
WHY DOES JUSTIFICATION MATTER?

BY

MATTHEW WEINER

Abstract: It has been claimed that justification, conceived traditionally in an internalist fashion, is not an epistemologically important property. I argue for the importance of a conception of justification that is completely dependent on the subject's experience, using an analogy to advice. The epistemological importance of a property depends on two desiderata: the extent to which it guarantees the epistemic goal of attaining truth and avoiding falsehood, and the extent to which it depends only on the information available to the believer. The traditional internalist notion of justification completely satisfies the second desideratum and largely satisfies the first.

1. Introduction

Once, epistemologists may have agreed on the broad outlines of what 'justification' meant, even if they disagreed about when beliefs were justified. When Gettier refuted the justified true belief account of knowledge, he had no need to defend a particular definition of justification.¹ The glosses on 'justification' that he discussed conformed to traditional internalist definitions, and it was obvious to all readers that the traditional internalist definition was the appropriate one.

Those days are gone. Today there is no guarantee of agreement even on the broadest outlines of the definition of 'justification'. One person's 'justification' may be another person's 'warrant'. Indeed, some philosophers are ready to abandon the idea that the word 'justification' picks out a property that is worth epistemological investigation. Alvin Plantinga, for instance, has argued that the early modern conception of justification rested on outmoded views of the nature of belief, and consequently that modern conceptions of justification do not capture an epistemologically

Pacific Philosophical Quarterly 86 (2005) 422–444

© 2005 University of Southern California and Blackwell Publishing Ltd. Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

important property.² These conceptions, Plantinga argues, are internalist, and there is no reason to value an internalist epistemic property.

Against Plantinga, I will sketch an outline of what justification might be and defend the epistemological importance of justification; but first I must say a word about what it is for a property to be epistemologically important. Epistemologists should worry about whether beliefs have a certain property only if that property is a good property for beliefs to have. If a property were not important in our evaluation of beliefs, there would be no point in worrying about which beliefs had that property. So the issue is whether the word 'justification' picks out a property that makes a belief in some respect a good belief.

In arguing that 'justification' can be used to pick out such a property, we will not arrive at a conception that is detailed enough to settle all questions about which beliefs are justified. For instance, it will not by itself settle whether we are justified in believing in the existence of external objects. Epistemological work will still need to be done in order to fill in details of the conception of justification and to delimit its extension. What I will argue is that, whatever justification is, it is what I call *meta-reliabilist*: A subject is justified in believing something when she has come to believe it by a method that she would be justified in believing reliable. Meta-reliabilism is enough to distinguish justification from the competing alternative of reliabilist warrant, on which a belief is warranted when it is obtained by a method that is *in fact* reliable. The argument for the epistemological importance of justification will allow us to counter an argument by William Alston that (an analogue of) reliabilist warrant is more epistemologically important than (an analogue of) meta-reliabilist justification.³

I will argue for meta-reliabilism by arguing for the importance of a conception of justification that is completely supervenient on the subject's experience, or (as I call it) completely experience-dependent. The nature of experience is a question beyond the scope of this essay; determining what experience is part of fleshing out the conception of justification. For instance, internalism as Plantinga describes it is a special case of complete experience-dependence. Only if experience itself is completely internal will justification be completely determined by our internal states. Nevertheless, my argument will be enough to demonstrate the epistemological importance of justification conceived as depending only on experience, as opposed to reliabilist warrant, which depends on factors that are not only external but completely beyond the believer's experience.

An epistemologically important property, as remarked above, is one that is important to our evaluation of beliefs. More controversially, I will take an epistemologically important property to be one that is important to our evaluation of beliefs *qua* beliefs. In particular, the property should provide a yardstick for evaluating beliefs rather than the actions that the believer took in arriving at those beliefs. A stupid person and a lazy

person may each arrive at beliefs that go against the available evidence; the stupid person because he is unable to evaluate the evidence properly, the lazy person because he does not bother to evaluate the evidence. Only the lazy person's actions can be criticized, because the stupid person cannot help believing what he believes. Nevertheless, each person's belief is equally bad as a belief. Epistemologically important properties will help with this sort of evaluation of beliefs rather than the actions that produce them.

When evaluating beliefs as beliefs, we should begin from the idea that beliefs aim at truth.⁴ When a belief is true, it is good as a belief, and when it is false, it is bad as a belief. Sometimes it may be conducive to a subject's other goals to believe something that is false, as when she harnesses the power of positive thinking, but this belief will not be good from the epistemic point of view. From the epistemic point of view, reliabilist warrant might seem clearly superior to meta-reliabilist justification. A warranted belief, arrived at by a procedure that is in fact reliable, is more likely to be true than a justified belief, arrived at by a procedure that may only be justifiably believed to be reliable. Indeed, this is Alston's argument for the importance of reliabilist warrant over meta-reliabilist justification.⁵ An analogy with advice, however, will establish that a completely experience-dependent property is important from the epistemic point of view. This completely experience-dependent property, I will show, is meta-reliabilist justification.

2. Definitions

In the face of varying uses of terms like 'justification', it is best to begin with an exact statement of definitions. After Plantinga, I will call justification's competitor 'warrant'.⁶ *Reliabilist warrant* will be defined as follows:

The belief that p is warranted for a subject S iff:

- (i) S arrived at the belief (or sustains the belief) by a procedure P ,
- (ii) P is in fact reliable, and,
- (iii) this warrant is not overridden.

The main idea of this conception of warrant is that whether a belief is warranted for S depends on the actual reliability of the procedure P , whether or not S has any reason to believe that P is reliable. Two points about the definition: For brevity's sake, I treat subjects as time-slices; this device could be avoided by quantifying over times as well as subjects throughout. Clause (iii), that the warrant is not overridden, entails that there is not an even more reliable procedure that would lead to the opposite belief, and rules out other similar possibilities.⁷

The difference between reliabilist warrant and the conception of justification that I will defend, *meta-reliabilist justification*, is that meta-reliabilist justification depends not on whether P is reliable but on whether S would be justified in believing that P to be reliable:

The belief that p is justified for a subject S iff:

- (i) S arrived at the belief (or sustains the belief) by a procedure P ,
- (ii) if S believed that P is reliable, that belief would be justified, and,
- (iii) this justification is not overridden.

Again, clause (iii) rules out the availability of a procedure that S would be justified in believing even more reliable and that would produce the opposite belief. (These caveats about overrides will be taken as implicit throughout most of this essay.) The definition of meta-reliabilist justification is circular, but this circularity will eventually be resolved. In section 6, I will defend a completely experience-dependent conception of justification, yielding a non-circular definition of justification. Then, in section 7, I will demonstrate that the property defined by this non-circular definition satisfies an analogue (stated just below) of the above definition. Thus justification as I define it will turn out to be meta-reliable.

The above definitions refer to the procedure that leads to or sustains a belief. Whether a belief that p is justified depends on the origin of the belief. Considering a belief's actual origin, however, will complicate our analysis excessively. Accordingly, in this essay I will abstract from the actual origin of beliefs; when I evaluate a belief that p , the evaluation will depend on p rather than on how the believer came to believe that p . Given this abstraction, we will be concerned with whether a subject's belief has certain *propositional properties*, where a propositional property will be defined as any function from subjects to sets of propositions. If S is a subject and D is a propositional property, then $D(S)$ is a set of propositions, and S 's belief that p is said to have the property D if and only if p is in $D(S)$.

As defined, propositional properties capture any quality whatsoever that can be ascribed to contents of beliefs, but many of these qualities will have no particular epistemological importance. 'Being about the Matterhorn', 'being most naturally expressed in English by sentences beginning with "G"', and many arbitrary functions are all propositional properties, but none of them are good or bad properties for beliefs to have. This essay will concentrate on evaluating the importance of the propositional properties that correspond to reliabilist warrant and meta-reliabilist justification:

The belief that p is warranted for a subject S iff there exists some procedure that is in fact reliable whose application would have led S to arrive at the belief that p .

The belief that p is justified for a subject S iff there exists some procedure whose application would have led S to arrive at the belief that p and that S would be justified in believing to be reliable.

(More precisely, these propositional properties are 'is potentially warranted' and 'is justifiable'; actual warrant and justification depend on the actual origin of the belief and are thus not propositional properties as I have defined them. To save space, I will continue to refer to the propositional properties as 'is warranted' and 'is justified'.) Again, the definition of meta-reliabilist justification is circular, but I will argue for a non-circular definition of justification that will be shown (in section 7) to yield a property that satisfies the above definition.

The last term to be defined is 'experience-dependence'. The details of experience-dependence will depend on what experience is, of course. I will not attempt to define experience exactly or to choose among competing conceptions of experience. As a matter of common sense, experience is the means by which we learn about the world, if there is a world; for instance, I have an experience of some sort when I look at a clock that reads 6 p.m., and it is this experience that leads me to believe that it is 6 p.m. Exactly what this experience consists in is a further debate. Some philosophers may take it that experiences are defined in terms of their phenomenal qualities, so that I would have exactly the same experience if I saw the clock or if I experienced a qualitatively identical hallucination of the clock. Other philosophers might take seeing the clock and hallucinating the clock to be different experiences, even if they are qualitatively identical. I will not choose between these or other conceptions of experience; we can give an outline of what justification is like without settling the details of what experience consists in. I will assume only a weak restriction on the notion of experience: Two subjects have the same experience if their experiences are qualitatively identical, they have the same background knowledge, and their immediately visible (or otherwise sensed) physical surroundings are the same.⁸ For instance, looking at a stopped clock and looking at an otherwise identical working clock will count as the same experience, because the physical differences are in the hidden workings of the clock.

We can take the experiences that a subject has had as a primitive. A *completely experience-dependent* propositional property can then be defined as one that supervenes on the experiences that the subject has had:

P is completely experience-dependent iff, whenever two subjects S and T have had exactly the same experiences, the set of propositions that are P for S is exactly the same as the set of propositions that are P for T .⁹

For instance, if experiences consist of sense-data that can be described propositionally such as "Red to the left now," the property 'being a

sense-data statement that is or has been given in experience' is completely experience-dependent; if "Blue to the right back then" has been given in one subject's experience and not in another's, the two subjects must have different experiences. Note that exactly which properties are completely experience-dependent depends on what experience is. The details of which beliefs are justified will depend on what experience is, too. We will, however, be able to demonstrate that justification is completely dependent on experience, whatever experience may be; this establishes what justification must be like in outline, even if the details depend in part on our conception of experience.

Some propositional properties supervene on the facts of the world and of the time at which they are evaluated, so that the world determines which propositions have this property for a particular subject, regardless of the subject's experience.¹⁰ I will call such properties *experience-independent*, which can be defined in terms of possible worlds:

P is completely experience-independent iff, whenever two subjects *S* and *T* occupy the same world at the same time, the set of propositions that are *P* for *S* is exactly the same as the set of propositions that are *P* for *T*.

Truth is a completely experience-independent property, if we exclude first-person indexical propositions. The propositions that are true for a subject are determined by the world and time that the subject occupies, rather than the subject's particular experiences.

Reliabilist warrant, on the other hand, is neither completely experience-dependent nor completely experience-independent. The output of procedures for arriving at beliefs will depend on the subject's experiences as well as on the facts of the world, so warrant is not completely experience-independent. The actual reliability of the procedures that produce beliefs will depend on the facts of the world as well as on the subject's experiences, so warrant is not completely experience-dependent. Consider three subjects: *S* has just looked at a working clock that reads six o'clock, *T* is in the same world as *S* but cannot see the clock, and *U* (in a different world) has had exactly the same experiences as *S*, but the clock she has just looked at is stopped; let us assume that *T* has no way of getting a look at the clock, and *S* and *U* have no way of telling the time other than by looking at the clock.¹¹ *S* is warranted in believing that it is six o'clock, because she can arrive at that belief by the reliable method of checking that clock. *T*, in the same world, is not warranted in believing that it is six o'clock, because her experiences do not allow her to use that method, nor any other reliable method. *U*, with the same experiences, can use the method of looking at the clock, but the method is not reliable for her because the clock is stopped; so she is not warranted in believing that it

is six o'clock. The difference between *S* and *T* shows that warrant is not completely experience-independent; the difference between *S* and *U* shows that warrant is not completely experience-dependent.

I will defend a completely experience-dependent conception of meta-reliabilist justification. Note that the difference between *S* and *U* is that the method *U* uses is in fact unreliable. Nothing yet said determines whether *U* (or, for that matter, *S*) is justified in believing that her method is reliable. So meta-reliabilist justification, for all we have said, may not discriminate between *S* and *U*. It may still be that they both, sharing the same experiences, are justified in believing that it is six o'clock. To show that meta-reliabilist justification indeed is completely experience-dependent, I will argue that there is a completely experience-dependent property that is epistemologically important, and then show that that property is meta-reliabilist justification. First, however, we should examine Alston's argument for the superiority of reliabilist warrant over meta-reliabilist justification, and Plantinga's arguments for the unimportance of justification in general. These arguments will show why the importance of a completely experience-dependent property needs defending.

3. *Criticisms of justification*

The properties that Alston considers are not quite the propositional properties of reliabilist warrant and meta-reliabilist justification. Alston considers the actual grounds on which the subject believes a proposition, as opposed to the proposition that the subject believes. So for Alston positive epistemic status requires not only that there exist a reliable method that could produce the belief, but also that the belief actually was produced by that method. Some such consideration is probably necessary for a thorough account of justification, but it will be easier to get clear on justification if we ignore the etiology of beliefs and focus on the proposition believed. Accordingly, I will consider Alston's arguments as they would apply to reliabilist warrant and meta-reliabilist justification, as I have defined them.¹²

As will I, Alston starts from the idea that, from the epistemic point of view, true beliefs are good and false ones bad. He then argues that, from the epistemic point of view, it is better to believe something on grounds that are in fact reliable than to believe it on grounds that one is justified in believing reliable:

[Positive epistemic status] depends . . . on whether [S's] believing that *p* is a *good thing* from an epistemic point of view. And however justifiably S believes that his grounds are adequate, if they are not then his believing that *p* is not a good move in the truth-seeking game. Even if he isn't to blame for making that move it is a bad move nonetheless (p. 74).¹³

If *S* is (in my terms) warranted but not justified in believing that *p*, then there is some reliable method that would lead him to believe that *p*, even though he would not be justified in believing that method reliable. If *S* is justified but not warranted in believing that *p*, then there is no reliable method that would lead him to believe that *p*, even though there is such a method that he is justified in believing to be reliable. Alston's point is that, in these cases, one is more likely to attain truth by believing the warranted proposition than the justified proposition. A belief that could be reached by a method that is in fact reliable is likely to be true. Accordingly, evaluating belief by the standard of truth-conduciveness, warrant would seem to be more epistemologically important than justification.

Alston's argument is in danger of proving too much, however. As he asks, "If goodness from an epistemic point of view is what we are interested in, why shouldn't we identify [positive epistemic status] with truth, at least extensionally?" (p. 70).¹⁴ One is even more likely to attain the truth if one believes what is true than if one believes what is warranted. If truth-conduciveness were the sole test for the epistemological importance of a propositional property, then truth would seem to be the only important property. As warrant is superior to justification, so is truth superior to warrant.

Alston's answer to his question is that 'justification' must be internalist in character.¹⁵ What he calls justification (and what I call warrant) depends at least partly on the subject's perspective; in my terminology, it is not completely experience-independent. Plantinga, however, casts doubt on the idea that internalism can make 'justification', even as Alston conceives it, into an epistemologically important concept. The seeming importance of internalism, Plantinga argues, is a historical relic.

As Plantinga tells it, Descartes and Locke held the deontological conception of justification: that it is literally a duty to believe only what was justified, and that unjustified beliefs arise from a misuse of free will in affirming a proposition that we should not. We could be blamed for an unjustified belief in exactly the same way that we could be blamed for an unjustified action. This presupposes that belief is under our voluntary control. The deontological conception also entails that justification is internalist, in that it "depends only on states, like experience or belief, that are in a recognizable if hard to characterize sense internal to the believer" (Plantinga, p. 49).¹⁶ To sketch Plantinga's account of this entailment, in order to be held accountable for unjustified beliefs the believer would have to be able to know with certainty whether or not a belief was justified. She can know with certainty only the internal states of belief and experience, so justification must depend entirely on these internal states (Plantinga, pp. 58–9).

Plantinga points out that the deontological conception of justification does not capture an epistemologically important property. Beliefs do not generally arise from a direct exercise of free will, so it is not appropriate

to ask whether someone is obeying her duty in believing that *p*. Many of our actions indeed affect beliefs; not only can we attempt to convince ourselves of something, we can choose to gather evidence that influences our beliefs. Nevertheless, evaluating the actions that influenced a belief will not produce an appropriate evaluation of the belief as a belief. For instance, when some cognitive incapacity means that a person cannot help but believe that *p* against the evidence, she will believe that *p* even if she obeys all her duties. This belief would be justified according to the deontological conception, and indeed the believer cannot be blamed for it, but the belief itself is obviously epistemically poor.

Plantinga concludes that there is no reason to think that an internalist property will turn out to be epistemologically important. The only reason for “widespread intuitions favoring an internalist requirement” on justification (p. 67)¹⁷ is that most people are still in the grip of a classical deontological conception of justification; “[c]ut off the deontology, and the internalism looks like an arbitrary appendage” (p. 68). Alston’s conception, Plantinga claims, may be the “closest coherent conception” (p. 68) to the contemporary use of ‘justification’, but that does not make it epistemologically important. We should focus instead on warrant, “that quantity enough of which is sufficient, with truth, for knowledge” (p. 51). On a reliabilist notion of knowledge, this quantity would be reliabilist warrant.¹⁸ Certainly reliabilist warrant is a more externalist concept than meta-reliabilist justification. If Plantinga’s argument is correct, we should doubt the epistemological importance of meta-reliabilist justification.

Justification’s epistemological importance, however, need not be rooted in deontology. In what follows I will present an account of what it is for a propositional property to be important in evaluating beliefs that does not presuppose that beliefs are evaluated deontologically. This account will be based on an analogy with advice on how to achieve a specific goal. Sometimes it is good to advise a person to do something that she can achieve with certainty, even if doing this will not guarantee achievement of her goal. Analogously, from the epistemic point of view, it may be good to base our beliefs on our experiences, even if in doing so we may fall short of the epistemic goal of believing truths and avoiding falsehoods. This account of epistemic goodness will prove the epistemological importance of a certain completely experience-dependent property.¹⁹ I will then show that this completely experience-dependent property is the property of meta-reliabilist justification.

4. *Helpful advice*

First we need an account of advice that is directed to some well-defined goal.²⁰ Ideally, what we advise someone to do will be within her power to

accomplish, and ideally, accomplishing what is advised will be enough to guarantee achievement of the goal. But, because people have limited capacities, advice may not be able to achieve both these ideals. There may be nothing that the advisee can accomplish with certainty and that will also guarantee achievement of her goal. Call advice completely *effective* when successfully doing what is advised is sufficient for achieving the goal, and completely *doable* when the advisee is capable of guaranteeing success at an attempt to do what is advised. We can also define comparative degrees of these qualities: The more likely it is that successfully doing what is advised will suffice for achieving the goal, the more effective the advice is; the more likely it is that the advisee will succeed if she attempts to do what is advised, the more doable the advice is. In some situations an advisor will not be able to give advice that is both completely effective and completely doable. She must choose one desideratum or the other, or balance the two.

For example, suppose we are advising Alice on how to bowl a strike. Like most bowlers, Alice is not capable of knocking down all the pins at will, nor is she capable of hitting the right side of the head pin at will, but she is capable of *aiming* the ball at the right side of the head pin. We could advise, "Knock down all the pins." This advice will be completely effective, because Alice will certainly bowl a strike if she knocks down all the pins; but it is not completely doable, because she cannot guarantee success at an attempt to knock all the pins down. Another piece of advice, "Aim at the right side of the head pin," is completely doable, because Alice can guarantee that she *aims* the ball there; but the advice is not completely effective, because aiming the ball at the right side of the head pin is not sufficient for bowling a strike (or even for hitting the right side of the head pin).

An intermediate case is the advice, "Hit the right side of the head pin." This is not completely doable, because Alice cannot guarantee that she will hit the pin; she can only aim and hope. Nor is it completely effective, because hitting the right side of the head pin is not enough to guarantee a strike. It is, however, more effective than "Aim at the right side of the head pin," because Alice is more likely to knock down all the pins if she does hit the right side of the head pin than if she merely aims there (possibly hitting it). On the other hand, it is more doable than "Knock down all the pins," because Alice is more likely to succeed at an attempt to hit the right side of the head pin than at an attempt to knock down all the pins. "Hit the right side of the head pin" can also be compared with "Hit the left side of the head pin." "Hit the left side of the head pin," we may suppose, is exactly as doable as "Hit the right side of the head pin"; Alice is equally likely to succeed no matter which side of the pin she attempts to hit. "Hit the left side," however, is less effective than "hit the right side," because Alice is less likely to bowl a strike if she hits the left side than if she hits the right side.

Neither effectiveness nor doability should be taken to be the primary desideratum on advice; they are simply two different criteria for evaluating advice. Depending on the circumstances and the purpose of the advice, it may be preferable to give completely effective advice that is not completely doable, or completely doable advice that is not completely effective, or advice that is neither completely doable nor completely effective. If, however, one piece of advice *dominates* another with respect to effectiveness and doability – it is more effective and at least as doable as the other, or more doable and at least as effective – then the dominant piece of advice is unequivocally better than the other. In our example, “Hit the right side of the head pin” dominates “Hit the left side of the head pin,” so there is no reason to advise Alice to hit the left side of the head pin. We may call advice *helpful* so long as there is no alternative advice that dominates it with respect to effectiveness and doability. In a given situation there may be many different pieces of helpful advice, some more doable and less effective than others, some more effective and less doable.

5. *Epistemic desiderata*

I have discussed advice in order to illuminate epistemologically important properties: properties that make beliefs good as beliefs. Specifically, I am concerned to show that being justified, on a completely experience-dependent conception of justification, makes a belief good from the epistemic point of view, and so that justification is epistemologically important. To do this I will draw an analogy between propositional properties and pieces of advice. Endorsing a propositional property *D* as epistemologically important is like telling all subjects, “You should believe all and only those propositions that are *D* for you.” We can then evaluate propositional properties according to epistemological analogues of effectiveness and doability.

We are interested in evaluating beliefs only with respect to the purely epistemic goals of attaining truth and avoiding falsehood. The ways in which true beliefs contribute to our practical goals and false beliefs hinder them does not concern us; from the epistemic point of view, attaining truth and avoiding falsehood are ends in themselves.²¹ In addition, it is worth repeating that we are interested in evaluating beliefs as beliefs, not in evaluating the actions that lead to the beliefs. Our analysis of epistemologically important properties will reveal the properties that it is good for beliefs to have; the best way for a thinker to deliberate so as to arrive at beliefs with these good properties is a separate question, beyond our purview.²²

“Believe all and only the propositions that are *D* for you” is not literally advice. As discussed in section 2, a belief is not an exercise of voluntary control, and so this principle does not literally advise an act. Accordingly,

effectiveness and doability as defined for advice cannot be literally applied to propositional properties. “Believe all and only the propositions that are D for you,” however, is enough like advice that it can be evaluated according to analogues of effectiveness and doability. Like advice, this principle is directed toward specific goals, those of believing truths and not believing falsehoods. The epistemic analogues to effectiveness and doability will be based on how a propositional property contributes to the fulfillment of these goals.

The epistemological analogue to complete effectiveness is fairly straightforward. A completely effective epistemic principle, if followed, will guarantee achievement of the goal of believing truths and not falsehoods. So the property D will be a *completely effective propositional property* iff any subject who succeeds in following “Believe all and only the propositions that are D for you” succeeds in believing every proposition that is true and no proposition that is false in her world and at her time. (Recall that a subject is a time-slice, so specifying the subject means specifying the world and time.) Thus, if D is completely effective, the propositions that are D for a subject are all and only the propositions that are true for her.

Defining degrees of effectiveness for propositional properties is less straightforward. We can begin with an analogy to effective advice. The advice “Do A ” is comparatively effective to the extent that it is likely that, if the advisee succeeds in doing A , she will attain her goal. The epistemic goals are to believe truths and avoid believing falsehoods. So the propositional property D is comparatively effective to the extent that it is likely a subject who believes a proposition that is D will believe a truth, and that it is likely that a subject who does not believe a proposition that is not D will avoid believing a falsehood.

This fails to yield a precise and general definition of comparative effectiveness, because it is not clear how to define the likelihood that a subject who believes a proposition that is D will believe a truth. The degree of effectiveness of D is meant to depend only on D , not on any particular subject or proposition. Accordingly, we must quantify over subjects and propositions: We seek the likelihood that, if we randomly choose a subject S and a proposition p , and p is D for S , p will be true for S (and also the likelihood that, if we randomly choose a subject S and a proposition p , and p is not D for S , p will be false for S). Since different subjects are in different worlds, this measure of likelihood must range across all worlds as well. In particular, it cannot depend on any contingent facts. This means that we cannot rely solely on objective propensity or frequency of truth interpretations of probability.

For instance, if we fix a world that contains a loaded die, it may be that, as a matter of the die’s objective propensity, there is a two in three likelihood that the next throw of the die will come up six. This propensity, however, will not tell us anything about the likelihood that a randomly

chosen subject inhabits a world in which the next throw of a die will come in six, because there is no guarantee that the dice in the world of the randomly chosen subject are loaded in the same way. Objective propensity tells us about likelihood across a restricted set of worlds. If we consider only worlds containing dice that are loaded in this way, there is a two in three likelihood that a subject randomly chosen from one of those worlds inhabits a world in which the next throw of that die comes up six.

It is difficult to come up with a notion of likelihood that would allow us to calculate, for an arbitrary propositional property D , how likely it is for random S and p that, if p is D for S , p is true for S (*mutatis mutandis* for propositions that are not D and false). Fortunately, I will show (in section 6) that we do not need to determine degrees of effectiveness for arbitrary propositional properties. To demonstrate the epistemological importance of justification, we will only need to consider degrees of effectiveness for the completely experience-dependent properties. This can be done by considering certain restricted sets of worlds that can be measured in terms of likelihood on the available information, much as the restricted set of worlds containing dice loaded in a certain way can be measured in terms of the dice's objective propensities. Accordingly, we can hold to our original definition of degrees of effectiveness: A propositional property D is effective to the extent that it is likely a subject who believes a proposition that is D will believe a truth, and that it is likely that a subject who does not believe a proposition that is not D will avoid believing a falsehood. Questions about how to apply the definition will be addressed in section 6 when we do apply it.

The effectiveness of a principle measures how conducive it is to truth. If effectiveness were the only desideratum for epistemic principles, then positive epistemic status would collapse to truth; this is Alston's problem as discussed in section 3. Accordingly, we need to define another desideratum, an epistemological analogue of doability; but this is much less straightforward than the epistemological analogue of effectiveness. Doability for advice was defined in terms of the advisee's ability to guarantee the success of an attempt to do what was advised. But beliefs are not exercises of voluntary control. Accordingly, it would be odd to consider whether a subject who attempts to believe a proposition can guarantee that she succeeds in believing the proposition. An epistemological analogue of doability will require a notion of the believer's capability of following an epistemic principle other than her ability to succeed at an attempt to follow it.

Cognitive capacity might seem appealing here; it is a sort of capability whose exercise does not require voluntary control. For instance, a subject may not choose whether to believe a proof, but for the proof to convince she must be capable of following it. Using cognitive capacity in an analogue of doability, however, would lead to problems like those faced by the

deontological conception of justification. If a desideratum of an epistemic principle is that it never exceed a subject's cognitive capacity, then a belief can attain positive epistemic status because the subject is incapable of believing anything else. This goes against the idea that epistemic status evaluates beliefs rather than the actions that produce the beliefs. Someone who believes something through cognitive capacity is personally blameless, but the belief itself is not as good as it would be if the believer had greater capacity.

A more useful epistemic analogue to doability, which I will call *knowability*, adverts to the information that the subject's experiences make available. The available information sets epistemic limits for the subject, as the subject's capacities set the practical limits that pertain to advice. Accordingly, we will define a *completely knowable propositional property* as one that supervenes on the available information: If *D* is completely knowable, and the same information is available to subjects *A* and *B*, then exactly the same propositions are *D* for *A* and for *B*. (In this essay we will not need to measure comparative degrees of knowability.)

What is it for an experience to make information available? Many answers are possible, and I will not plump for a specific one. For instance, I will not take a position on whether an experience that has been forgotten makes any information available; the answer would affect the details of what is justified, but not the overall shape of the argument.²³ Our account, however, must not beg epistemic questions or push them back a level. Take the proposal that an experience carries the information that *p* when it puts the subject in a position to know that *p*.²⁴ This would mean that our conception of knowledge will determine what properties are epistemically important. This begs the question as to whether knowledge is epistemically important. To bring out the question-begging character of this proposal, consider the analogous proposal in which the information carried by an experience is determined by what the experience justifies a subject in believing; this obviously would beg the question in favor of the importance of justification.

On the other hand, we must exclude a reliabilist notion of information such as Dretske's, on which all and only reliable processes transmit information, regardless of whether the subject has any reason to believe that her experience results from a reliable process.²⁵ This would push back epistemic questions to the question of how and whether the subject can know that certain information is available to her. I will take it that if two subjects have had the same experiences, the same information is available to them.

We may think of the available information as everything that a person with no cognitive incapacities would be able to use in coming to a belief, if she had had the same experiences. A propositional property that is not completely knowable may depend on information that the subject lacks.

No matter how well a subject reasoned, she would need guesswork to determine whether a certain belief had the property. This seems analogous to the flaw in undoable advice. We may fail to believe every proposition that has the unknowable property, through no failure of our cognitive capacities; while we may fail to follow undoable advice, through no failure of our will to follow it.

On the view that knowability is an epistemic desideratum, a belief may attain positive epistemic status because the believer has not had the experiences that would show it to be false. It may be objected that this does not reflect well on the belief as a belief, any more than cognitive incapacity would; a believer who lacks the appropriate experience is not to blame for believing as she does, but I have already argued with respect to cognitive incapacity that the person's freedom from blame does not make her belief a good one. In response, experience is what allows a subject to learn about the world. It reflects well on a belief if the belief is somehow constrained by the subject's experiences. A subject who lacks the appropriate experience is not incapacitated in the same way as one who cannot draw the appropriate conclusions from her experiences; her belief is inaccurate not because of her infirmities but because she lacks the basis for an accurate belief.²⁶ Accordingly, one of the epistemic desiderata must be a measure of experience-dependence. Otherwise the only epistemologically important property would be truth, which is completely effective though completely experience-independent.

Knowability and effectiveness, then, are the two desiderata for propositional properties, analogous to doability and effectiveness for advice. As doability and effectiveness determine whether advice is helpful (in our technical sense), so knowability and effectiveness determine whether a propositional property is worth endorsing as epistemologically important. If a property is to be epistemologically important, there can be no other propositional property that dominates it: that is more knowable and at least as effective, or vice versa. If a property were dominated, the dominating property would be unambiguously better than the original unimportant property as a guide to beliefs, and there would be no reason to evaluate beliefs according to whether they have the dominated property. In the next section, we will look at two propositional properties that are not dominated. One is truth; the other will turn out to be meta-reliabilist justification.

6. *Epistemologically important properties*

The epistemological importance of a propositional property depends on its effectiveness and knowability. The most knowable of all completely effective properties will be epistemologically important; no property can be more effective, and no equally effective property can be more knowable.

This property is truth. That truth is completely effective follows from the definition of a completely effective property; indeed, any completely effective propositional property must coincide extensionally with truth. So truth turns out to be epistemologically important, as it should.

Truth, however, completely satisfies only one of the two desiderata on propositional properties. It is completely effective, but it is not completely knowable, because which propositions are true is not completely determined by the information that the subject's experiences make available. Two subjects in different circumstances may have had exactly the same experiences and have exactly the same information available, yet different propositions may be true for those subjects. Consider two subjects, in two different worlds, each of whom has had identical experiences of living in the twentieth and twenty-first centuries; yet in one world Caesar shaved the morning he crossed the Rubicon and in the other he did not. The subjects have had the same experiences, yet the proposition "Caesar shaved the morning he crossed the Rubicon" is true for one and false for the other, proving that truth is not determined by the information available to the subject, and so is not completely knowable.

Thus, there is room for another propositional property to attain epistemological importance by being more knowable than truth, albeit less effective. Consider the propositional properties that are completely knowable; each is more knowable than truth is. Let us give the name *J* to the most effective of all these completely knowable properties. No other propositional property can be more knowable than *J*, because *J* is completely knowable; and any property that is more effective than *J* will be less knowable, because *J* is the most effective of all completely knowable properties. *J* thus is not dominated and meets the criterion for an epistemologically important property.

What is this property *J*; what is the most effective of all the completely knowable properties? In section 7 I will show that *J* meets the definition of meta-realist justification, but first we require a more intuitive description of *J*: A proposition is *J* for a subject if and only if it is likely on the information that is available to the subject. This description almost falls out of our definitions thus far: We have defined how effective a propositional property is by how likely propositions with that property are to be true, and which propositions are *J* for a subject is determined entirely by the subject's experiences, so the notion of likelihood at issue must be likelihood on the available information. Nevertheless, I will give a more detailed derivation of this description of *J*, in order to clarify exactly what notion of likelihood is at issue.

From the definition of a completely knowable property, any property *D* is completely knowable so long as, for any *p*, if the same information is available to two subjects, then either *p* is *D* for both subjects or *p* is not *D* for both subjects. The set of all subjects may be partitioned into what I

will call *information classes*, so that two subjects are in the same information class if and only if the same information is available to them. Then whether p is D for a subject is determined by the subject's information class. To find the most effective of all completely knowable properties, we may consider each proposition individually on each information class I , asking: Will D be more effective if p is D for every subject in I or if p is not D for every subject in I ? When we have determined the answer for every p and every I , we will have found the most effective of all completely knowable properties.

In section 5, we defined the effectiveness of a propositional property D in terms of the likelihood, for a randomly selected subject S and proposition p , that D gets p right for S : that p is D for S if it is true for S , and that p is not D for S if it is false for S . As discussed in section 5, this definition requires a notion of likelihood. We have now narrowed our focus to a single proposition p and information class I . In order to determine whether or not D is more effective if p is D for every S in I , we must ask: If we select a random subject S from I , how likely is it that p is true for S ? Since I contains every subject whose experiences have made certain information available, this likelihood is the likelihood of p on the information available in I . For example, the subjects in I may all have watched (or seemed to watch) a die being rolled many times, coming up six two-thirds of the time; now they are about to roll what seems to be the same die again. On the information available to them, there is close to a two-thirds chance that the die will come up six. Not every subject will be in a world in which the die has an objective propensity to come up six two-thirds of the time; in some worlds the die will be a fair one that has gone on an improbable run of sixes, while in others a fair die or an even more biased one may have been substituted for the die that was tested. Still, if one randomly selects a subject from the information class, there will be a two-thirds chance that that subject is in a world in which the die is about to come up six.

So if J is the most completely effective of all knowable proposition properties, then determining whether a proposition p is J for a subject S will require determining what information is available to S and determining whether p is likely on that information. Since J is a completely knowable property, a proposition that is J for a subject must be J for any other subject to whom the same information is available; so factors such as the objective propensities of the objects around the subject cannot affect J unless they are reflected in the available information.²⁷ This means that empirical investigation cannot help determine these likelihoods. Empirical investigation involves gathering new information by means of new experiences, the experience for instance of observing the outcome of some experiment. Though it may reveal whether a belief was truth-conducive for a particular subject who once was in a certain information class, it cannot reveal whether that belief would have been truth-conducive across

the entire information class, because the newly available information excludes some of the subjects who were in the original information class. Likelihood on the available information hence is to be determined by aprioristic methods. The question is, "How likely would this proposition be for someone who had exactly this information?"; not a question to be answered by gathering further information.

Note that this treatment pushes back some epistemological problems onto the question of how likely a certain proposition is on certain information. That will frequently be a question to be addressed by further argument about whether particular information makes a particular hypothesis likely. Consider Goodman's new riddle of induction.²⁸ When a subject has examined many emeralds and found them all to be green/grue, do her experiences make it overwhelmingly likely that all emeralds are green or that all emeralds are grue? Answering this question is philosophical work that goes beyond the question of what justification is. Similarly, scientists and philosophers of science may dispute whether or not certain experiments support a certain hypothesis even after the experiments have been done. The dispute could be framed indifferently as a dispute concerning whether the available information justifies belief in the hypothesis or about whether the available information makes the hypothesis likely. The conception of justification at issue here will not settle these questions; it only states that to be justified *is* to be likely on the available information. Nor does the conception of justification require that every hypothesis be assigned a definite probability on every set of experiences. All that is needed is judgments that certain experiences make certain hypotheses very likely, or fairly likely, or some such; then the hypotheses will be very or fairly justified on those experiences.

To return to *J*, the most effective completely knowable propositional property: Whether or not it is more effective for *p* to be *J* in information class *I* depends on how likely *p* is on the information available in *I*. Effectiveness measures conduciveness to two separate goals: attaining true beliefs and avoiding false ones. The definition of effectiveness given in section 5 did not say how these goals were to be weighted; let us pretend for the moment that avoiding falsehood is exactly three times as important as attaining truth. Then it will be more effective for *p* to be *J* in *I* than for *p* not to be *J* in *I* if and only if *p* has a probability of greater than three-fourths, on the information available in *I*. Of course, it is implausible that there is a single right answer to the question, "How important is attaining truth relative to avoiding falsehood?" The proper weighting will depend on non-epistemic factors such as the severity of the consequences of false belief. These non-epistemic factors will determine how likely a belief must be in order to count as justified. This is a necessary effect of treating justification as an all-or-nothing property when it is in fact a matter of degree. Nevertheless, this account indicates how to treat

justification as a matter of degree: The more likely a proposition is within a subject's information class, the more justified that belief would be for the subject.

We have defined justification as *J*, the most effective of all completely knowable properties; a belief is justified if it is likely on the available information. On this definition, justification is completely experience-dependent. What a subject is justified in believing depends on the information that her experience makes available; if two subjects have had exactly the same experiences, then the same information is available to them, and they are justified to exactly the same extent in believing exactly the same things. So we have established the epistemological importance of a completely experience-dependent propositional property, without relying on the kind of deontological considerations that Plantinga argues motivate the traditional conception of justification (see section 3). In the next section, I show that justification so defined is meta-reliabilist justification. We will then be prepared to answer the Alstonian argument (section 3) that reliabilist warrant is more important than meta-reliabilist justification.

7. *Justification vs. warrant*

Consider again the difference between meta-reliabilist justification and reliabilist warrant. A belief is warranted reliabilistically if it could be reached by a procedure that is in fact reliable (and there are no overrides); a belief is justified meta-reliabilistically if it could be reached by a procedure that the subject would be justified in thinking reliable (and there are no overrides). If a method is reliable but the subject would not be justified in thinking it so, then the beliefs it produces are warranted but not justified. When a method is unreliable but the subject would be justified in thinking it so, the beliefs it produces are justified but not warranted.

To show that *J*, the conception of justification sketched in the previous section, is the same as meta-reliabilist justification, I will show that the differences between *J*-justification and reliabilist warrant are the same as the differences between meta-reliabilist justification and reliabilist warrant. That is, I must show that a belief is warranted but not *J*-justified when the method that produced it is in fact reliable, but the subject would not be *J*-justified in believing it to be reliable. (The converse case, in which the belief is justified but not warranted, will be left to the reader.) This will establish the epistemological importance of meta-reliabilist justification and will also resolve the circularity of its definition (see section 2). *J*-justification is an important property defined in a non-circular way, and it satisfies the definition of meta-reliabilist justification: The belief that *p* is likely on the information available to a subject *S* iff there exists some procedure whose application would have led *S* to arrive at the belief that

p , it is likely on the information available to S that this procedure is reliable, and nothing would override justification obtained from this procedure.

Suppose, then, that the subject sees that her watch reads six o'clock, but unknown to her the watch stopped fifteen minutes ago. Her stopped watch no longer provides a reliable method for learning the time, so she is not warranted in believing that it is six o'clock. Her experiences, however, would have been more or less the same if her watch had not stopped. So on the information that her experiences make available, it is likely that her watch is working. If not, then the subject would not be J -justified in believing her watch to be reliable. This means that, for most of the subjects who have had the same experiences, their watches give the right time. Since one of these experiences is seeing that the watch reads six o'clock, it is indeed six o'clock for most of the subjects who have had the same experiences. This means that the subject is indeed J -justified in believing it to be six o'clock. Whenever a subject is J -justified in believing a method reliable (and there are no overrides), she will be J -justified in believing whatever results from the method.²⁹ Even if the method is in fact unreliable in the subject's world, that does not affect the J -justification of the beliefs it produces. We must consider all subjects, in any world, who have had the same experiences, and evaluate whether that method would be reliable for them.

We are now ready to meet the Alstonian argument for reliabilist warrant as opposed to meta-reliabilist justification. Alston argued that (in our terms) it is better from the epistemological point of view to believe what is warranted than to believe what is justified, because warranted beliefs are more likely to be true. This amounts to pointing out that warrant is a more effective property than justification. As we have seen, however, taking effectiveness as the sole criterion of epistemological importance would make truth the only epistemologically important propositional property. This is why, on Alston's argument, positive epistemic status threatened to collapse to truth. Alston postulated an internalism requirement to avert this collapse, but this left him vulnerable to Plantinga's objections to internalism, as discussed in section 3.

I have argued that knowability, as well as effectiveness, is a criterion of epistemological importance. We learn about the world through experience, so it is desirable that our beliefs should be governed by our experiences. The knowability of a property captures this desirability, by measuring the extent to which the property depends on experience. Epistemological importance requires not only truth-conduciveness, as measured by effectiveness, but also dependence on experience, as measured by knowability. This does not require a deontological conception of justification (or any other knowable property), so it is not vulnerable to Plantinga's criticisms of internalism.

The joint importance of effectiveness and knowability averts the collapse of positive epistemic status to truth. Truth is not the only epistemologically

important property, because truth is not completely knowable. But once we take knowability into account, Alston's argument for warrant over justification no longer goes through. Warrant is superior to justification with respect to effectiveness, but inferior with respect to knowability. The epistemic principle "Believe p if and only if it is warranted for you" is not unequivocally better than the principle "Believe p if and only if it is justified for you." If you succeed in believing what is warranted, you will be more likely to attain the truth than if you believe what is justified. Believing what is justified, however, requires no information other than what your experience has made available, while believing what is warranted will require guesswork. It is this complete experience-dependence that makes justification epistemologically important.³⁰

Department of Philosophy
Texas Tech University

NOTES

¹ Edmund Gettier (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, pp. 121–3.

² Alvin Plantinga (1990) "Justification in the 20th Century," *Philosophy and Phenomenological Research* 50, pp. 45–71. All subsequent references to Plantinga are to this article.

³ William P. Alston (1985) "Concepts of Epistemic Justification," *Monist* 68, pp. 57–89. Except where noted, subsequent references to Alston are to this article.

⁴ Pragmatists such as Rorty might object to the focus on truth as the main criterion for evaluating belief. On the other hand, it is hard to see how such a pragmatist could object to the completely experience-dependent notion of justification that I will defend, since there will be nothing outside the subject's experience that could serve as a check on belief. Other views (see e.g. Timothy Williamson (2000) *Knowledge and Its Limits*, Oxford: Oxford University Press) might take knowledge to be the primary goal of belief, but to do so would beg the question against the importance of justification.

⁵ Alston uses 'justification' for the concept that I am calling 'reliabilist warrant'; I have changed the terminology for clarity's sake. In later writings, Alston suggests the abandonment of the term 'justification', though not for the same reasons as Plantinga. See William P. Alston (1993) "Epistemic Desiderata," *Philosophy and Phenomenological Research* 53, pp. 527–551.

⁶ I do not mean to imply that Plantinga must take reliabilist warrant as justification's competitor. See the discussion of Plantinga on warrant versus justification, in section 3.

⁷ Here and throughout, I will ignore the generality problem: that it is not always possible to pick out the method by which the subject arrives at a belief. The conception of justification that I eventually defend (in section 6) does not depend on individuating methods of arriving against belief; if the generality problem is insoluble, that tells in favor of this conception of justification as opposed to reliabilist warrant.

⁸ This must be a sufficient condition for sameness of experience, but on most accounts it will not be a necessary condition. For instance, even if we count qualitatively identical veridical perceptions and hallucinations as different experiences, we may wish to allow that two people with different background knowledge who see the same object can have the same experience.

⁹ In section 5 we will consider the possibility that justification supervenes on a subset of experience rather than on all a subject's experience. This possibility would require a technical change in the definition of complete experience-dependence, but it will not make a difference to the spirit of the concept. See note 23, below.

¹⁰ Throughout this discussion I use 'determined' in the sense in which the parameters of a function determine its value, rather than in the sense of causal determination.

¹¹ Thanks to the referee for pointing out an unclarity in the original statement of this example.

¹² Alston makes clear, however, that his counterpart to reliabilist warrant (which he calls 'justification') cannot be identified with reliability, because of the importance of the actual grounds of a belief. See Alston, p. 79.

¹³ I have changed Alston's J_e for 'evaluative justification', to 'positive epistemic status'. This avoids terminological confusion between Alston's use of 'justification' and mine, which is reserved for meta-reliabilist justification.

¹⁴ As above (see note 13), I have replaced 'justification' with 'positive epistemic status'.

¹⁵ "Concepts of Epistemic Justification," p. 71; see also William P. Alston (1988) "An Internalist Externalism," *Synthese* 74, pp. 265–283.

¹⁶ As discussed in section 2, I am not committed to the idea that experiences are internal states of the believer. If they are not, then experience-dependent justification will be completely experience-dependent but not internalist in Plantinga's sense.

¹⁷ See also Alston, "An Internalist Externalism," p. 272.

¹⁸ Plantinga's criticism of Alston's conception of 'justification' has to do with Alston's insistence that adequate grounds, which Alston identifies with a reliable belief-producing method, must somehow be accessible to the believer. It is not clear that Plantinga would offer similar criticisms against what I describe as reliabilist warrant, which is what Alston's conception of justification becomes when we abstract from the grounds of the belief.

¹⁹ Whether complete experience-dependence entails internalism in Plantinga's sense depends on whether we conceive experience as internal. See note 16, above, and the discussion of different conceptions of experience in section 2.

²⁰ This is as opposed to advice concerning what to do *tout court*, without reference to one particular goal. The kind of advice that we are concerned with could be given in response to the question, "What should I do to get this stain out of my pants?" The kind of advice that we are not concerned with might be given in response to the question "What should I do with my life?"

²¹ Indeed, from other points of view it may be advantageous to avoid certain true beliefs and attain certain false beliefs. Someone who does not know a dictator's dirty secrets may be safer.

²² This is not to say that the process of belief formation is entirely irrelevant from the epistemic point of view. The way in which a belief is formed, which will involve involuntary processes as much as voluntary actions, may be evaluated from the epistemic point of view for its contribution to truth and falsehood. It is only for simplicity's sake that we are ignoring the process of belief formation to concentrate on the proposition that is believed; see the discussion of propositional properties in section 2.

²³ If an experience that has been forgotten does not make any information available, then two subjects who have had the same information may have different information available, if one has forgotten some of the experiences. Justification would then turn out to depend on the experiences that the subject has not forgotten rather than on the subject's experiences *tout court*. This would require a change in the definition of experience-dependence, but it would not affect the spirit of this account of justification. See note 9 above.

²⁴ This conception resembles Timothy Williamson's position that our evidence is what we know. See Timothy Williamson (1997) "Knowledge as Evidence," *Mind* 106, pp. 717–741, and Chapter 9 of *Knowledge and Its Limits*.

²⁵ See Fred I. Dretske (1980) *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press. For Dretske the process in question must be infallibly reliable, given appropriate background conditions, but we must also exclude accounts on which information that *p* can be transmitted by less than perfectly reliable processes.

²⁶ Someone who believes something because she has failed to gather the evidence that would disprove her belief may be culpable for her lack of experience. This reflects poorly on the actions that led up to her belief, not necessarily on the belief itself. To determine when someone is epistemically responsible for gathering more evidence, we would have to consider the costs of gathering that evidence, and this will depend on practical as well as purely epistemic considerations.

²⁷ The claim here is not that a property that takes into account the objective propensities of objects around you will be less effective than *J*, which depends on an aprioristic notion of likelihood on the available information. On the contrary, proportioning one's beliefs to objective propensities is more truth-conducive than proportioning it to likelihood on the available information. But objective propensities are not completely knowable; they are not determined by the available information. Any completely knowable property will depend on the available information only, and not on further empirical investigation. Aprioristic likelihood on the available information is the most effective of *these* properties. (Thanks to the referee for pressing me on this point.)

²⁸ Nelson Goodman (1983) *Fact, Fiction, and Forecast*, Cambridge, MA: Harvard University Press, 4th edn. An object is *grue* if it is green and observed before time *t* (say, January 1, 2050), or blue and observed after time *t*.

²⁹ If the speaker's experience makes it likely with probability *x* that her method is reliable, and if the reliability in question is less than 100%, then the speaker will not be as justified in believing the output of the method as she will in believing that the method is reliable. This will be a feature of any account of meta-reliabilist justification: Multiplying justifications weakens them.

³⁰ Earlier versions of this paper were presented at the University of Pittsburgh and the University of Wisconsin, Milwaukee. Thanks to the audiences; Mark Kaplan and Kieran Setiya in particular raised helpful questions, though I am certain I am overlooking other audience members who deserve thanks; my apologies. Nuel Belnap was particularly helpful in many ways. Thanks also to an anonymous referee for *Pacific Philosophical Quarterly*, who made many helpful suggestions over and above the ones noted in the text.