake I will here focus on a case in which the target proposition is false. Let *A* be an agent who, in familiar circumstances, has the ability to perceptually identify sheep from a wide variety of viewpoints. Presently, *A* sees what looks like a sheep to him next to an alpine hut and forms a corresponding belief. Unbeknownst to *A*, however, there is a sheep-themed party going on inside the hut and the animal *A* sees is actually a dog in fancy dress. Now contrast *A* with another agent, *B*, who comes by a belief with the same content on the basis of 'evidence' provided by crystal balls, tea leaves and the like. Surely, in our epistemic assessment of the two beliefs, it makes sense to mark the distinction between *A* and *B*'s beliefs.¹¹ A concept of justified belief with application conditions distinct from those gov-

¹¹For those who, unlike me, do not find it immediately obvious that it makes sense to mark this distinction, here is a sketch of an argument in support of this claim (inspired by Sosa). Sosa (2007, 22-3) argues that all performances that have an aim are assessable in terms of accuracy (reaching the aim), adroitness (manifesting competence), and aptness (accuracy through adroitness). In the above case, aptness of cognitive performance is out of the question. Still, *A*'s cognitive performance differs from *B*'s in that *A*'s performance is adroit, whilst *B*'s isn't. After all, *A* manifests a disposition that, in normal circumstances at least, would ensure/make highly likely that the belief he forms is true (qualifies as knowledge). As a result, *A*'s disposition qualifies as a competence (see Sosa 2007, 29). (Notice also that it does so no matter whether the cognitive aim is construed standardly as true belief or in the way envisaged by Sutton as knowledge.) As opposed to that, the disposition *B* manifests is not such as to ensure or even make highly likely that, in normal circumstances, *B*'s belief is true (qualifies as knowledge). As a result, *B*'s disposition does not qualify as a competence (again, no matter whether we construe the cognitive aim in the standard way or as envisaged by Sutton). To the extent that it makes sense to distinguish between accuracy, adroitness and aptness of performances, it makes sense to distinguish, in our epistemic assessments, between *A* and *B*'s beliefs.

One might worry that we still haven't answered the question concerning the reasonableness of distinguishing between *A* and *B*'s beliefs to any degree of satisfaction lest we have elucidated why it makes sense distinguish between accuracy, adroitness and aptness of performances. Providing a fully satisfactory answer to this question seems too tall an order for the purposes of this paper. One obvious reason for making these distinctions, however, is that they facilitate improvements in future performances. Once we have figured out in what way a performance was wanting, we can assess what needs to be done to ensure reaching the aim in the future and take appropriate steps.

erning the concept of knowledge allows us to make this distinction.

One might object that this does not demonstrate the need for a distinct concept of justified belief. The distinction between *A* and *B* can be explained in terms of the concept of blameless belief. According to this suggestion, then, whilst neither *A* nor *B* are justified in their beliefs, *A*'s belief is blameless in a way in which *B*'s belief isn't. However, this move won't do the trick here. To see why not, consider a case in which someone, *C*, is hypnotised to form beliefs about the presence of sheep in response to certain cues (that are entirely unrelated to the presence of sheep). When such a cue is presented and *C* forms the corresponding belief, *C* is surely blameless in believing as he does. Yet, it makes as much sense to distinguish in our epistemic assessments between the beliefs of *A* and *C* as it makes sense to distinguish between the beliefs of *A* and *B*.¹² That is to say, even when knowledge is not attainable, it still makes sense to distinguish in our epistemic assessments between good belief, blameless belief and blameworthy belief. A concept of justified belief with application conditions distinct from those governing the concept of knowledge plays a part in allowing us to make just these distinctions.

But perhaps even this is not enough to show that it makes sense for us to have a distinct concept of justified belief. Following Sutton once more, one could maintain that the distinctions can be explained in terms of what the agent knows to be *probable*. More specifically, the thought is that, while neither *A* nor *C* knows, unqualifiedly, that there is a sheep next to the hut, *A* but not *C* knows this to be *likely*. It seems to me that there is little mileage in this move. To begin with, it may be that *C* also knows that it is likely that there is a sheep next to the hut. Suppose, for instance, that *C* has just been told by a knowledgeable source that there will be exactly one animal next to the hut and that it will likely be sheep. If so, he can still know that it's likely that there is a sheep next to the hut he is walking past even though his unqualified belief about the presence of a sheep does not qualify as knowledge since it is false. On the other hand, *A* may also fail to know that it is likely there is a sheep next to the hut. To see this, suppose that *A* does not even have the concept of probability. Since *A* does not have this concept, he is in no position to acquire beliefs about what is probable. So, when, in the present case, *A* forms a belief about the presence of a sheep, he not only fails to know, unqualifiedly, that there is a sheep next to the hut, but also

¹²Again, *A*'s but not *C*'s belief is adroit.

fails to know that, likely, there is a sheep next to the hut. Still, it makes sense to distinguish in our epistemic assessments between the unqualified beliefs of *A* and *C*. The distinction cannot be recovered in terms of a distinction in knowledge of probability. A distinct concept of justified belief plays a part in allowing us to make this distinction. If this is correct, of course, the identity thesis can only be expected to fail.*

REFERENCES

- ADLER, J. 2002, *Belief's Own Ethics*. Cambridge/MA.: MIT Press
- CHANDLER, J. 2010, "The Lottery Paradox Generalised?", *British Journal for the Philosophy of Science* 61, pp. 667–679
- DOUVEN, I. 2002, "A New Solution to the Paradoxes of Rational Acceptability", *British Journal for the Philosophy of Science* 53, pp. 391–410
- DOUVEN, I. 2006, "Assertion, Knowledge, and Rational Credibility", *Philosophical Review* 115, pp. 449–485
- EVNINE, S. 1999, "Believing Conjunctions", *Synthese* 118, pp. 201–227
- KAPLAN, M. 1996, *Decision Theory as Philosophy*, Cambridge: Cambridge University Press
- LEHRER, K. 1990, *Theory of Knowledge*, London: Routledge
- NELKIN, D. 2000, "The Lottery Paradox, Knowledge, and Rationality", *Philosophical Review* 109, pp. 373–409
- POLLOCK, J. 1990, *Nomic Probability and the Foundations of Induction*, Oxford: Oxford University Press
- RYAN, S. 1996, "The Epistemic Virtues of Consistency", *Synthese* 109, pp. 121–141
- SOSA, E. 2007, *A Virtue Epistemology. Apt Belief and Reflective Knowledge, Vol. 1*, Oxford: Oxford University Press
- STALNAKER, R. 1984, *Inquiry*, Cambridge/MA.: MIT Press.
- SUTTON, J. 2005, "Stick to What You Know", *Noûs* 39, pp. 359–96
- SUTTON, J. 2007, *Without Justification*, Cambridge/MA.: MIT Press

*Many thanks to the audiences of the *Justification Revisited* conference at Geneva and the *Philosophy Colloquium* at Purdue University. Thanks also to an anonymous referee for a number of very helpful comments on this paper. This work was funded by a postdoctoral fellowship with *Research Foundation – Flanders*.