

Truth as the aim of epistemic justification

Forthcoming in T. Chan (ed.), *The Aim of Belief*, Oxford University Press.

Asbjørn Steglich-Petersen
Aarhus University
filasp@hum.au.dk

Abstract: A popular account of epistemic justification holds that justification, in essence, aims at truth. An influential objection against this account points out that it is committed to holding that only true beliefs could be justified, which most epistemologists regard as sufficient reason to reject the account. In this paper I defend the view that epistemic justification aims at truth, not by denying that it is committed to epistemic justification being factive, but by showing that, when we focus on the relevant sense of ‘justification’, it isn’t in fact possible for a belief to be at once justified and false. To this end, I consider and reject three popular intuitions speaking in favor of the possibility of justified false beliefs, and show that a factive account of epistemic justification is less detrimental to our normal belief forming practices than often supposed.

1. Introduction

A popular account of epistemic justification holds that justification, in essence, is a means that we rely upon to achieve the ultimate aim of believing the truth about the propositions we are interested in forming beliefs about – in short, that epistemic justification aims at truth. An influential objection against this account points out that it is committed to holding that only true beliefs could be justified, which most epistemologists regard as an obvious falsehood, and thus as sufficient reason to reject the account. In this paper I attempt to defend the view that epistemic justification aims at truth, not by denying that it is committed to the factivity of epistemic justification, but by showing that, when we focus on the sense of ‘justification’ relevant to the aim account’s purposes, it isn’t in fact possible for a belief to be at once justified and false. In Section 2, I introduce the view that epistemic justification aims at truth, present the argument that it is committed to the factivity of justification, forestall certain initial objections to this argument, and point out some consequences of the argument to the related thesis that *believing* aims at truth. In Section 3, I consider and reject three popular intuitions speaking in favor of the possibility of justified false beliefs. In Section 4, I consider whether an account of epistemic justification that makes justified false beliefs impossible might be altogether too demanding, before concluding in Section 5.

2. The problem

Many philosophers have been attracted to a theory of belief, according to which believing *aims at truth*. Some understand this claim in *normative* terms, as expressing a constitutive norm of correctness for belief, for example that believing P is *correct*, in some distinctively normative sense, only if P is true.¹ I have argued that the claim is best understood in *teleological* terms, as the claim that when a person believes that P, she has the aim of believing P truly, or at least has some sub-intentional surrogate of such an aim. This interpretation gives rise to something similar to the norm of correctness, namely a criterion of *success*. Trivially, if believing P involves having the aim of believing P truly, someone believing P will be successful in doing so only if P is true.² However, for the purposes of this paper, I shall more or less ignore the dispute between ‘normativist’ and ‘teleological’ interpretations of the truth aim, and for convenience use ‘norm of truth’ and ‘aim of truth’ interchangeably, according to what seems most natural in the context. The problem I wish to discuss is relevant for proponents of both normative and teleological accounts.

As compelling as the truth norm may seem, almost all who defend it recognize that truth cannot be the only relevant consideration when evaluating beliefs. The main problem is that the truth norm, on its own, seems to provide little *guidance* for someone wishing to form a belief about some subject matter. It seems that the truth of a proposition doesn’t by itself make it the case that I ought to believe it, nor does its falsity make it the case that I ought to disbelieve it, even if I am interested in forming a belief about it.³ The additional relevant consideration when evaluating beliefs is often expressed in terms of *epistemic justification*, or in terms of a requirement that beliefs must be adequately backed by *epistemic reasons*.

¹ Or something similar – it is debatable what exact form the norm is best understood as having, but this will not matter a great deal for the purposes of this talk. For proponents of the normative interpretation, see N. Shah (2003), N. Shah and D. Velleman (2005), R. Wedgwood (2002), P. Boghossian (2003), among others. For discussion of the correct form of the truth norm, see K. Bykvist and A. Hattiangadi (2007), K. Glüer and Å. Wikforss (2009), among others.

² See D. Velleman (2000) and A. Steglich-Petersen (2006, 2009, 2011) for this interpretation.

³ I shall accept this point for the purposes of this paper, although I think that it is often exaggerated, for example as presented by K. Glüer and Å. Wikforss (2009). See Steglich-Petersen (2010) and Glüer and Wikforss (2010) for further discussion.

Whatever the details of the correct account of justification turn out to be, the mere fact that truth isn't the only relevant consideration when evaluating beliefs raises the question of how to understand the relationship between the two kinds of evaluation. It would be odd, at least *prima facie*, if the two kinds of evaluation turned out to be completely distinct from one another, and many philosophers have consequently tried to account for the nature of justification in terms of the aim of truth. One influential version of this type of account, which I shall from now on refer to as the 'instrumentalist' account of epistemic justification, tries to understand epistemic justification as, in essence, a *means* to achieving the ultimate aim of believing the truth with respect to the relevant propositions.⁴ As the account goes, since we cannot achieve the aim of believing the truth directly, we must instead rely on norms of epistemic justification as 'instruments' to forming true beliefs, and this instrumental role of epistemic justification is what accounts for its essential nature. Apart from being strikingly simple, one of the main virtues of this account is that it assimilates the normativity of justification, i.e. the nature of the 'oughts' that the requirement of epistemic justification gives rise to, to a kind of normativity that many philosophers find relatively unproblematic, not least from a naturalistic point of view, namely instrumental normativity. According to the instrumentalist account, whether a belief in a particular proposition is justified is simply a matter of whether it is formed in a way that, in a suitable fashion, advances the aim of believing the truth with respect to that proposition, and this property seems amenable to naturalistic analysis.

As a number of people have pointed out, however, there is a potential problem with this instrumentalist picture of the nature and normative force of justification: it seems committed to denying the possibility of justified false beliefs. There are a number of versions of this objection floating in the literature, developed most prominently perhaps by Stephen Maitzen (1995) and Richard Fumerton (2001; 2002), but the argument can be stated quite simply: Suppose that epistemic justification gets its rationale, or internal aim, from advancing the ultimate aim of believing the truth with respect to the propositions we are interested forming beliefs about. It follows that a belief is justified, or successful vis-à-vis the internal aim of epistemic justification,

⁴ For prominent examples of this type of account, see W. Alston (1989) and (1991), R. Foley (1987), R. Nozick (1993), P. Kitcher (1992), R. Giere (1989), L. Laudan (1990), and D. Velleman (2000).

only if the belief advances the aim of truth. But trivially, only true beliefs advance the aim of believing what is true – a false belief will detract from it, not advance it. So a belief can be justified, or successful vis-à-vis the internal aim of epistemic justification, only if it is true. Many would regard this as an absurd consequence of the account, and therefore a sufficient reason to reject it.⁵

The obvious initial response to this argument is to change focus from the justification of individual beliefs, to the methods or norms of justification by which beliefs ought to be formed. When evaluating whether justification serves the truth aim, we shouldn't evaluate whether individual justified *beliefs* serve that aim, but whether the aim is served by the general *methods* by which the beliefs are formed. When we ask whether a belief is justified, we are asking whether the belief is formed according to a method or norm which, when followed, *tends* to serve the aim by resulting in true beliefs. And it is clearly compatible with a certain method of belief formation *tending* to result in true beliefs, that not *every* belief formed in accordance with that method be true. So one might be tempted to conclude that the instrumentalist can allow justified false beliefs, after all. Cases of instrumental action like the following also support this line of thought.⁶ Suppose that I want to choose a strawberry-flavoured chocolate from the box, but I don't know which chocolates are strawberry-flavoured. However, I do know which chocolates are strawberry-shaped. If it is probable on my evidence that there is a high correlation amongst the chocolates in the box between strawberry flavour and strawberry shape, then it seems reasonable for me to aim at a strawberry-shaped chocolate as a means to the end of a strawberry-flavoured chocolate. So if a philosopher inspired by Maitzen and Fumerton objects that my plan is incoherent because I have forgotten that strawberry shape is not a sufficient condition for strawberry flavour, the objection may seem less than compelling.

But this reply on behalf of the instrumentalist account is incompatible with the ambitious explanatory goals of that account, for two reasons. The first reason is that the instrumentalist account attempts to identify the *internal* aim and not merely an

⁵ There are some notable exceptions to this. E.g. D. Armstrong (1973), T. Williamson (2000), and J. Sutton (2007) all seem to hold that a belief can be justified only if true.

⁶ I owe this example to Timothy Williamson.

external aim of epistemic justification. An internal aim of some type of aim-directed state or activity is an aim that plays a role in defining the essential nature of that state or activity, such that one couldn't be in the state or engage in the activity without having that aim. An external aim, by contrast, is any subsidiary aim one might have in some state or activity, as a contingent matter, for example when one performs the activity as a means to achieving something else, but could have performed it without that instrumental aim. For example, the aim-directed activity or action type of *dishwashing* plausibly has clean dishes as its internal aim – one couldn't be dishwashing unless one aimed at this. But one might have any number of external aims in engaging in dishwashing – making someone happy that the dishes have been washed, for example. The contingent nature of external aims means that one can be successful as far as the internal aim of some type of activity goes, without being successful in some external aim with which it happened to be performed. For example, I can successfully wash the dishes without being successful in thereby making someone happy. The example above involving instrumental action relies on this type of case: I can be successful in picking a strawberry shaped chocolate without being successful in the external aim of thereby coming to pick one that is strawberry flavored. But it is clearly not possible to be successful vis-à-vis the internal aim of an action, and not achieve that aim. I cannot be successful in my effort to pick a strawberry shaped chocolate, unless I actually end up with a chocolate that strawberry shaped, just as I cannot successfully wash the dishes without the dishes actually becoming clean. So if the internal aim of relying on epistemic justification is to believe the truth of the proposition one relies on justification to form a belief about, one clearly cannot be successful in doing so without that belief being true. And since 'justified' just means 'successfully justified', it follows that a belief cannot be justified unless it is true.

The second reason that the above initial reply on behalf of the instrumentalist fails, is that if we take seriously the claim that the essential nature of epistemic justification is that of being a means to achieving our aim of believing truly, the validity of any candidate norm of justification must be judged by whether beliefs formed in accordance with it advance or detract from achieving the aim of truth. The ambitious instrumentalist cannot help herself to a predetermined set of norms of justification, and, upon observing that following them in most cases advances the aim

of truth, decide that they are essentially means to advancing that aim, even if they sometimes allow false beliefs. Instead, she begins with the guiding assumption that the essential nature of epistemic justification is to advance the aim of truth, and on that basis decide what the valid norms or methods of belief formation ought to be. And, trivially, if a particular norm of justification allows that beliefs formed in accordance with it are on occasion false, that norm will, at least sometimes, condone beliefs that detract from the aim of truth, in which case the ambitious instrumentalist must reject the norm in favour of a norm which doesn't condone false beliefs. Maitzen (1995) illustrates this point by reference to a similar problem facing act utilitarianism in ethics. According to the act utilitarian, the morally right action to take in any given situation is the action that maximizes utility. But that doctrine faces the problem of making it very difficult to decide what to do – it is not always obvious which alternative will maximize utility, and sometimes it will be outright impossible to find out. So the act utilitarianism yields a decision rule, which is in many cases difficult or impossible to follow. This problem is in relevant respects analogous to the problem facing someone wanting to apply the truth norm directly to beliefs. A popular response is to adopt an alternative doctrine, rule utilitarianism, telling us to follow the general action rules, which, if followed, will *tend* to maximise utility. These general rules seemingly have the advantage of being more readily applicable in choice scenarios than the simple rule of act utilitarianism. Suppose, for example, that you are considering whether to lie in a particular situation. Act utilitarianism requires of you that you lie only if that is the action that results in the greatest utility in that particular situation. But how are you to determine this? Rule utilitarianism, on the other hand, tells you to lie if lying generally tends to maximise utility, which, it appears, is easier to determine. So while you may not be able to predict on a particular occasion whether utility is best served by lying or by telling the truth, you might know that lying *tends* not to maximize utility, and thus decide to tell the truth. This strategy is analogous to the method-based instrumentalist explanation of justification. But making this move from act to rule utilitarianism results in a well-known dilemma: either the adopted rules will be inconsistent with the basic tenet of utilitarianism, namely the ultimate aim of maximizing utility, or the new doctrine of rule utilitarianism will have to recommend the very same actions as act utilitarianism, thereby doing nothing to alleviate the initial concern about inapplicability. For suppose that the rules get their rationale from serving the ultimate aim of maximizing

utility. Then a rule will be valid only if following it will result in actions, which maximize utility. But that would effectively make the rule equivalent in its recommendations to the simple rule stating that one should always choose the act, which maximizes utility. If the ultimate aim of maximizing utility is to be served by the rules, the rules cannot differ in its recommendations from those made by the simple act utilitarian decision rule. Maitzen claims that a similar dilemma will face any attempt to save the instrumentalist account of justification by invoking imperfect methods of belief formation. Either the justifying norms, methods, or mechanisms for belief formation will be inconsistent with the basic aim of truth, or the justificatory norms will be satisfied only in when the simple truth norm is satisfied too.

Having rejected this initial objection to the argument that instrumentalist accounts are committed to the factivity of justification, it is important to limit the scope of this argument in a crucial respect. The problem as I have set it up was that if the internal aim of justification is to end up believing the truth, there couldn't be justified false beliefs, i.e. beliefs that are successful vis-à-vis the internal aim of justification, and yet false. But there might be another and more serious problem arising from reasoning parallel to that leading to the first problem: if the internal aim of justification is to end up believing the truth, there couldn't be *unjustified true beliefs*, i.e. beliefs that are successful vis-à-vis truth, but unsuccessful vis-à-vis the internal aim of justification. If the internal aim of justification were to end up believing the truth, a true belief would necessarily satisfy the internal aim of justification, and hence couldn't be unjustified. But one could clearly have a true but unjustified belief!⁷

To see why the instrumentalist isn't committed to denying the possibility of unjustified true beliefs in the way that she is committed to denying the possibility of justified false beliefs, it is instructive to compare the present problem to the so-called 'swamping problem' facing instrumentalist theories of the *value* of justification.⁸ If we suppose that the value of justification is to be understood as the instrumental value of promoting the intrinsic value of believing the truth, a problem arises when we wish

⁷ I am grateful to Timothy Williamson for urging me to consider this.

⁸ For a helpful recent discussion of the swamping problem, see D. Pritchard (2011). In my (2011), I have defended a teleological account of epistemic value.

to explain why a justified true belief is more valuable than an unjustified true belief, since the value to be promoted by justification would then already obtain in a true belief, regardless of its justification. The value of justification is 'swamped' by the value of truth, as it were. But there is a crucial difference between claiming that truth is the *value* to be promoted by epistemic justification, and claiming that truth is its internal *aim*. Here's the difference, starting at the general level: If S aims at bringing about some condition C by ϕ -ing, the obtainment of C does not by itself make that instance of ϕ -ing successful in bringing about C, even though the obtainment of C by itself is sufficient for S to get what he valued, and therefore aimed at bringing about by ϕ -ing. In short, it is possible to get what one values while being unsuccessful in bringing it about. Suppose, for example, that S aims at making his wife happy by washing the dishes. The fact that his wife already *is* happy does not mean that S successfully made her happy by washing the dishes, even though it does mean that S in the end obtained what he valued, and aimed for in washing the dishes, namely for his wife to be happy. This point holds regardless of whether the obtainment of C is the internal aim of ϕ -ing in the sense of being what ϕ -ing as an action-type necessarily aims at, or whether C is a mere external aim one adopts in ϕ -ing. Suppose, as above, that the internal aim of dishwashing is clean dishes. The fact that the dishes are already clean, or would get clean by magic regardless of anyone washing them, is not sufficient for an act of dishwashing to be successful in its internal aim of clean dishes. For that act to be successful in its internal aim, the dishes would have to get clean as a result of the act of washing them, and this holds even if we value the clean dishes equally regardless of how they got clean. So the fact that C could obtain without being the result of ϕ -ing does not present any obstacle to the obtainment of C being the internal aim of ϕ -ing. The internal aim of dishwashing is clean dishes, even if clean dishes by itself is insufficient to make an act of dishwashing successful. These considerations apply to the case of justification as well. Even if an instance of relying on epistemic justification aims at achieving the state of believing the truth about the relevant proposition, the fact that the proposition is truly believed does not by itself make the reliance on epistemic justification successful. For that to be the case, the true belief would have to be a result of the reliance on epistemic justification. So the fact that a belief could be true without this being the result of being justified, or indeed without being justified at all, is no obstacle to regarding the obtainment of a true

belief as the internal aim of justification. This means that the argument against the possibility of justified false beliefs as stated above doesn't establish the parallel conclusion against the possibility of unjustified true beliefs. The argument concerns the conditions under which an instance of justification is successful in its internal aim, and not merely the conditions under which the state of affairs aimed at happens to obtain regardless of reliance on justification.

If the above is correct, there are compelling reasons for thinking that the instrumentalist account of epistemic justification must deny the possibility of false justified beliefs. Insofar as one regards this as an unfortunate consequence, this doesn't just threaten a certain attractive account of the point of justifying one's beliefs. As pointed out recently by Kathrin Glüer and Åsa Wikforss (2009), it also threatens the plausibility of supposing that belief aims at truth in the first place. One reason for this is that many philosophers have been attracted to the notion of beliefs aiming at truth precisely because this seemed to provide a promising explanation of the nature of justification. If the truth aim cannot provide such an explanation, there will thus be one less reason for supposing that beliefs aim at truth. Another and perhaps more serious problem is that the norm of truth is sometimes defended from the charge of not providing guidance, and thus of being without normative significance, by arguing that we can conform to the norm of truth in an indirect way by way of ensuring that our beliefs are justified (e.g. Shah 2003). But if justification cannot be seen as essentially being a means to true beliefs, this defense will fail, in which case the thesis that beliefs aim at truth might itself begin to look less plausible.

In the rest of this paper, I shall attempt to defend the instrumentalist account of epistemic justification, and thus, in the process, defend the idea that beliefs aim at truth. However, I shall not doubt the soundness of the above argument, but instead focus on a single crucial presupposition that it makes, namely that a belief can be at once justified and false. I shall argue that the notion of justification relevant to the aim-theorists' concerns does *not* allow beliefs to be at once justified and false. Even so, it should be mentioned that the reasons I offer against the possibility of justified false beliefs do not depend on a prior acceptance of Maitzen and Fumerton's argument. Although their argument is my initial motivation for reevaluating the possibility of justified false beliefs, it plays no role in my arguments against it, and

anyone unconvinced by the Maitzen and Fumerton's argument can thus treat what follows as an independent defence of the factivity of justification.

3. The case in favor of justified false beliefs

In this section, I shall consider three popular intuitions speaking in favor of the possibility of justified false beliefs. I shall argue that they are either false, or concern a sense of 'justification' that is irrelevant as a guide when an agent considers what she ought to believe. But before I go on to this, I'd better define the sense of 'justification' relevant to our purposes a little more carefully.

The main complaint against the truth norm was that it doesn't provide guidance, in the sense of informing the agent what she ought to believe in regard to some proposition on any given occasion. In other words, it doesn't tell the agent what she has *reason* to believe. It is the sense of justification that can play this role that is interesting from the point of view of supplementing and defending the truth norm. Consequently, I shall restrict my question concerning justification in the following way. In asking whether some type of fact or consideration could act as justification for S in believing P, I shall focus on whether that fact or consideration could act as an *adequate epistemic reason for which* S believes that P. Some type of fact or consideration justifies S in believing P, in this sense, if mentioning that a fact or consideration of that type obtains, would be an adequate, i.e. sufficient, answer on S's behalf to the question 'what gives you epistemic reason to believe that p?'⁹ Call this the 'reasons constraint' on epistemic justification. Although this may seem like not much of a restriction at all, it will transpire later on that certain compelling intuitions concerning the justification of belief actually concern a different sense of 'justification', which, I shall argue, is not directly relevant to agents when considering what to believe.

⁹ The qualification 'epistemic' is important since, as I have argued elsewhere, if sufficient reason for belief is taken to imply that one ought to take up the belief, epistemic considerations alone never suffice. One must also have a practical reason to form beliefs about the relevant proposition. See Steglich-Petersen (2008) and (2011) for more on this.

3.1 Probabilism

The first popular intuition in favor of justified false beliefs I shall consider is what I call *Probabilism*:

Probabilism: When the evidential probability for S that P is sufficiently high, S has justification for believing that P.

I take this thesis to encompass not only theories of justification that explicitly understand justification as a matter of one's evidence raising the probability of a proposition being considered for belief, but also theories that rely on the probabilistic relation to ground other justifying properties. For example, on some interpretations of *reliabilism*, being formed by reliable processes justify beliefs simply because the relative frequency with which beliefs resulting from those processes are true, makes it *probable*, to some suitable degree, that the resulting beliefs are true. My discussion of Probabilism should therefore be taken to apply to these forms of reliabilism as well.

Before considering its merits, three remarks about the above statement of Probabilism are needed, which will also apply to the subsequent statements of intuitions in favor of justified false beliefs. The first thing to note is that Probabilism is a statement concerning the *sufficiency* of the antecedent for the consequent. So, on the definition of 'justification' given above, it says that the fact that a person's evidential probability for P is sufficiently high, would be an adequate reason for that person to believe that P. The second thing to note is that Probabilism is a statement about *propositional* justification rather than *doxastic* justification. In order for the evidential probability for S that P to provide doxastic justification for S to believe that P, S must base her belief that P on this. But whether S does this or not, we can ask whether the relevant evidential probability *would* justify S in believing P, were S to base her belief on that, which is what I shall mean by 'propositional' justification. Finally, Probabilism does not state anything about the cognitive relation S must stand in to the antecedent condition, in order for S to rely on it in justifying the relevant belief. I shall, as far as it is possible, hold this issue open (I return to it briefly in Section 4). However, there is a clear sense in which we can ask about the adequacy of reasons independently of any considerations having to do with the epistemic access we have to those reasons. When evaluating whether an agent, in ϕ -ing for reason R,

was ϕ -ing for an adequate reason, one line of inquiry would proceed by asking ‘Well, *suppose* that R (or the fact constituting R) is the case, and that the agent had adequate access to R, would R then be a good reason to ϕ ?’ It is a question of this kind I shall ask concerning Probabilism, and we do not need to suppose anything specific about the agent’s access to the antecedent condition at the present stage – even granting adequate access to the facts held by Probabilism to justify beliefs, I shall argue that the thesis fails.

Returning now to the main issue: Probabilism speaks in favor of justified false beliefs in the rather straightforward way that it is possible for a proposition to have a high probability given a person’s evidence, and yet be false. Lotteries provide an especially clear example of this: suppose that S holds a single ticket in a million-ticket fair lottery, and S knows these facts. In that case, the evidential probability for S that she will loose is extremely high. Yet, despite the high evidential probability that S will loose, it is possible that she won’t. Nevertheless, it has seemed plausible to many that in this situation, S is justified in believing outright that she will loose.¹⁰

My argument against probabilism has two main components. First I shall argue that probabilism is committed to permitting irrational practices of belief formation, and that the only obvious way of dealing with this problem leads to new problems. Later on, in Section 4, I shall argue that the intuitive appeal of probabilism can be accounted for in a way that does not allow false justified beliefs.

The irrational practice of belief formation that probabilism is committed to permitting, is what I shall call *doxastic risk-taking*. The notion of risk involved is the standard decision theoretic notion, according to which a decision is taken under risk whenever the agent knows the probabilities of the states of nature relevant to the outcome of the action, but does not know for certain which of the states will ensue. When forming a belief that P, the state of nature relevant to the success of doing so is that in which P is true, so we can define doxastic risk as follows:

¹⁰ Prominent examples include B. Russell (1948), Chisholm (1957), A. Goldman (1986), J. Hawthorne (2003), J. Pryor (2004), and D. Pritchard (2005).

Doxastic risk: When S forms a belief that P while having merely probabilistic knowledge that doing so will result in believing P truly, S forms the belief under risk.

Why is doxastic risk-taking always irrational, if ordinary risk-taking in action is frequently rational? One way of seeing this is by simply reflecting on what it means to take a risk in action. When an action is performed under risk, or at least *perceived* risk, the agent performs the action *in the hope* that a particular outcome will ensue. She doesn't *believe* that it will ensue, since if she did that, she wouldn't regard the action as a risky one. So the very notion of risk seems to exclude belief that the action will have a particular outcome. This means that there is something paradoxical about the notion of doxastic risk, since in this case, the risky 'action' is exactly a belief that the success-condition of the 'action' will ensue, which contradicts the idea that the belief was formed under risk. One cannot regard oneself as taking a risk in some particular act if one knows that taking the act will entail not regarding it as risky.

One might think that this paradoxical situation can be resolved by observing that 'under risk' refers to the agent's doxastic situation *prior* to completing the act. As this objection goes, performing an action under risk means that the agent is unsure about the outcome of the act prior to, and perhaps during its completion, but this is of course compatible with the agent coming to believe that a particular outcome ensued once the act is completed. For example, my bringing an umbrella on a walk may be a risky action in the sense that I, prior to and perhaps during portions of the walk, am unsure about whether it will rain and thus be worthwhile to carry the umbrella. But this is compatible with me coming to believe that it rains, for example in case that it indeed starts to rain. But the fact that bringing my umbrella was a risky action is not perturbed by this subsequent belief.

How would such an explanation go in the case of doxastic risk? Perhaps we could say that all it takes for a belief that P to be risky is that it is *taken up* or *formed* while the agent is unsure as to whether P, but that this is compatible with the agent no longer being unsure as to whether P, once the belief is formed. This would mean that the risky belief is essentially a revision of the initial doxastic state, but perhaps this is a consequence we can live with – after all, in the unproblematic umbrella case above,

the agent also comes to revise her doxastic state concerning the possibility of rain, from a state of being unsure to a state of being sure. But there are several crucial differences. First, whereas in the umbrella case, the revision in the agent's doxastic state happens in response to new evidence about the relevant state of nature (e.g. seeing that it rains), no new evidence would motivate the revision in the case of doxastic risk. The revision would simply happen in response to the state of being uncertain as to whether P. Second, the probabilist should be wary about any such revisions in doxastic state. After all, the initial doxastic state of ascribing a certain probability to P, is the state on which the probabilist proposes that we *base* the belief that P. So defending the rationality of doxastic risk by reference to the possibility revising in the initial doxastic state upon forming the risky belief is hardly an option for the probabilist.

Given this, it seems, in fact, that forming a belief under risk will necessarily result in rationally incompatible beliefs. Suppose that S, upon evaluating her evidence, finds that the probability of P upon her evidence is .6 (call this her 'evidential belief'). Suppose that .6 is the threshold above which believing P is deemed justified by Probabilism (this supposition is of course false, but assume it for the sake of argument). Since one can form a belief in response to an evidential belief only if one doesn't thereby revise that evidential belief, Probabilism then recommends that S both believe that her evidential probability of P is .6, and believe outright that P is true. But these seem like rationally incompatible beliefs for S. In order to resolve this, the obvious move for the probabilist is to raise the evidential probability required for justifying belief. But there does not seem to be any value of evidential probability short of 1, such that the evidential belief concerning the evidential probability of P, and the outright belief that P, are rationally compatible. So forming a belief under risk will always lead to rationally incompatible beliefs.¹¹

It may be objected that when forming beliefs in response to evidence, we rarely form beliefs in response to beliefs concerning evidential probabilities, but rather in response to beliefs in the propositions, which, if true, raise the probability of

¹¹ This does not exclude that we can and may take risks when adopting the attitude of *accepting* as opposed to *believing* some proposition. For this distinction, see especially J.L. Cohen (1992).

the target proposition to the appropriate level.¹² Suppose, for example, that you consider whether to believe that Gordon Brown has resigned as Prime Minister. The relevant evidence here may include various writings in the news media, and perhaps a televised speech in which he announces his resignation. According to the probabilist, such facts can justify believing that Gordon Brown has resigned if they jointly raise the probability of that proposition to an appropriate level. But it seems that even the probabilist could grant that no explicit belief in this latter probabilistic proposition is needed in order for the news media stories and the televised speech to justify the belief, in which case she would avoid commitment to the rationally incompatible beliefs. However, while it may well be the case that no explicit belief in the probabilistic proposition is needed, this would do little to help the probabilist. First of all, even if the probabilist can grant that no explicit belief in the probabilistic proposition is needed, she must still hold that it is in virtue of raising the evidential probability of Gordon Brown having resigned, that the evidential propositions justify believing that. So even if the probabilist can avoid commitment to explicit beliefs that are rationally incompatible, she is still committed to something just as bad, namely implicit reliance on the truth of some set of propositions, which cannot be rationally believed at once.

A further objection might be that since P, and the proposition that the probability of P upon a person's evidence is some value short of 1, can both be true at once, that person can believe both propositions without contradiction, and thus without irrationality. In reply, observe first that the fact that two propositions can be true at once doesn't entail that they can be believed at once without irrationally. The most famous example of such a pair of propositions is the one involved in 'Moorean absurdities', where a person judges that P and that she does not believe that P. Clearly, these propositions can be true at once, and it is thus also possible to *believe* that P and that one does not believe that P without any contradiction in content. Yet, most people would find such a pair of beliefs somehow irrational.¹³ So the mere compatibility of two propositions doesn't make the pair of those propositions available for rational belief. Still, some might feel that declaring the relevant beliefs

¹² I am grateful to Ralph Wedgwood for raising this objection.

¹³ Explaining exactly *why* it would be irrational has proven difficult (hence 'Moore's paradox'). For a recent collection devoted to this, see M. Green and J. Williams, eds., (2007).

rationally incompatible simply begs the question against the probabilist. For those, it might be more convincing to consider the role such beliefs would play in motivating rational action.

Suppose that an agent must choose whether to ϕ , where the success of ϕ -ing is P-dependent. Suppose further that the agent believes both that P and that the probability of P upon her total evidence is .9. It is not difficult to imagine situations in which she should ϕ in response to the one belief, but not ϕ in response to the other, thus leaving the agent with contradictory recommendations. It is no help for the probabilist to claim that the evidential belief is relevant to the rationality of action only via rationalizing an outright belief, since in the context of an action, the rationality of which depends upon the exact probability that it will be successful, it should strike us as especially implausible to let an evidential probability short of 1 justify an outright belief that the action will be successful. Nor is it any help to point to decision-rules that take probability intervals or fuzzy probability regions rather than simple probabilities as inputs, since by hypothesis, no such intervals or regions exist in the case at hand.¹⁴ So it seems that allowing the rationality of both believing that P and that one's evidential probability of P is some value short of one, will come at the cost of allowing conflicting action recommendations.

The obvious probabilist reply (that some readers no doubt have been waiting for) is to retreat to a principle concerning the justification of *degrees of subjective credence*. A plausible version might be this:

Credence Probabilism: When the evidential probability for S that P is D, S has justification for adopting credence in P to degree D.

As plausible as this principle may be, however, it doesn't immediately tell us about the justification of *outright belief*, which was our initial interest. If Credence Probabilism is to tell us anything about the justification of outright belief, some principle must thus be provided which links the justification of degrees of credence to

¹⁴ For an example of decision rules involving probability intervals, see P. Gärdenfors and N. Sahlin (1982). For the notion of fuzzy regions of probabilities, see Sturgeon (2008).

the justification of outright belief. One such principle is known as the Lockean Thesis:¹⁵

The Lockean Thesis: Outright belief that P = Any degree of credence that P above threshold T.

The Lockean Thesis together with credence Probabilism entails that when the evidential probability for S that P is above a certain threshold (the threshold of evidential probability which justifies a degree of credence above the threshold for outright belief), S will be justified in an outright belief that P. This seems to resolve the problem with rationally incompatible beliefs, since the belief supported by the evidence ‘really’ just is a degree of credence, compatible with the evidential belief.

However, the combination of Credence Probabilism with the Lockean Thesis gives rise to another problem. The problem has to do with the conditions under which one is justified in making *assertions*. The following principle seems plausible:

The Belief-Assertion Principle: When S is epistemically justified in outright believing P, S is justified in asserting P.

The Belief-Assertion Principle together with Credence Probabilism and the Lockean Thesis entails that when the evidential probability for S that P is above a certain threshold (the threshold of evidential probability which justifies a degree of credence above the threshold for outright belief), S will be justified in asserting that P. But it seems false that a mere evidential probability of P above some threshold short of 1 could justify a person in asserting that P. It could at most justify that person in asserting that P is quite likely, or in asserting that she is quite confident that P. Not in baldly asserting that P. If the Belief-Assertion Principle is correct, either Credence Probabilism or the Lockean Thesis must thus be false. But if either of those principles

¹⁵ For recent discussions of the Lockean Thesis, see S. Sturgeon (2008) and R. Foley (2009). The term ‘Lockean Thesis’ is sometimes used for the claim that the rationality of outright belief is a part of the rationality of degrees of credence, but I shall use it for the identity claim below.

is false, the retreat from Probabilism to Credence Probabilism cannot solve the original problem of permitting doxastic risk, while accounting for the justification of outright belief, at least not in lieu of some other principle connecting the justification of degrees of credence with that of outright belief. In sum, the prospects for Probabilism as an intuition concerning justification of outright belief are beginning to look less promising, and the thesis therefore doesn't lend support to the possibility of justified false beliefs. This does not mean, of course, that evidential probabilities have no role to play in justifying beliefs. What that role might plausibly be, I will return to in Section 4.

3.2 Blamelessness

The second main intuition in favor of justified false beliefs is what I shall call *Blamelessness*:

Blamelessness: When the epistemic situation of S is such that she couldn't be blamed for believing P even if P is false, S has justification for believing P.

The possibility of justified false beliefs does not follow immediately from this intuition, but it does when the intuition is coupled with the very commonplace observation that we frequently are blameless for our false beliefs.

Especially vivid examples of blameless yet false beliefs are to be found in skeptical scenarios, such as those in which a person's perceptual inputs are systematically manipulated to misrepresent the person's surroundings. In such cases, we are inclined not to blame the person for her false beliefs about her surroundings. And for many, this is tantamount to regarding those beliefs as justified. But much more commonplace examples are also easy to come by; for example, people are often misled by persons they have no reason to distrust, without being blameworthy for their false beliefs.

However, in the sense of 'justification' we are interested in here, namely that living up to the reasons constraint, Blamelessness should strike us as false. We never do, nor ought we, form beliefs in response to the antecedent condition that we would

be blameless for doing so. The fact that one would be blameless for believing P is not a good reason for believing P.

We can back up this *prima facie* verdict in the following way. Note first that for any norm N, we can distinguish between the conditions under which an agent conforms to N, and the conditions under which an agent can be held blameless for failing to conform to N. There is a real difference here: trying to conform with a norm N is not the same as trying to satisfy the conditions under which one can be excused for not conforming to N. Further, given this distinction, if Blamelessness is a genuine norm that we ought to conform to when forming beliefs, there must also be a set of conditions under which one can be excused for not conforming to Blamelessness. But that is absurd: there is no such thing as the set of conditions under which one can be held blameless for failing to be blameless in forming a belief. So it is doubtful that Blamelessness should guide us when forming beliefs.

This does not entail that it is uninteresting to ask whether someone can be blamed for a false or unjustified belief. It is often very important to know whether someone exerted proper care when forming his or her beliefs. It may even be the case that there is a sense of ‘justification’ that would be the proper label for this kind of status (I doubt that many epistemologists have a very good grip on their pre-theoretical linguistic intuitions concerning the term ‘justification’). But if what we are interested in is the norm we ought to comply with when forming beliefs, Blamelessness is not what we are looking for.

3.3 Modal Conditions

So far, I have considered and rejected two intuitions one might have thought favored the possibility of justified false beliefs. I argued that Probabilism is false, and that Blamelessness concerns a different sense of ‘justification’ from the one we’re interested in. An important part of what makes Probabilism so implausible, is that it allows probabilistic evidence for P to justify outright belief that P, even when the possible error scenarios in which P is false are probabilistically ‘on a par’ with the scenarios in which P is true. This holds in lottery cases where the possibility of one’s ticket being the winning ticket is a highly unlikely one, but not any stranger or farther removed from actuality than the possibilities in which any of the other tickets win.

But perhaps there are other justifying conditions of a broadly modal kind, that allow for the possibility of error without the error possibilities being probabilistically on a par with those in which the belief is true.

There are many candidates in the literature for modal conditions linking belief to truth. Most prominent of these are so-called ‘safety’ conditions, which obtain roughly whenever a belief is true not only in the actual world, but also in some class of relevantly similar non-actual worlds, with error possibilities farther removed in modal space. These are typically presented as necessary conditions for a belief to count as *knowledge*, rather than as conditions or analyses of epistemic justification. But most importantly for our purposes, safety conditions don’t typically allow for false beliefs, since they require the belief to be true in the actual world. So safety conditions don’t speak in favor of the possibility of justified false beliefs. If there is some modal condition on justification, which speaks in favor of justified false beliefs, it must thus be some weaker condition than safety.

A compelling illustration of a condition linking belief to truth in a seemingly stronger sense than a merely probabilistic one, while also allowing for falsity, can be found in Smith (2010), who adapts the example from Nelkin (2000). Suppose that Bob has configured his computer screen such that, whenever it is turned on, the background color is determined by a random selection mechanism. One in a million times, the screen will be red; the rest of the times it will be blue. Suppose that Bob turns on the screen and leaves the room before anything appears on it. It now seems that on Bob’s evidence, there is a .999999 chance of the screen being blue. Compare now Bob’s epistemic situation to Bruce’s: when Bob walked out of the room, Bruce walked in and now *sees* that the screen is blue. It seems that from a purely probabilistic point of view, Bruce’s evidence is *worse* than Bob’s – after all, for all he knows, there’s a chance that he is hallucinating or being played a clever trick with colored lights, and it is plausible enough to suppose that the probability of any of these scenarios is higher than .000001. In spite of this, it seems that Bruce is better justified than Bob in believing the screen to be blue, despite being worse off probabilistically. Bruce also seem justified in *asserting* that the screen is blue, while we would expect Bob to hedge an assertion about the color of the screen in a way that indicates that he merely knows it to be highly probable that the screen is blue.

There clearly is a sense in which Bruce's belief is justified, despite the possibility of falsity, and the comparison to Bob shows that Bruce's justification does not come down to a matter of mere evidential probability. Because of this, the case also seems to exemplify a set of circumstances in which one can be justified in believing P outright, despite the evidential probability being short of 1. The question is whether the sense in which Bruce is justified is the one we're interested in. Here I think matters are less clear. Given the discussion above, we have (at least) two rival interpretations of Bruce's justification. One possibility, of course, is that the sense in which Bruce is justified *is* the guiding sense of justification relevant to the instrumentalist account of epistemic justification, in which case that account would be in trouble. But another possibility is that Bruce's belief is justified merely in the sense of being blameless, in which case his example would not cause trouble for the instrumentalist account. In order to decide between these rival interpretations, a more detailed diagnosis of the case of Bruce and Bob is needed.

The diagnosis proposed by Smith himself does not decide this question either way. Smith's explanation of the difference between Bruce and Bob is that whereas falsity in the case of Bruce's belief would require an explanation, no explanation would be required in the case of Bob's belief, since the possibility of falsity is probabilistically 'on par' with the any of the scenarios in which Bob's belief turns out true. Just like we wouldn't feel that winning the lottery would require a special explanation, despite being so unlikely, we wouldn't feel a need to explain the falsity of Bob's belief about the color of the screen if it turned out to be red, despite being so unlikely. Bruce, on the other hand, would feel puzzled if the screen turned out to be red, despite appearing blue – some explanation would be needed: How could he get it wrong? But does the fact that error would require an explanation in Bruce's case and not in Bob's indicate that Bruce is justified in a stronger sense than being blameless for his belief? To answer that, we must look deeper into Bruce's situation.

One important clue comes from reflecting on how Bruce himself would think of the possibility of his belief being false. When Bruce formed his belief that the screen is blue in response to *seeing* that the screen is blue, Bruce, if he is anything like an average believer, most likely didn't even consider the possibility of him

hallucinating or being played a clever trick. And we would not blame him for not thinking of these error possibilities. One plausible explanation of us withholding blame for this, is that it would take considerable imagination and ingenuity to conjure up ways in which evidence as solid as Bruce's could allow for error. And, in general, we are inclined to withhold blame for people failing at difficult tasks requiring imagination and ingenuity. Bob's evidence, on the other hand, came equipped with an inbuilt and obvious error possibility, however unlikely, in as much as his evidence was that of knowing about the probabilities encoded in the random selection mechanism determining the color of the screen. So it would *not* put any particular strain on Bob's imagination to think of ways he could be wrong about the color, and we would consequently not withhold blame from Bob were he to disregard these very obvious possibility of error. If prompted to consider the question, or helped to imagine the relevant scenarios, Bruce might of course have come to believe that him having a perceptual experience of the blue screen *did* leave open the possibility of him having this experience because of a clever trick, rather than because of the screen actually being blue. But crucially, this would presumably make Bruce think differently of his grounds for believing that the screen is blue. He would then be taking into consideration one or more error possibilities, whereas before he didn't take into consideration any such possibilities at all. But once Bruce takes into considerations these error possibilities, it becomes unclear how his epistemic situation differs from Bob's, especially if we've already conceded that the probability of error in Bruce's case may well be as high as or higher than in Bob's case. If Bruce takes seriously these possibilities of error, he should find that he is on a par with Bob after all. But nothing has changed about Bruce's grounds for believing that the screen is blue – his evidence is the same, namely that of seeing that it is blue. But by coming to realize that the evidence leaves open certain hard to imagine error possibilities, Bruce thinks differently of the support his evidence provides for his belief about the color of the screen. He realizes that he was wrong about the level of support provided by his evidence for the belief that the screen is blue.

If this diagnosis is plausible, it suggests that the sense we have of Bruce being better justified than Bob comes down to a matter of Bruce's error possibilities being more difficult to imagine than Bob's, thus making it more forgivable for Bruce to fail to imagine them. The fact that Bruce would think differently of his justification if the

error possibilities were brought to his attention, shows that Bruce's justification is superior to Bob's only in the sense of it exempting him from blame in case he ends up with a false belief, to a higher degree than Bob's justification would exempt Bob from blame. But in that case, the sense in which Bruce's belief is better justified than Bob's, is not a sense that should be a worry to the instrumentalist account of epistemic justification.

This concludes the negative part of the paper, arguing against some of the most common intuitions in favor of the possibility of justified false beliefs. I do not pretend to have refuted these intuitions conclusively, or to have discussed every possible consideration counting in favor of justified false beliefs. But I do think the above should at least give reason to explore the plausibility of accounts of justification, which do not allow false beliefs to be justified. I shall not attempt to flesh out a detailed positive account of epistemic justification with that feature here. In the remaining part of the paper, I will instead concentrate on an objection that is no doubt on the readers mind by now: If neither evidential probabilities short of 1, nor additional circumstances such as those Bruce found himself in, could justify belief in a sense relevant to the instrumentalist accounts purposes, won't the instrumentalist account of epistemic justification be impossibly demanding? I shall argue that it needn't be as demanding as it initially may seem.

4. The Demandingness Objection

To fix matters in discussing whether the instrumentalist account is committed to an excessively demanding account of epistemic justification, it will be useful to consider a slightly more precise thesis than that of justification being factive. Supposing that the instrumentalist account of justification will let one's justification depend on one's evidence (on some suitable construal of that property), it seems that the instrumentalist, insofar as he is committed to the factivity of justification, must at the very least be committed to evidence being factive, and to one's evidence only justifying what it ensures the truth of. This minimal constraint can be summarized in what we might call *Implication*:

Implication: S has justification for believing P only if the truth of P is implied by S's evidence.

It is important to stress that Implication is not intended as an account of epistemic justification. For one thing, it only indicates a necessary, not a sufficient condition for a belief to be justified. P's truth being implied by one's evidence could not plausibly be seen as a sufficient condition for being justified in believing P, since we would then have justification for believing things we clearly don't have reason to believe, such as the conjunction of all of the propositions that form part of one's total evidence, or, even worse, all of the logical consequences of those propositions, however complicated. So further constraints on what it takes to have justification for believing P, in addition to P being implied by one's evidence, are clearly needed. Nevertheless, the Implication thesis is useful to consider for the purposes of finding out whether the instrumentalist account of justification is overly demanding, because even Implication by itself will seem much too strict for many, and the troublesome demandingness associated with the instrumentalist account being factive will stem from its commitment to something like Implication. But I think that this sense of Implication being excessively demanding can be dispelled at least in part by considering three mitigating considerations.

The first mitigating consideration is that in many cases where a person's evidence doesn't imply the truth of P, Implication allows that the evidence may nevertheless justify the person in believing a proposition concerning the evidential *likelihood* that P, by implying the truth of this probabilistic proposition. This is not *always* the case. There are cases in which the fact that the evidential probability for S that P is D will not be implied by S's total evidence, since even if E is S's total evidence, E does not imply that S's total evidence includes E or that E includes S's total evidence. But in cases where the agent has access to those latter facts, the probabilistic proposition is implied by the evidence. In cases where no precise evidential probability is yielded by the evidence, or we are unable to assess the precise probability, we may be justified in believing an equally imprecise probabilistic proposition. For example, seeing Gordon Brown announcing his resignation in a televised speech may make it *quite* likely or *very* likely that Gordon Brown has resigned, without yielding a precise evidential probability for this. In such cases,

Implication allows that we believe a correspondingly vague proposition, for example that it is quite likely, given the evidence, that Gordon Brown has resigned. But however vague this proposition is, the evidence does imply its truth.

So even if we are never justified in believing a falsehood, it is compatible with the Implication thesis that we are often justified in believing some proposition to be likely, which nevertheless turn out false. I suspect that many have found it plausible that we can be justified in believing a falsehood because they confuse the subtle distinction between being justified in believing some proposition to be very likely, and being justified in simply believing that proposition to be true. The psychological reaction upon discovering that a proposition is false is also likely to be similar, given the two kinds of prior attitude. Even if the belief that some proposition is likely is not strictly speaking falsified by the falsity of the proposition believed to be likely, one is nevertheless likely to feel surprised by this, given one's evidence. What is more, justified beliefs, or knowledge, concerning such likelihoods are sufficient as grounds for rational action. There is no reason to suppose that the role played by rational degrees of credence in decision theory couldn't be equally well played by justified outright beliefs concerning probabilities.

The second mitigating consideration is that even in cases where we lack epistemic justification for *believing* a proposition, it is possible to be justified in *accepting* a proposition for some restricted range of purposes (Cohen 1992), where the latter justification will often involve a combination of epistemic and pragmatic considerations. There may be cases, for instance, where the computing costs of relying on probabilistic knowledge in deciding what to do exceeds the potential benefits of doing so, in which case it is reasonable to instead rely on acceptances of the truth of some simple non-probabilistic propositions, and act *as if* some proposition is simply true rather than just likely. Although this is not the same as having adequate epistemic reason to *believe* such simple propositions, we may sometimes confuse this with the justification of acceptances.

The third mitigating consideration in reply to the demandingness objection is that Implication says nothing about the conditions under which one can be held *blameless* for believing what one is not justified in believing. It is possible for a

person's belief to not satisfy Implication, and nevertheless be blameless. Plausibly, one condition for being blameless is that one made an honest attempt at forming a justified, and thus true belief, exerting as much care and skill as one could muster in assessing one's evidence. As suggested above, it may also make one less blameworthy if the relevant error possibilities upon one's evidence were hard to imagine or understand. No doubt, there are many more complex and important things to be said about these conditions under which one can be considered blameless for a false belief, and there is no harm done in continuing to use the term 'justification' to speak of them, as long as one keeps in mind that those conditions cannot act as antecedent conditions in a norm that guides agents when forming beliefs.¹⁶ But the feeling that Implication is too demanding may just come down to confusing justification in the sense that are important for our present purposes, with the conditions for being blameless for one's beliefs.

There are no doubt many other objections to the claim that justification is factive that I cannot consider here. I wish to close, however, by briefly addressing a different sort of concern, of a more methodological nature, which the comments made in this section may have raised for some.¹⁷ The concern is that the meager conceptual resources of belief, truth, and the instrumental relation, which the ambitious instrumentalist proposes to explain epistemic justification in terms of, won't be rich enough to yield an adequate account of epistemic justification, since it leaves out the idea that justification depends on some property more *accessible* to the agent than truth, without providing an obvious way of accounting for what that property might be. In the discussion above, I simply took for granted an understanding of evidence, which, as on any reasonable construal of that term, must be somehow accessible to the agent having it. I did that with the sole purpose of being able to frame the discussion of demandingness in the more concrete context of the Implication thesis. But it hasn't yet been shown how evidence, and the access we have to it, can be explained in terms of the basic elements of the instrumentalist account, and one may worry that such an explanation isn't forthcoming. If the instrumentalist resolves this by simply helping himself to an understanding of evidence and accessibility, one might worry that this

¹⁶ For an interesting study devoted to the conditions under which we are blameless for our beliefs, see N. Nottelmann (2007).

¹⁷ Again, I am indebted to Timothy Williamson for urging me to consider this.

will make the instrumentalist account obsolete, since all of the interesting epistemological questions will be settled in advance by the presupposed understanding of those terms, leaving no substantial explanatory work to be done by the instrumentalist account. Or so the worry goes.

While I agree that there are large and unresolved questions about the nature of evidence and accessibility for the instrumentalist to answer (I hope to do so on another occasion), they are not as intractable as the above suggests. The above worry presupposes that an eventual instrumentalist account of justification can invoke no epistemic resources other than belief, truth, and the instrumentalist relation, without becoming obsolete. But it is far from clear why this should be so. The instrumentalist account is an account of the essential aim or rationale of epistemic justification. It thus provides the normative back foil necessary for deciding between theories of the more specific epistemic properties that play a role in justifying beliefs. Suppose, for example, that we are considering a substantial question about evidence, such as the question of whether evidence is factive. It seems that we can decide this question only if we already have an idea of the role evidence is to play in justifying beliefs, which in turn requires that we have settled on a view of the essential point of epistemic justification. If we thought that epistemic justification was essentially a matter of rendering epistemic agents *blameless* for their beliefs, rather than being a matter of guiding them to believe the truth, we would presumably take a very different view of the nature of evidence – in particular, there would be no motivation for regarding it as factive. So a theory of the essential point of epistemic justification will in many cases decide, and be needed to decide, between rival accounts of more particular epistemic properties. This means that even if our theory of particular epistemic properties such as evidence will invoke additional resources than those relied upon by the instrumentalist account of the aim of epistemic justification, the instrumentalist account will be far from obsolete. It is, at least in part, our opinion about the basic aim of epistemic justification that gives us reason to prefer one theory of some lower level epistemic property over another. This doesn't mean, of course, that no other consideration can influence our theories of lower level epistemic properties. If the instrumentalist account of epistemic justification is workable only if we accept a theory of evidence that seems independently very implausible, that will of course count against the instrumentalist account. But that just shows that no part of the

overall account is isolated from critique, which, it seems, is a characteristic of any philosophical theory. It does not follow from this that the basic picture of the aim of epistemic justification cannot play a role in deciding between theories of lower level epistemic properties.

5. Conclusion

I have argued that an influential objection to the instrumentalist account of epistemic justification, namely that it precludes the possibility of justified false beliefs, is not in fact an objection, but a plausible consequence of the theory. Of course, much more must be said before we have a completely convincing case against the possibility of justified false beliefs, but I hope that the reader will find it plausible that if we by ‘justification’ mean an adequate guiding reason to take up belief in a proposition, the intuitions usually cited in favour of the possibility of justified false beliefs are less compelling than they seem initially. There are no doubt other reasons why someone might reject the instrumentalist account of epistemic justification. But I hope to have shown that at least one objection to it is less compelling than often supposed.¹⁸

References

- Alston, W. P. (1989), *Epistemic Justification* (Ithaca: Cornell University Press).
- Armstrong, D. M. (1973), *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press).
- Boghossian, Paul A. (2003), 'The Normativity of Content', *Philosophical Issues*, 13, 31-45.
- Bykvist, Krister and Hattiangadi, Anandi (2007), 'Does thought imply ought?' *Analysis* 67, 277-85.

¹⁸ The ideas developed in this paper originate in my 2007 PhD dissertation, and owe much to discussion with my supervisor Jane Heal, and examiners Quassim Cassam and Alan Millar. Previous versions of the paper were presented at Oslo University, Aarhus University, University of Calgary, and University of Copenhagen. I am grateful to Timothy Chan, Nishi Shah, Pascal Engel, Åsa Wikforss, Kathrin Glüer, Ralph Wedgwood, J. Adam Carter, Paul Horwich, Anandi Hattiangadi, Anders Nes, Olav Gjelsvik, Stephen Mumford, Johanna Seibt, Lars Bo Gundersen, Eline Busck Gundersen, Anne Meylan, Carl Erik Kühl, Asger Steffensen, and two anonymous referees for Oxford University Press for useful comments. I am especially indebted to Timothy Williamson, who provided extremely helpful written comments on an earlier draft, and presented a reply to the paper at the 2009 CSMN Aim of Belief conference, which improved the paper substantially. Finally, I wish to thank the Danish Research Council for Culture and Communication for financial support.

- Chisholm, R. (1957), *Perceiving* (Ithaca, NY: Cornell University Press).
- Cohen, L. Jonathan (1992), *An Essay on Belief and Acceptance* (New York: Clarendon Press).
- Foley, R. (1987), *The Theory of Epistemic Rationality* (Cambridge: Harvard University Press).
- Foley, R. (2009), 'Belief, degrees of belief, and the Lockean Thesis', in F. Huber and C. Schmidt-Petri (eds.), *Degrees of Belief. Synthese Library Series*. (Synthese Library, Dordrecht: Springer).
- Fumerton, Richard (2001), 'Epistemic Justification and Normativity', in Matthias Steup (ed.), *Knowledge, Truth and Obligation: Essays on Epistemic Responsibility and the Ethics of Belief* (Oxford: Oxford University Press).
- (2002), 'Theories of Justification', in Paul K. Moser (ed.), *Oxford Handbook of Epistemology* (Oxford: Oxford University Press), 204-33.
- Gärdenfors, P. and Sahlin, N. (1982), 'Unreliable probabilities, risk taking, and decision making', *Synthese* 53, 361-86.
- Giere, R. (1989), 'Scientific Rationality as Instrumental Rationality', *Studies in History and Philosophy of Science* 20, 377-384.
- Glüer, K. and Wikforss, Asa (2009), 'Against Content Normativity', *Mind*, 118, 31-70.
- Glüer, K. and Wikforss, Asa (2010), 'The Truth Norm and Guidance: A Reply to Steglich-Petersen'. *Mind*, 119, 757-61.
- Goldman, A. (1986), *Epistemology and Cognition* (Cambridge, MA: Harvard University Press).
- Green, Mitchell and Williams, John N. (eds.) (2007), *Moore's paradox: new essays on belief, rationality, and the first person* (Oxford: Clarendon Press) X, 247 s.
- Hawthorne, J. (2003), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Kitcher, P. (1992), 'The Naturalist's Return', *The Philosophical Review* 101, 53-114.
- Laudan, L. (1990), 'Aimless Epistemology?', *Studies in History and Philosophy of Science* 21, 315-322.
- Maitzen, Stephen (1995), 'Our Errant Epistemic Aim', *Philosophy and Phenomenological Research*, 55 (4), 869-76.
- Nelkin, D. (2000), 'The lottery paradox, knowledge and rationality', *Philosophical Review* 109, 373-409.
- Nottelmann, N. (2007), *Blameworthy Belief, Synthese Library Series* (Dordrecht: Springer).

- Nozick, R. (1993), *The Nature of Rationality* (Princeton: Princeton University Press).
- Pryor, J. (2004), 'What's wrong with Moore's argument?', *Philosophical Issues 14: Epistemology*.
- Pritchard, D. (2005), *Epistemic Luck* (Oxford: Oxford University Press).
- (2011), 'What is the swamping problem?' in A. Reisner and A. Steglich-Petersen (eds.), *Reasons for Belief* (Cambridge: Cambridge University Press), 244-59.
- Russell, B. (1948), *Human Knowledge: Its Scope and Limits* (London: Unwin).
- Shah, N. and Velleman, D. (2005), 'Doxastic Deliberation', *Philosophical Review*, 114, 497-534.
- Smith, M. (2010), 'What else justification could be', *Noûs* 44, 10-31.
- Steglich-Petersen, A. (2006), 'No norm needed: on the aim of belief', *Philosophical Quarterly*, 56, 499-516.
- (2008), 'Does doxastic transparency support evidentialism?' *Dialectica* 62, 541-47.
- (2009), 'Weighing the Aim of Belief', *Philosophical Studies*, 145, 395-405.
- (2010), 'The Truth Norm and Guidance: a Reply to Glüer and Wikforss', *Mind* 119, 749-56.
- (2011), 'How to be a teleologist about epistemic reasons', in A. Reisner and A. Steglich-Petersen (eds.), *Reasons for Belief* (Cambridge Cambridge University Press), 13-33.
- Sturgeon, S. (2008), 'Reasons and the grain of belief', *Noûs* 42, 139-65.
- Sutton, J. (2007), *Without Justification* (Cambridge, MA: MIT Press).
- Vahid, H (2003), 'Truth and the aim of epistemic justification', *Teorema* 20, 83-91.
- Velleman, J. David (2000), *The Possibility of Practical Reason* (Oxford: Oxford University Press).
- Wedgewood, Ralph (2002), 'The Aim of Belief', *Philosophical Perspectives*, 16, 267-97.
- Williamson, Timothy (1998), 'Conditionalizing on knowledge', *British Journal for the Philosophy of Science*, 49, 89-121.
- (2000), *Knowledge and Its Limits* (Oxford: Oxford University Press).