SURVEY ARTICLE

Highlights of Recent Epistemology
James Pryor

ABSTRACT

This article surveys work in epistemology since the mid-1980s. It focuses on (i) contextualism about knowledge attributions, (ii) modest forms of foundationalism, and (iii) the internalism/externalism debate and its connections to the ethics of belief.

The period from the mid-1970s through the mid-1980s was a period of bold new positions in epistemology. We can group the major developments of that period under three headings: (i) relevant alternatives theories of knowledge came on the scene; (ii) foundationalism made a come-back, and the debate between foundationalists and coherentists intensified; and (iii) we saw the development of externalist theories of knowledge and justification, most prominently Goldman's reliabilist account of justification ([1979], [1986], Chs 4–5) and Nozick's account of knowledge ([1981], Ch. 3).[1]

This paper surveys what has been happening in epistemology *since* those developments. I will not attempt to address *all* the important recent work in epistemology; but will focus on the issues that seem to me to have occupied

---

[1] I won't be able to discuss Nozick in this survey, but his account of knowledge has been an important topic during the past twenty years. (See Forbes [1984]; McGinn [1984]; Goldman [1986], Ch. 3; Luper-Foy [1987]; Plantinga [1988]; Fumerton [1995], Ch. 4; and DeRose [1995] for some of the critical literature.) Kripke has a well-known discussion of Nozick's account of knowledge, but unfortunately this remains unpublished.

center stage. We can group those issues under three headings, too, corresponding to the three developments cited above: (i) contextualist theories of knowledge came on the scene and are now starting to be widely discussed; (ii) modest forms of foundationalism have become more familiar, and are now widely believed to be viable; and (iii) epistemologists have begun more vigorously to explore issues in the ethics of belief, and their connection to the internalism/externalism debate. I will discuss each of these in turn, and will devote the most attention to issue (iii).

I should note one important omission: I will say little about 'naturalized epistemology' in this survey, though that has received a good deal of attention in recent years. In part, this is because it is hard to find a single set of theses one can claim to be definitive of naturalism. But largely it is because of space constraints. Since naturalism has been well-surveyed elsewhere (see Maffie [1990b]; Kitcher [1992]; Goldman [1994]; and Kornblith [1995]), I thought it would be better to concentrate on the issues I do take up below.[2]

1 Contextualism

1.1 Contextualism and relevant alternatives theories

We'll understand *Contextualism* to be the doctrine that the truth of a knowledge-ascription can vary with the context in which that ascription is made or assessed. For example, according to the Contextualist, the sentence 'Reagan knew that the Iran/Contra dealings were taking place' might express a hard-earned truth when uttered in a history classroom, yet express something false when uttered in a discussion of skepticism.[3]

Contextualism grew from a number of different roots. Firstly, it was common in discussions of skepticism to invoke a distinction between the high standards for knowledge that the skeptic seems to demand and the lower standards for knowledge that seem to be in place when we're making

---

[2] Other issues that have been prominent recently, but which I lack the space to discuss, include: renewed interest in the topic of *a priori* justification (see fn. 14 for references); Sosa's work on virtue epistemology (see [1991]) and the development of this field (for surveys, see Axtell [1997]; Greco [1993], [1999]); how externalism about content affects one's capacities for self-knowledge (this is discussed in some recent collections: Ludlow and Martin [1998]; Wright *et al.* [1998]; Boghossian and Peacocke [2001]); and a greater attention to testimony and to social aspects of knowledge more generally (for an overview, see Fuller [1996] and Schmitt [1999]).

[3] For recent defenses of Contextualism, see: Unger ([1986]); Cohen ([1987], [1988], [1999]); Heller ([1989], [1999a]); DeRose ([1992], [1995]); and Lewis ([1996]). For criticism, see Schiffer ([1996]); Feldman ([1999a]); and Vogel ([1999]). DeRose ([1999]) gives a useful overview of the field. Contextualism in the sense defined here is importantly different from the view Annis ([1978]) and Williams ([1996], [1999b]) defend, which makes a subject's epistemic status depend on facts about *the subject's* circumstances: e.g. what epistemic standards his community endorses, what kinds of activity he's engaged in, and so on. Contextualism makes a subject's epistemic status depend on facts about *the context in which that epistemic status is being appraised*. This might be different from the context the subject occupies. See the discussion of RA-Theory without Contextualism, below.

ordinary epistemic evaluations. The Contextualist can allow that both of these standards are legitimate, in their proper context. (And the Contextualist is not limited to just *two* kinds of standards: there may be *many* different standards, appropriate to different contexts.)

Another early root of Contextualism was the discussion of Gettier-cases where evidence one does not possess—e.g. a newspaper headline one hasn't read yet—seems to defeat one's claims to knowledge (see Harman [1968], [1973], Ch. 9, [1980]).[4] It is very natural to think that the extent to which this kind of evidence affects one's epistemic standing will vary with context. Suppose your friend in the White House tells you that the President is in Mexico, and suppose you truly believe on that basis that the President is in Mexico. But suppose also that, *unbeknownst to you*, the *Washington Post* has just printed a (false) story that claims that the President is not in Mexico but rather vacationing in Maine. In some contexts we might want to count your justified true belief as knowledge despite the existence of the misleading newspaper report. In other contexts, though, we'd take the fact that the newspaper report is widely disseminated to preclude your belief about the President's whereabouts from counting as knowledge. The Contextualist can make sense of these shifts.

A third root of Contextualism is the Relevant Alternatives Theory of knowledge put forward in the 1970s by Dretske ([1970], [1981]) and Goldman ([1976]). We should discuss this at length.

As we'll understand the RA-Theory, it says that, if q is an irrelevant alternative to p, then knowing p doesn't require you to have evidence which would enable you to rule q out; in fact, so long as q is irrelevant, you don't need to have *any* independent evidence against q. For example, if you see a zebra in a zoo pen, and the possibility that the animal you see is a cleverly disguised mule is an *irrelevant* alternative, then you would not need to have evidence against that alternative, in order to know that the animal is a zebra.

This RA-Theory combines naturally with Contextualism. The combined view says that what alternatives count as relevant is a function of the context in which the knowledge ascription is made or assessed (see Stine [1976]; DeRose [1992], [1995]; and Lewis [1996]).

Now, early proponents of the RA-Theory, like Dretske, denied that knowledge is closed under known logical consequence. But contemporary Contextualist proponents of RA-Theory argue that we should retain some kind of Closure Principle. Let q be some alternative you cannot rule out. Their idea is that, in a given context where q is irrelevant, then if you know p and you know that p implies not-q, you count as knowing not-q, too. Some

---

[4] Cohen ([1987]) discusses a similar phenomenon: the extent to which evidence one *does* possess but doesn't appreciate the force of can defeat one's claims to knowledge.

contextualists say that when we explicitly *raise the question* whether you know not-*q*, that makes *q* relevant, and hence, since you can't rule *q* out, you don't count as knowing not-*q*, and you don't count as knowing *p* either. But as they see it, there is no *single* context in which you both know *p* and fail to know not-*q*. In contexts where *q* is relevant, you count as knowing neither. In contexts where *q* is irrelevant, you count as knowing both.[5]

Although RA-Theory is nowadays often combined with Contextualism, it is important to recognize that these two theories don't *have to be* combined. Some RA-Theorists believe that what alternatives are relevant is wholly determined by facts about the putative *knower* and his environment (perhaps including facts about his epistemic community). It does not vary from one context of knowledge-*ascription* to another. That is a form of RA-Theory without Contextualism.[6]

One can also hold Contextualism without RA-Theory. This is the best way to understand Cohen ([1998], [1999]). Let's call this view *Strength-of-Evidence Contextualism*. (Vogel [1999] calls it 'Plain Contextualism'.) The Strength-of-Evidence Contextualist thinks that what you count as knowing varies from context to context, but he denies that you can *ever, in any context*, know a proposition while having *no* evidence against alternatives to that proposition. Facts about the context in which we're ascribing knowledge to a subject never excuse the subject from the requirement to have *some* evidence against alternatives to what he believes. For instance, the Strength-of-Evidence Contextualist says that until you have at least *some* evidence against the possibility that the animal you see is a cleverly disguised mule, you can't know that animal to be a zebra—even if the possibility that the animal is a cleverly disguised mule is an irrelevant alternative. What context does affect is *how strong your evidence has to be*, what newspapers you have to take into account, what defeating evidence you need to appreciate the force of, and so on. In some contexts, you can know that the animal you see is a zebra, while having only *very slender* evidence against the possibility that it's a cleverly disguised mule; in other contexts, you need a lot more evidence. By the same lights, the Strength-of-Evidence Contextualist says that in order to count as knowing anything about the external world on the basis of perception, you'd need to have at least *some* reasons for rejecting hypotheses about brains in vats and the

---

[5] Stine ([1976]); Cohen ([1988]); DeRose ([1995]); and Lewis ([1996]) all argue for Contextualist views which retain some form of Closure. Heller ([1999b]) defends Dretske's rejection of Closure. For further discussion of the relations between Contextualism and Closure, see Vogel ([1990]).

[6] This appears to be Dretske's view ([1991]). It is also formulated by Goldman ([1976]), though Goldman doesn't endorse the view. He prefers a Contextualist version of RA-Theory. On the difference between RA-Theory with Contextualism and RA-Theory without Contextualism, see Cohen ([1991]) and DeRose ([1992], [1999]). DeRose calls RA-Theory without Contextualism 'subject contextualism', contrasting it to the 'attributor contextualism' he advocates.

like, even in contexts where those hypotheses are irrelevant. Strength-of-Evidence Contextualism does not itself offer any story about where those reasons come from. We need some further, independent account of that.[7]

1.2 Ruling an alternative out

In discussions of RA-Theory and Contextualism, most attention has been paid to the question of what makes an alternative *relevant*. (See Vogel [1999], section 3 for discussion.) Comparatively little attention has been paid to how we manage to *eliminate* or *rule out* an alternative. So let's say a few words about this issue.

On Dretske's account, ruling q out is just a matter of knowing not-q. However, this account of ruling out would not suit those RA-Theorists who wish to retain Closure. For suppose that q is some alternative you can't rule out. According to the RA-Theorist, you can sometimes know p despite being unable to rule q out—so long as q is irrelevant. If q were to become relevant, then your inability to rule it out would prevent you from knowing p. Now, if we retain some form of Closure, then in those contexts in which you count as knowing p you'd also, by Closure, count as able to know not-q. This despite the fact that you're not able to rule q out. So if we retain Closure, we should understand ruling an alternative out as requiring *more* than just knowing that alternative to be false.

A better account might be an *evidential account* of ruling out. We could say that ruling q out requires you to have independent evidence against q, evidence which does not beg any of the questions for which q is an alternative.[8] An RA-Theorist who adopts this account of 'ruling out' can say that, when the possibility that you are a brain in a vat is irrelevant, you can know by perception that you have hands, and hence by Closure also know that you're not a handless brain in a vat. But you're not in a position to *rule out* the possibility that you are a brain in a vat, because you have no *independent* evidence against that possibility. Whatever justification you have for believing that you're not a brain in a vat crucially rests on things you

---

[7] As I said, Cohen ([1988], [1999]) is a Strength-of-Evidence Contextualist. He does not try to give any story about what justifies our rejection of run-of-the-mill alternatives, like the cleverly disguised mule. He does, however, think we have some inductive evidence that those alternatives do not obtain. More problematic are global skeptical hypotheses, like the hypothesis that one is a brain in a vat. Cohen thinks no possible *evidence* could tell against these alternatives. But he thinks they are in some sense *a priori* irrational, that we're entitled to reject them without evidence (compare BonJour [1985], sections 8.3–8.4).

Stine, on the other hand, rejects Strength-of-Evidence Contextualism in favor of a Contextualist version of RA-Theory. She says you can know that the animal in the pen is not a cleverly disguised mule merely by virtue of the fact that the disguised-mule hypothesis is irrelevant. You do not need any specific evidence that tells against it (Stine [1976], pp. 257–58). Some Contextualists like Lewis ([1996]) follow Stine here. Others like Heller ([1999b]) follow Cohen.

[8] See the discussion of antecedent justification in Pryor ([2000]); and the discussion of 'transmission failure' in Wright ([1985], [2000]) and Davies ([1998], [2001]).

purport to know about the external world by perception. It's for this reason that when the possibility that you are a brain in a vat becomes relevant, you no longer count as knowing you have hands.

A different account construes ruling out in *counterfactual* terms. According to this account, you rule an alternative out by having your belief be counterfactually sensitive to whether that alternative obtains. For example, when it's your knowledge of p that is in question, such an account could say that being able to rule q out is a matter of being such that, if q were true, you would no longer believe p.[9]

The accounts of ruling out mentioned so far attempt to be explanatory in a certain way. They appeal to some independent facts about your relation to your evidence in order to explain facts about what that evidence enables you to know. An alternative way to proceed would be to take as given the facts about what your evidence would enable you to know in certain contexts, and use those facts to explain the notions of 'being relevant' and 'ruling out'. This is Cohen's initial approach; he defines an alternative to p as 'relevant' just in case it has to be ruled out in order for you to know p ([1988], pp. 100–1). However, Cohen then goes on to discuss more substantial criteria for being 'relevant'. A similar view would say that you're in a position to rule some alternative to p out just in case you'd be able to know p even if that alternative were relevant. These construals of 'relevant' and 'ruling out' are perfectly legitimate, as far as they go. They just won't be very useful for giving explanations of what we know when. They take the facts about what we know when for granted, and ride piggy-back on them.

Though this issue of how to understand the notion of 'ruling out' needs more attention, Contextualism is currently a very active research program, and it has rapidly become one of the most seriously regarded responses to skepticism.

2 Modest forms of foundationalism

Another active source of resistance to skepticism comes from new foundationalist accounts of perceptual justification. Foundationalism's prospects quite generally look much better today than they did twenty-five years ago. In this section, we'll take a brief look at this development.

Say that you are 'immediately justified' in believing p iff you're justified in believing p, and this justification doesn't rest on justification you have for believing other supporting propositions. According to the foundationalist,

---

[9] This proposal is akin to Nozick's account of knowledge; however, the current proposal only requires your beliefs to be counterfactually sensitive to the obtaining of *relevant* alternatives. Goldman ([1976]); DeRose ([1995]); and Heller ([1989], [1999a], [1999b]) favor this account of ruling out. A related view says that you are in a position to rule q out iff, if q were true, you would not *have the same evidence* you in fact have for believing p. Lewis ([1996]) employs an account similar to this.

the justification for all our beliefs ultimately traces back to a set of 'basic beliefs', which we have immediate justification for believing.

Traditionally, foundationalists thought that these basic beliefs could only concern the nature of one's current thoughts and experiences. (And perhaps some *a priori* matters as well.) All our other empirical beliefs, including beliefs about our perceptual environment, had to rest inferentially on this austere foundation of beliefs about our current mental states. These basic beliefs were thought to be exceptionally secure. They were often claimed to be infallible and indubitable. According to many foundationalists, the reason why these beliefs were so secure was that we had a non-propositional 'direct apprehension' of the mental states the beliefs were about.

There are many reasons to be unhappy with that traditional cluster of views. It's generally felt that a foundation consisting only of infallible beliefs about our own mental states would be *too* austere to support the rest of our empirical beliefs. And many philosophers also criticized the traditional foundationalists' views about 'direct apprehension'.[10]

In addition, BonJour put forward an influential argument which seemed to show that, at least on internalist conceptions of justification, it was impossible for there to be any basic beliefs. BonJour defended an account of justification according to which, to be justified in believing a proposition *p*, you have to be justified in believing that your belief has some features which make it likely to be true. This account makes it impossible for your belief in *p* ever to be *immediately justified*. Your justification for your belief in *p* would always rest in part on your justification for those further beliefs about the belief's features.[11]

---

[10] See the anti-Given arguments in Sellars ([1956]); Davidson ([1981]); BonJour ([1978], [1985]); and Williams ([1999a]). BonJour has changed his mind about these arguments; see below.

These anti-Given arguments deserve a re-examination, in light of recent developments in the philosophy of mind. The anti-Given arguments pose a dilemma: either (i) direct apprehension is not a state with propositional content, in which case it's argued to be incapable of providing us with justification for believing any specific proposition; or (ii) direct apprehension *is* a state with propositional content. This second option is often thought to entail that direct apprehension is a kind of *believing*, and hence itself would *need* justification. But it ought nowadays to be very doubtful that the second option does entail such things. These days many philosophers of mind construe perceptual experience as a state with propositional content, even though experience is distinct from, and cannot be reduced to, any kind of belief. Your experiences represent the world to you as being a certain way, and the way they represent the world as being is their propositional content. Now, surely, *it's looking to you as if the world is a certain way* is not a kind of state for which you need any justification. Hence, this construal of perceptual experience seems to block the step from 'has propositional content' to 'needs justification'. Of course, what are 'apprehended' by perceptual experiences are facts about your perceptual environment, rather than facts about your current mental states. But it should at least be clear that the second horn of the anti-Given argument needs more argument than we've seen so far. To be fair, most of the foundationalists who appealed to a notion of 'direct apprehension' understood this notion along the lines of horn (i); and that is where their critics have focused most of their attention.

[11] See BonJour ([1978], [1985], Ch. 2). BonJour's account of justification is a version of what I call 'Access Internalism' in Section **3.1**, below. BonJour attempts to motivate this account by appeal to considerations about epistemic responsibility. We will discuss the connection between justification and notions like epistemic responsibility in Section **4**, below.

In recent years, though, foundationalism has staged quite a successful comeback.[12] There are two main reasons for this. In the first place, foundationalists nowadays think that the implausible elements in traditional foundationalism are inessential extras that aren't required by the reasoning that primarily motivates foundationalism. Newer forms of foundationalism are quite 'modest', in that they allow basic beliefs to be fallible, revisable, less than maximally justified, and so on. They only require that these basic beliefs be immediately justified—that is, justified in a way that doesn't rest on justification one has for other, supporting propositions. Some of these new foundationalist views allow beliefs about our perceptual environment to qualify as basic. Some of them even grant important epistemic roles to facts about coherence: e.g. they may allow facts about coherence to defeat or to strengthen one's justification for a belief, even a basic belief. These theories count as foundationalist because they say that, for every justified belief a subject has, *at least some* of his justification for that belief is immediate or ultimately rests on other justification that is immediate.

In addition, many philosophers have come to think that BonJour-style arguments against the possibility of basic beliefs are unsuccessful. The now-standard reply to BonJour's argument is that it builds into justified beliefs *about the world* conditions which really ought only apply to our higher-level, reflective beliefs that *our beliefs about the world are justified*. For you simply *to have* a justified belief about the world, this reply says, that belief may need to *have* features that make it reasonable or likely to be true, but you need not, in addition, *be justified in believing* it has such features. You can have a justified belief 'There is a hand', without also being justified in believing 'I am justified in believing that there is a hand' or 'I believe that there is a hand on grounds which make that belief likely to be true'. Acquiring justification for the latter beliefs would be a further epistemological accomplishment. Once this move is made, it is hard for BonJour's argument to get going.[13]

A notable recent development is that BonJour, who was formerly a leading coherentist about empirical beliefs, has given up coherentism and

---

[12] For some recent defenses of foundationalism, see Alston ([1989], Essays 1-3); Pollock and Cruz ([1999], Ch. 5, Section 4.2 and Ch. 7, Sections 1–3); and Audi ([1993], esp. Chs 3, 4, 10, and 12). I advocate a form of modest foundationalism about perceptual beliefs in Pryor ([2000]). McGrew ([1995]) argues against modest foundationalism, in favor of a more traditional view with infallible foundations. Recent years have also seen a number of attempts to revive the notion of direct apprehension. See Fumerton ([1985], [1995]); Moser ([1985], [1989]); Fales ([1996]); and BonJour ([1999b]). Finally, one might count Foley ([1987]) and Haack ([1993]) as foundationalists, though their views differ in important ways from the ones just cited. A good survey of foundationalist work in the late 1970s and 1980s is Triplett ([1990]).

[13] This is emphasized in various articles in Alston ([1989]), esp. Essay 6; and in Audi ([1993]), esp. Chs 4 and 11.

now himself advocates a form of foundationalism (see BonJour [1999a], [1999b]).[14]

3 The internalism/externalism debate

3.1 Simple Internalism and Access Internalism

Let's turn from debates about the structure of justification to debates about whether justification is internalist or externalist. First, let's ask what we mean when we call an account of justification 'internalist'.[15]

Internalism in epistemology is a different sort of issue than internalism in philosophy of mind. Philosophers of mind are concerned with whether certain mental properties (e.g. the property of having a belief with such-and-such content) supervene on *one's intrinsic make-up*. Epistemologists are more concerned with the connection between one's justificatory status and *facts to which one has a special kind of access*—e.g. facts about what one believes, what experiences and memories one has, facts about what one's beliefs are based on (see below), what one's goals are, what one is attempting to do, and so on. Clearly, not every state that is 'internal' in the philosophy of mind sense will be among these states. Sub-personal states of one's visual processing system are 'internal' in the philosophy of mind sense, but they are not states to which we have any special first-person access. In addition, it is a hotly debated question what sorts of special first-person access we have to the 'wide' contents of our thoughts (see the collections cited in fn. 2). So it is not obvious that every state which is 'internal' in the epistemologist's sense will be 'internal' in the philosophy of mind sense. Hence, it is certainly possible to combine externalism in philosophy of mind with internalism in epistemology, and vice versa.

What kind of 'special access' is at play in the debate about internalism in epistemology? Different accounts of this are possible. One could understand this 'special access' in terms of the *strength* of one's access: one could understand it as meaning infallibility, say, or as meaning that non-culpable mistakes are impossible. Alternatively, one could understand this 'special access' in terms of the *route* by which one has access: one could understand it as meaning that one

---

[14] BonJour has also been at the forefront of another recent movement, which is the development of modest accounts of *a priori* knowledge, and a revival of interest in rationalism (see BonJour [1998]; see also Edidin [1984]; Casullo [1988]; Burge [1993]; Plantinga [1993b], Ch. 6). These accounts of *a priori* knowledge often incorporate many elements of the new modest forms of foundationalism. George Bealer has written a series of papers on these issues, starting with Bealer ([1987]), and summarized in Bealer ([1999]). These will be incorporated in a forthcoming book. Two useful collections, including some of the works just cited, as well as other important papers, are Moser ([1987]) and Casullo ([1999]).

[15] For other discussions of this question, see Alston ([1989], Essay 8); Fumerton ([1988], [1995], Ch. 3); and Feldman and Conee ([forthcoming]).

can *know by reflection alone* whether one is in one of the relevant states. (By 'reflection' I mean *a priori* reasoning, introspective awareness of one's own mental states, and one's memory of knowledge acquired in those ways.) It would then be a further question whether one's reflective knowledge was infallible or especially certain in any remarkable way. Different choices about how to understand the notion of 'special access' will affect one's understanding of the entire internalism/ externalism debate. Most epistemologists understand the notion in the last way I described, and I will follow them.

Now, there are two dominant ways in which epistemologists understand the notion of 'internalism'. One is in terms of a supervenience thesis. The other is in terms of whether the subject has special access to her epistemic condition. We'll start with the supervenience construal of 'internalism'.

I will take the following supervenience thesis to express the most minimal internalist position:

> Simple Internalism: Whether one is justified in believing *p* supervenes on facts which one is in a position to know about by reflection alone.

I will take Externalism to be the denial of Simple Internalism.

Now, sometimes when we're evaluating your epistemic standing, we're only interested in whether you have justification for believing certain *propositions*—regardless of whether you actually do believe those propositions. Other times, though, we're interested in more than that. We want to know whether you *do believe* the propositions you have justification for believing, and if so, whether your belief is *based on* that justification. You may have very good reasons for believing *p*, but base your belief in *p* on bad reasons. In such cases, your *belief* is epistemically defective, even though it's a belief in a proposition you have justification for believing.[16]

The internalism/externalism debate bears on both of these kinds of epistemic evaluation. We can ask whether the facts that you're in a position to know about by reflection alone determine that you are *justified in believing some proposition*. And we can also ask whether those facts determine that your belief is *based on something that justifies it*. One normally sees externalist positions formulated in terms of the second kind of evaluation. For example, process reliabilism is usually formulated in terms of whether the process which led to

---

[16] In such cases, we say that your belief is 'ill-founded', though it's a belief in *a proposition* you have good reason to believe. See Feldman and Conee ([1985]) on the contrast between 'well-founded' and 'ill-founded' beliefs. A belief counts as well-founded iff the subject *has* grounds that support the proposition he believes, his belief is *based on* those grounds, and the subject isn't ignoring any evidence he has which defeats those grounds. See also Audi's discussion of the contrast between 'impersonal' and 'personal' justification, and the contrast between reasons *one has* for a belief and the reasons *for which one holds* a belief (in Audi [1993], Chs 3, 7, and 8). Korcz ([1997]) is a useful survey of literature on the basing relation.

Sometimes one will in *forming* a belief base it on one ground, but then later come to base the belief on another ground. This later ground will then be what sustains one's continued acceptance of the belief. For our purposes we can ignore such complications.

one's *actually forming* a given belief is reliable. This makes it look like the debate between process reliabilists and their opponents solely concerns the evaluation of beliefs. But Goldman pointed out that the internalism/externalism debate can be conducted in terms of either sort of evaluation (see his discussion of *ex post* vs. *ex ante* justification in Goldman [1979]). The focus in these discussions on the reliability of our *belief-forming* processes is sometimes just an indirect way of raising questions about the reliability of certain *sources* of justification, like our senses. Consider a community whose visual experiences are reliable guides to their environment, but who have a cognitive defect that leads to them never basing any beliefs on their visual experiences. One can imagine a reliabilist saying that these people's experiences give them *justification for believing* certain things, even though they have no well-founded *visual beliefs*. That would be a reliabilist view about the evaluation of propositions, rather than the evaluation of actual beliefs.

Whichever sort of evaluation we focus on, we should be careful to distinguish Simple Internalism from the following, more demanding internalist position:

> Access Internalism: One always has 'special access' to one's justificatory status.

A strong version of Access Internalism would say that whenever you're justified in believing *p*, then you're justified in believing that you are so justified. This is akin to the 'KK Principle', which says that whenever you know *p*, you also know that you know. (Nozick [1981], p. 281 construes 'internalism' as the acceptance of such a KK Principle.) Weaker versions of Access Internalism are also possible. For instance, one might claim that whenever you're justified in believing *p*, then the fact that you are so justified is one *you're in a position to acquire* justification for believing, by a certain route. This does not require that you actually already be justified in believing those facts. Like Simple Internalism, Access Internalism can be formulated either as a view concerning the evaluation of propositions or as a view concerning the evaluation of beliefs.[17]

'Access Internalism' is now a somewhat-established name for the view(s) just described. However in some ways this nomenclature is unfortunate, for *both* Simple Internalism and Access Internalism have *something* to do with considerations of access. The Simple Internalist says that the accessible facts form a supervenience base for whether we're justified; the Access Internalist says that whether we're justified is itself an accessible fact. It is important to recognize that Simple Internalism does not by itself entail that stronger claim. Even if the fact that one *is* justified is wholly determined by facts of which one is

---

[17] Chisholm and BonJour (in his coherentist period) were paradigm Access Internalists. (See e.g. Chisholm [1977], Ch. 6, Section 5; BonJour [1985], Chs 1–2.)

reflectively aware, that does not guarantee that one *can tell* just by reflection that one is justified. After all, one might be epistemologically untrained, or, worse, have false beliefs about what it takes to be justified. There's no guarantee that one would be able to correct these defects by mere armchair reflection.[18]

The way I formulated Access Internalism, it imposes an access requirement on *all* of one's justification. One might feel that this is too strong, especially for basic, unreflective kinds of justification like we get in perception. Access Internalism seems to say that our perceptual beliefs are justified only if we're justified in believing that they are justified. But whether and how our perceptual beliefs are justified is a matter that epistemologists are constantly investigating. It seems harsh to say that no one's perceptual beliefs are justified unless he has a justified position on that matter. In light of this, some philosophers prefer to impose access requirements only on the kind of justification we seek from *inferences*. Other philosophers require us to have access only to *the presence* of the stuff that justifies our beliefs (e.g. our experiences), and not to *the fact that* our beliefs are justified by that stuff. Let's look at these possibilities briefly.

3.2 Internalism about inferences and grounds

When you base a belief on some (genuine or alleged) body of supporting justification, call that justification your 'grounds' for your belief. Different things might count as grounds. In some cases, you will believe *p* on the basis of *other beliefs*; these other beliefs are *inferential grounds* for your belief in *p*. When you base your belief in *p* on an inferential ground, say the belief *q*, it is widely agreed that your belief in *p* is justified only when your belief in *q* is also justified. But there is disagreement about what more is required for your belief in *p* to be justified. (i) Do you need to have *good reason to believe* that your ground *q* constitutes good evidence for *p*? Or (ii) does *q* merely have to *in fact* be good evidence for *p*? Call someone who affirms (i) an Inferential Internalist. An Inferential Externalist denies (i) and affirms (ii) instead.[19]

---

[18] If we understand the relevant kind of 'special access' in more demanding ways, it will be even clearer that simple forms of internalism don't entail Access Internalism. For example, the fact that one's justificatory status is determined by facts of which one is *currently conscious*, or about which one is *infallible*, does not entail that one is currently conscious of, or infallible about, what one's justificatory status is.

[19] I borrow these terms from Fumerton ([1995], esp. Ch. 3). Fumerton defends something like the view I'm calling 'Inferential Internalism' (but see Cohen [1998] for some important distinctions which I'm glossing over here). See also the discussion of the 'connecting belief' requirement in Audi ([1993], Ch. 8). Feldman and Conee ([forthcoming]) are Simple Internalists who endorse Inferential Externalism.

Some Inferential Internalists require not merely that *q is justifiably believed to be*, but also that *q is in fact*, good evidence for *p*. Their commitment to (i) is what makes them Inferential Internalists.

Inferential Internalism is a kind of Access Internalism, limited just to one's inferential justification. (It says nothing about what kind of access we have to our immediate justification.) Denying Inferential Internalism does not force one also to deny Simple Internalism. Whether an Inferential Externalist accepts or rejects Simple Internalism will depend on what he thinks is required for q to be good evidence for p. The Simple Internalist will argue that whether q is good evidence for p is a necessary matter, or at least, that it's wholly determined by facts to which one has the relevant kind of 'special access'; the externalist will argue that it depends on contingent facts about one's environment, to which one has no 'special access'. These positions are both compatible with Inferential Externalism.

We've been discussing cases where your belief in p is based on other beliefs. In other cases, you will believe p but not on the basis of other beliefs; rather, your justification for believing p will come from the fact that you have certain experiences or memories, or from the fact that you understand certain propositions, or from some *a priori* intuition. These sorts of things are *immediately justifying grounds* for your belief in p. We have to proceed carefully here. To say that certain experiences are grounds for your belief does not mean that your belief is based on further *beliefs about* those experiences. Rather, on the views I'm envisaging, *merely having* those experiences gives you immediate justification for believing certain things— e.g. for believing that there is a table. When that's so, you needn't base your table-beliefs on any *beliefs about* your experiences.[20] However, we still need to distinguish between subjects who base their table-beliefs on the immediate justification their experiences afford them, and subjects who, despite having that experiential justification, none the less base their table-beliefs on other evidence (e.g. the stories their grandmother told them about tables). That is why we need the notion of basing a belief on an immediately justifying ground. Philosophers who deny that any beliefs are immediately justified will of course deny that a belief can be justified by virtue of being based on an immediately justifying ground.

Even when we count things like experiences as a kind of ground, it is still controversial whether justified beliefs are always based on some ground. For instance, it's unclear whether self-evident beliefs and some types of memorial belief are based on any grounds (see Plantinga [1993b], pp. 187ff.). But let's restrict our attention to beliefs that *are* based on grounds. We can ask several questions about the access we have to those grounds: Must a ground be

---

[20] I defend such a view in Pryor ([2000]). A more common view is that one's justification for beliefs about tables derives from one's justification for beliefs about one's experiences. On that view, one's experiences only serve as immediately justifying grounds for *beliefs about one's experiences*, and not for any beliefs about external objects. Acceptance of the view that there are immediately justifying grounds need not, though it may, be combined with the view that you have any kind of 'direct apprehension' of the ground, from which your justification for the belief derives.

something to whose *presence* one has the 'special access' which is operative in the internalism/externalism debate? Must one also have special access to whether the ground is an *adequate* one, that is, to whether it's possible to *justifiably* base a belief on that ground? Endorsing the weaker access requirement—namely, that grounds have to be things to whose presence one has special access—does not by itself commit one to Simple Internalism. For it could be that one tells an externalist story about what makes the ground adequate.[21] Alternatively, one could tell a Simple Internalist story about *why* a given ground is adequate, but refrain from claiming that all subjects will be in a position to know that the ground is adequate. That would be Simple Internalism without Access Internalism. If one endorsed the strongest access requirement—namely, that subjects must have special access both to the presence and to the adequacy of their grounds—then one would be a full-fledged Access Internalist (at least, concerning beliefs which *have* grounds).[22]

3.3 A spectrum of views

With these distinctions in hand, we can sketch out a spectrum of internalist and externalist views.

  (i) The strongest form of internalism advocates Access Internalism about all of one's beliefs. (BonJour in his coherentist period is an example of someone who held this view.)

 (ii) An intermediate form of internalism holds Inferential Internalism, but denies that one needs special access to one's *immediate* justification. (Fumerton holds this view.)

(iii) The weakest form of internalism eschews all forms of Access Internalism. On such a view, the fact that one's grounds are adequate must be an internally *determined* affair, but one need not have *special access to whether* those grounds are adequate. (Feldman and Conee hold this view.)

All of the preceding views count as internalist because they endorse Simple Internalism. The following views reject Simple Internalism:

 (iv) A weak form of externalism requires every justified belief to be based on some ground whose *presence* is accessible, but tells an externalist story about what it is for grounds to be adequate. (This is the view held by Alston and some other externalists; see fn. 21.)

---

[21] That is Alston's view (see [1989], Essays 4 and 9). See also Goldman ([1980]); Swain ([1981], Ch. 4, [1985]); and Conee ([1988]). Such views are examples of what Pollock and Cruz call 'norm externalism' ([1999], Ch. 5).

[22] A further question about access is: Do we have special access to facts about *whether* a belief is based on a given ground? For discussion of this question, see Audi ([1993], Ch. 11) and Kim ([1993]).

    (v)   The strongest form of externalism not only denies Simple Internalism, but also denies that justified beliefs are always based on accessible grounds. (Plantinga advocates a view of this sort about *warrant*; I will discuss the difference between justification and Plantinga's notion of 'warrant' in Section **4.2**, below.)

Of these views, (i), (iv), and (v) have been most prominent in the literature. (ii) and (iii) have had fewer champions. These latter two positions are *conceptually* important, however, for if one overlooked them, one might take criticisms of the other positions to have more force than they actually do. For instance, arguments against (i) cannot by themselves compel one to accept (iv) or (v); nor can arguments for Simple Internalism by themselves compel one to accept (i).

Now we understand the internalism/externalism debate a bit better, and appreciate the variety of forms it can take. Let's turn next to the question whether assumptions about the ethics of belief force us to take any particular position in the internalism/externalism debate.

4 The ethics of belief

4.1 What does it mean to say that justification is 'normative'?

It is often argued that the fact that justification is a certain kind of *normative* notion has implications for the internalism/externalism debate. I want to examine this claim. To begin, we need to know what's meant by saying that justification is 'normative', or that epistemology is a 'normative' discipline. Here we face the problem that there are different ways of being 'normative', important to different epistemological debates. We need to separate out the ones that are relevant here.[23]

The first debate about 'normativity' concerns the question of how relevant empirical research is to epistemology. Some naturalists argue that epistemology should be conceived as an extension of descriptive psychology; their opponents argue that epistemology is more than merely a descriptive enterprise, and that the proper methodology for doing epistemology is an a priori one. Various intermediate views are possible.[24]

---

[23] In constructing this chart of different debates about 'normativity', I've been much helped by: Alston ([1989], Essays 4, 5, and 8); Kornblith ([1983]); Feldman and Conee ([1985]); Feldman ([1988], [2000]); Goggans ([1991]); Plantinga ([1993b], pp. 45ff.); Montmarquet ([1993], Appendix 2); Foley ([1994]); Goldman ([1994]); Fumerton ([1995], Ch. 1); and discussions with J. Cruz, T. Kelly, and R. Wedgwood.

[24] Quine ([1969]) is standardly thought to have favored *replacing* epistemology with descriptive psychology. (See Foley [1994] for a careful account of what Quine's actual views were.) Other naturalists like Kornblith don't go so far; they believe there are genuine normative facts about which modes of inference are *good modes of reasoning*; but they think we can't tell what those good reasons are without studying how we *actually do* reason and acquire beliefs. For some useful orientation to this debate, see Kim ([1988]); Maffie ([1990b]); Goldman ([1994]); BonJour ([1994]); Kornblith ([1995], [1999]); and Feldman ([1999b]).

The second debate about 'normativity' in epistemology concerns the question whether epistemological recommendations are merely *hypothetical*, or whether they're somehow non-optional or *categorical*. Some philosophers argue that it's only insofar as you have an independent interest (perhaps a prudential interest) in having 'epistemically rational' beliefs, or getting at the truth in a reliable way, that you should care what beliefs are proscribed or recommended by epistemology.[25] Their opponents argue that epistemology's recommendations apply to all agents, independently of what independent interests those agents may have. (Of course, an agent can be *unaware* of some epistemological prescription that does in fact apply to him.)

A third debate about 'normativity' focuses around the *kinds* of recommendations that epistemology should be seen as issuing. To some extent, we can affect what we believe by *the actions* we take.[26] So the question arises: does epistemology merely issue guidelines for evaluating the beliefs we end up with, or does it also issue guidelines for how to go about forming beliefs? For instance, does it instruct us not to perform actions that would result in our having Pascal's-wager-type beliefs, beliefs which are beneficial but evidentially unsupported?[27] And does epistemology only concern itself with what we ought to believe, given a certain body of evidence? Or does it also issue guidelines on *how to behave as an inquirer*? For instance, does it ever tell us how and when to go out and gather more evidence? Does it tell us to consult multiple sources, use controls in our experiments, refrain from destroying or manufacturing evidence, listen carefully to criticism, avoid mind-altering drugs, and so on?[28] And if epistemology does issue guidelines of this sort, the question arises: in what ways does a subject's conforming or failing to conform to those guidelines affect the justificatory status of his beliefs? (I place this last question under the heading of the fifth debate, discussed below.)

A fourth debate about 'normativity' concerns the connection between being justified and *following* the norms or guidelines that epistemology issues. Is the notion of justification supposed to play a role in *guiding* and *regulating* our beliefs? That is, ought we to say that a belief is justified when and only when the

---

[25] Some relevant literature here is Laudan ([1987]); Foley ([1987]); Giere ([1988]); Stich ([1990]); Maffie ([1990a]); Kornblith ([1993]); and Nozick ([1993]).

[26] There is an extensive literature on what sorts of *choice* or *voluntary control* we have in forming our beliefs. See e.g. Williams ([1970]); Winters ([1979]); Heil ([1983b], [1984]); Naylor ([1985]); Cook ([1987]); Alston ([1989], Essay 5); Feldman ([1988], [2001]); Bennett ([1990]); Montmarquet ([1993], Ch. 5); Plantinga ([1993a], Ch. 2); Scott-Kakures ([1994]); Radcliffe ([1997]); Steup ([1988], [1999]); and various essays in Steup ([2001]).

[27] This issue is connected to the question discussed above, whether epistemological recommendations have more than hypothetical force. See also Meiland ([1980]); Heil ([1983a], [1992]); Kelly ([2000]); and Feldman ([2000]).

[28] See Goldman ([1978], [1991]); Kornblith ([1983]); Heil ([1983b]); Feldman and Conee ([1985]); Foley ([1987]); Laudan ([1996]); Hall and Johnson ([1998]); Feldman ([1988], [2000]); and DeRose ([2000]) for discussion of these kinds of epistemic guidelines. Some of the literature on virtue epistemology is also relevant here; see the surveys cited in fn. 2.

belief is permitted or recommended by some proper belief-guiding recipe, which the agent adhered to in forming the belief? Or ought we to view the notion of justification as merely playing a role in third-person evaluations of belief, e.g. when we're deciding what propositions some agent counts as knowing? It is controversial whether the notion of justification should be thought of as a belief-guiding notion.[29] One's answer to that question will depend in part on how much choice or control one thinks we have over what we believe (see fn. 26). Some philosophers have argued that the question whether justification is a belief-guiding notion profoundly affects the internalism/externalism debate. We will consider the extent to which it does affect that debate below. But we need a bit more conceptual apparatus before we can do that. So let's press on.

The fifth and last debate about 'normativity' in epistemology concerns the connection between justification and the notions of epistemic praise and blame. Some philosophers believe that justification should be understood in terms of epistemic blamelessness and epistemic responsibility. They say that whether one is justified depends on how well one has met one's epistemic obligations, whether one is 'doing the best one can' as a believer, or at least the best that can reasonably be expected of one, and so on. In the literature, these are called *deontological conceptions* of justification, because of the primacy they usually attach to meeting one's epistemic obligations. We might also speak of *praise-and-blame* conceptions of justification, or *responsibilist* conceptions of justification, if we wanted to emphasize different elements. These views are all closely related to each other.[30]

Philosophers who oppose those views acknowledge that 'justification' is an evaluative notion, a term of praise or criticism; but they deny that there is a tight connection between one's justificatory status and how well one has met one's obligations, whether one merits praise and blame, and so on. Some of these opponents deny that *we have* any epistemic obligations. Others agree that we have epistemic obligations—and that we can assign a kind of epistemic praise and

---

[29] Goldman is an example of someone who construes justification as a merely evaluative notion, not a belief-guiding one. (Goldman does sometimes discuss regulative questions, though: see [1978], [1980]. Note that the 'rules' he talks about in [1986] are not meant to constitute a belief-guiding recipe. Goldman emphasizes this on pp. 25–26 and 59.)

[30] For discussions of the conceptual connections between justification, epistemic obligation, and epistemic praise and blame, see: Goldman ([1980], esp. Section V, [1988]); Plantinga ([1988], [1993a], Ch. 1); Steup ([1988]); Alston ([1989], Essays 4, 5, and 8); Moser ([1989], Section 1.3); Fumerton ([1995], Ch. 1); and Haack ([1997]). Some philosophers who think of justification in deontological terms are: Ginet ([1975], Ch. 3); Chisholm ([1977], Ch. 1, Section 5, [1982]); Kornblith ([1983]); BonJour ([1985], Ch. 1); and Steup ([1999]). See the start of Goldman ([1999]) on the connection between deontologism and the belief-guiding conception of justification described above.

In discussions of deontological conceptions of justification, it's common to distinguish between *objective* obligations and *subjective* obligations, and to tie blamelessness to conformity to one's subjective obligations. See Goldman ([1980], pp. 36ff., [1986], pp. 73ff.); Alston ([1989], Essay 4 pp. 86ff.); Plantinga ([1993a], pp. 15ff.); and Pollock and Cruz ([1999], pp. 140ff.) for this distinction.

blame by reference to them—but they deny that one's justificatory status is always a function of the epistemic praise or blame one merits. On such views, the question whether one's belief is genuinely *justified* is not settled by the fact that one is blameless in holding the belief. As they see it, there is a one kind of epistemic standing, which involves meeting all of one's epistemic obligations and proceeding in a responsible and blame-free manner; and there is another kind of epistemic standing, having a genuinely justified belief; and the first does not suffice for the second. (We will discuss such views further below.)

In what follows, we will focus on whether justification is 'normative' in the ways operative in the fourth and fifth debates, and how one's views about that affect the internalism/externalism debate.

4.2 Justification, warrant, and epistemic blamelessness

These days, unfortunately, the terminological situation with respect to 'justification' is rather messy. A bit of stipulation will help.

If you believe that Bing is a male canary, and that 9 out of 10 male canaries can sing, and you have no other relevant evidence, then regardless of what you're *inclined* to believe on the basis of that evidence, there's an intuitive sense in which your evidence *supports* the conclusion that Bing can sing. Similarly, if you have a visual experience as of something yellow, and no countervailing evidence, what your evidence supports is the belief that there is something yellow in front of you. Let's call a belief that is supported in this way a *justificatorily supported* belief. For now, let's leave it an open question whether this support relation is internalist or not; and let's also leave it open whether this support relation should be understood in deontological terms.

I find it most natural to use the term 'justification' to refer to the support relation just articulated, however that relation ultimately gets to be analyzed.[31] However, as we'll see, some philosophers use the term 'justification' in a more restrictive way.

Following Plantinga, we can reserve the term 'warrant' for a second notion, the notion of what has to be added to true belief to yield knowledge.[32]

---

[31] Alston ([1993]) and Cohen ([1995]), on the other hand, doubt whether *there is* any common pre-theoretic notion of justificatory support for internalists, externalists, deontologists, etc., to be giving competing analyses of.

[32] See Plantinga ([1993a], [1993b]). Other philosophers use the term 'warranted' to express different notions (see e.g. Meiland [1980]; Pollock [1983]; and Wright [1991]).

  Burge ([1993], [1996]) draws a contrast between 'entitlement' and 'justification'. In Burge's usage, justification is a kind of positive epistemic quality whose presence and nature is readily accessible to the subject upon reflection. Entitlement is a lower-grade positive epistemic quality whose presence and nature is not as readily accessible. On Burge's view, basic unreflective capacities like perception give us entitlements for believing things, but not justifications. Burge uses the term 'warrant' to describe the entire class of positive epistemic qualities, comprising both entitlements and justifications. In my terminology, what Burge calls 'warrants' are relations of justificatory support, and what he calls 'justifications' and 'entitlements' are two different species of justificatory support. Burge does not discuss the Plantingean notion of warrant at all.

It is most natural to think that warranted beliefs are beliefs for which one has justificatory support, which are appropriately based, and which have some extra features necessary to handle the Gettier problem. But there is much disagreement about this. Some philosophers deny that warrant can be analyzed into justificatory support plus some extra features.[33] And some philosophers believe that it's possible to know things for which one has no justificatory support (see e.g. Alston [1989], Essay 7); if they're right, it follows that one can sometimes have warranted beliefs for which one has no justificatory support.

As I said above, not every philosopher acknowledges the existence of epistemic obligations. But as I also pointed out, even if one *does* acknowledge the existence of such obligations, it's a further question what the relationship is between meeting one's epistemic obligations and having a justificatorily supported belief. So let's stipulate that talk of 'epistemic blamelessness' is to concern how well one meets one's epistemic obligations, behaves in an epistemically responsible manner, and so on. Our stipulation will leave it open whether being justified and being epistemically blameless always go hand-in-hand. It needs *argument* to show that there is any close connection between these properties. Deontologists think such an argument can be given; and hence that being epistemically blameless suffices for being justified. We'll look at some reasons for disagreeing in a moment. But however that debate turns out, we can all agree that *there is* a legitimate and useful notion of epistemic blamelessness.

As I said, some philosophers deny that we have any real epistemic obligations. Hence they will not think that the notions of epistemic praise and blame I'm describing have any application. They might hold this view because they think we lack the kind of voluntary control over our beliefs that the notions of praise and blame demand (see fn. 26, above). My point is just that one *can* think there is a legitimate and useful notion of epistemic blamelessness without thereby committing oneself to a deontological conception of justification.

For example, Montmarquet describes a notion of 'justification' that has close ties to epistemic praise and blame, but he wants this notion of justification to ground our ascriptions of moral praise and blame for actions, not to play any important role in the theory of knowledge (see [1993], esp. Ch. 6). In my terminology, Montmarquet is giving an account of epistemic blamelessness, and not taking a stand on what the relation is between justification and blamelessness.

Later in his career, Chisholm also advocated the view that we have epistemic obligations, but the notion of justification should not be explicated in terms of them.[34]

---

[33] This is a guiding idea in Williamson ([1995], [1997]) and Plantinga ([1993a], [1993b]).
[34] Plantinga calls this 'post-classical Chisholm' (see Plantinga [1993a], Ch. 3).

Plantinga reserves the term 'justification' for the notion I'm calling 'epistemic blamelessness'; and Alston is tempted to do the same.[35] They claim that the etymology of 'justification' supports understanding it in this deontological way. I think it's unclear what the etymology supports; but in any case, in natural languages usage tends not to be a slave to etymology. And given the current usage of the term 'justification' among epistemologists, it seems to me better to use the term 'justification' for our first notion, the relation of justificatory support whose nature is being debated. Many past and some present-day epistemologists *believe* that justification should be explicated in deontological terms. But we shouldn't build it into our definition of what we mean by 'justification' that this is so.

4.3 How should one's ethics of belief affect one's views about the internalism/externalism debate?

Now we are finally in a position to explore the connections between the ethics of belief and the internalism/externalism debate. These connections have been the subject of much recent discussion.

There are good reasons for thinking that justification and epistemic blamelessness can come apart. Doing all that can be reasonably be expected of you, epistemically, is compatible with holding beliefs that have little or no genuine support. Imagine a hapless subject who is taught bad epistemic standards. For instance, suppose he makes mistakes when engaging in tricky statistical reasoning. Let's say the standards he's been taught fail to distinguish between the likelihood that a test will yield a false negative and the likelihood that the test will yield a false positive. Our hapless subject does his best to apply these standards, but unfortunately he is not intellectually capable of discerning their defects. Furthermore, he has lost all memory of his childhood and so doesn't recall how it was he first acquired these standards. Cases of this sort naturally prompt two judgments. First, the subject's beliefs are epistemically defective in some important way. But second, the subject cannot be held culpable for those defects. You can say 'The hapless subject *ought* to believe the things that his evidence *really* supports', but this 'ought' doesn't seem to express any epistemic *obligation*. Our subject is doing the best he can, and the best that can reasonably be demanded of him. That

---

[35] See Plantinga ([1993a]) and Alston ([1989], Essay 4, fn. 21 and Essay 5, p. 143). Goldman ([1988]) calls the notion of epistemic blamelessness 'weak justification'.

seems a good reason to say he's violating no epistemic obligations.[36] Yet the beliefs he forms on the basis of statistical reasoning will be seriously defective; and for that reason it's natural to regard those beliefs as *unjustified*. Hence it appears that a subject can sometimes be blameless for holding unjustified beliefs.[37]

I myself think that justification and epistemic blamelessness do come apart in this way, and that this militates against any deontological conception of justification. But for the sake of argument, let's suppose that the line of reasoning I sketched can be resisted, and that the deontologist is right to forge a strong connection between justification and epistemic blamelessness. What then would follow? One often sees it claimed that deontologism and internalism go hand-in-hand. However, there are some dissenting voices about this.

Let's first consider whether viewing justification as a kind of belief-guiding notion, or viewing it in deontological terms, commits one to internalism.[38]

---

[36] Feldman (Feldman and Conee [1985]; Feldman [1988], [2000]) and Wolterstorff ([1997]) employ a very thin notion of 'obligation', which merely requires some epistemic 'ought' -claim to be true of a subject. This thin notion is insensitive to issues about what can reasonably be demanded of a subject, and whether the subject could have done otherwise than he did. I think this notion is too thin to capture our concept of a genuine obligation. (See Goggan's [1991] useful contrast between 'purely valuative' norms and 'deontic' norms. Only the second kind of norm corresponds to a genuine kind of epistemic obligation.) In any case, regardless of whether we regard this thin notion as genuine kind of obligation, it should be clear that there is no direct route from what one is 'obliged' to do in that sense to what praise and blame one merits. When we talk of epistemic praise and blame, I assume we're talking about a kind of genuine *culpability* on the part of the subject, and not merely complimenting or criticizing the subject's beliefs. Praise and blame in this sense are sensitive to what can reasonably be demanded of one, and whether one could have done otherwise. In ([2000]), Feldman acknowledges that the notion of 'obligation' he's employing may come apart from assessments of praise and blame.

[37] Alston ([1989], Essay 4); Goldman ([1988]); and Feldman and Conee ([forthcoming]) all defend the claim that being epistemically blameless does not suffice for being justified. (Goldman calls blameless beliefs 'weakly justified', and calls genuinely justified beliefs 'strongly justified'.) Alston and Goldman believe that the notion of being genuinely justified, which goes beyond mere blamelessness, should be analyzed in externalist terms. However, I take that to be a further question. For a nice critical summary of Goldman's views on this topic, see Riggs ([1997]).

Plantinga ([1988], [1993a], Chs 1–2) argues that epistemic blamelessness doesn't suffice for *warranted* belief; but many of the cases he discusses also show that being epistemically blameless doesn't suffice for having what I'm calling a genuinely *justified* belief. (Be careful: Plantinga uses the term 'justified' differently than me; he uses it as a synonym for 'epistemically blameless'.)

In the text, I presented an argument that one can have unjustified beliefs that are epistemically blameless. Another question is whether it's possible to have blameworthy beliefs that are none the less justified. That is a more delicate question.

[38] Plantinga ([1993a], Ch. 1) argues that it would commit one to internalism. Alston ([1989], Essay 8) discusses arguments from deontologism to different internalist theses. See also Greco ([1990]). Pollock and Cruz ([1999], Ch. 5) argue from a regulative or 'procedural' conception of justification to a kind of 'internalism'. However, their definition of 'internal state' (see pp. 132–33) is different from the more standard one being used here; as a result, we do not mean the same thing by 'internalism'. Also, I'm not convinced that their understanding of what it is for behavior to be 'regulated by a norm' corresponds to the subject's genuinely *following*, or *being guided by* a norm. They seem to be more concerned with the underlying mechanisms that *causally control* a subject's behavior.

Let's assume that the relevant belief-guiding recipes or norms or obligations are of the form: 'In circumstances C, believe *p*'. The question then arises whether C has to be the kind of circumstance to which the subject has the 'special access' that is operative in the internalism/externalism debate. For instance, must the subject always be able to tell whether C obtains by reflection alone? Deontologists often argue in the following way:

> If a belief-guiding recipe is to be *usable* in deciding what to believe, then the circumstances C it refers to must be circumstances such that the subject *can tell* whether they obtain, when he's following the recipe. In particular, they must be such that the subject can tell by reflection alone whether they obtain.

However, Goldman argues that this last claim is gratuitous. Perhaps the subject needs to be able *to tell* whether C obtains, if he's to follow the recipe, but Goldman denies that he has to be able to tell whether C obtains *by reflection alone* (see Goldman [1980], esp. Section IX, [1999]; Jacobson [1997]). So the extent to which we have 'special access' to C is a debated question.

Another question that arises is whether the fact that a given belief-guiding recipe *is correct* will be an internally determined matter, or a matter that one can ascertain by reflection alone. These issues are also controversial. For instance, Goldman at one point defended a reliabilist story about what makes belief-guiding recipes correct (again, see Goldman [1980] and Jacobson [1997]).

Finally, Kornblith ([1982]) argues that thinking that justification is to be explained in terms of epistemological responsibility does not commit us to any form of Access Internalism.

So there does not seem to me to be any straightforward, uncontroversial route from a deontological or belief-guiding conception of justification to any of the internalist views we described earlier.

Let's instead consider whether there are any interesting entailments in the opposite direction. Plantinga argues that internalism is well motivated only when one is working with a deontological conception of justification ([1993a], esp. Chs 1 and 3). Is that right?

It is true that some prominent rationales for internalism do employ deontological conceptions of justification. For instance, arguments from thought-experiments involving brains in vats often include remarks like this:

> There's no way for the brain to discover his plight; so he can't be *blamed* for forming the beliefs he does. He forms the same beliefs anyone else would form on the basis of that evidence. So even though many of his beliefs are false, he is at least *justified* in having those beliefs.

Here it's being assumed that being epistemic blameless is sufficient for being justified.

But while arguments for internalism often *do* employ deontological conceptions of justification, I am not persuaded that they *have to*. For instance, Cohen ([1984]) argues for internalism by appeal to thought experiments involving brains in vats, and his arguments never appeal in any essential way to deontological assumptions. Many philosophers share Cohen's intuition that it's possible for a brain in a vat, if he conducts his affairs properly, to have many justified (albeit false) beliefs about his environment. And—at least in my case—this intuition survives the recognition that being *epistemically blameless* does not suffice for being justified. It doesn't merely seem to be the case that the brain in a vat can form beliefs in a way that is epistemically blameless. It also seems to be the case that he can form beliefs in a way that is epistemically *proper*, and that the beliefs he so forms would be fully justified—despite the fact that they're reliably false.

To illustrate this, consider three brains in a vat. The first brain is epistemically quite reckless. He believes whatever he feels like believing, so long as he doesn't have incontrovertible evidence that it's false. The second brain is like the hapless subject we considered above. He's doing the best he can, epistemically, and the best that can be reasonably expected of him. He's very scrupulous and careful in following the epistemic standards he picked up as a child. Unfortunately, these epistemic standards are defective in a way he's not capable of recognizing. (Perhaps, as before, they fail to distinguish false negatives from false positives.) The third brain carefully follows the same epistemic standards we follow.

It seems natural to say that, although there's a sense in which the second and third brains are both epistemically blameless, only the third is conducting his epistemic affairs *properly*, and so only he is justified in his beliefs. This indicates that there's a kind of positive epistemic status which goes beyond mere epistemic blamelessness, but which one can possess even if one's beliefs are reliably false, as the brains' beliefs are. Such considerations incline me to an internalist, non-deontological account of justification.[39]

Now, not every philosopher shares these intuitions about brains in vats. So by themselves they are not likely to settle the debate about internalism. Other arguments are needed to do that. But these intuitions about brains in vats do cast doubt on the claim that it would *only* be possible to motivate or defend internalism by appealing to deontological assumptions. I have so far seen no

---

[39] This style of argument from multiple brains in vats derives from Cohen ([1984], p. 283). There are a number of other non-deontological internalists. See e.g. Fumerton ([1988], [1995], Ch. 1); Audi ([1993], Ch. 10); and Feldman and Conee ([forthcoming]). See also the works of 'post-classical Chisholm', referred to in fn. 33, above.

convincing argument for that claim. So there do not seem to me to be any clear entailments from internalism to a deontological conception of justification.

In my own view, then, it remains to be seen whether there is any tight connection between the ethics of belief and the internalism/externalism debate. I certainly have not shown that there *could be* no such connection. But I hope that the present discussion and the papers I've pointed to make it clear that this question is really less settled than it sometimes appears.

Acknowledgements

Department of Philosophy
Harvard University
Cambridge, MA 02138, USA
jpryor@fas.harvard.edu

References

Alston, W. [1989]: *Epistemic Justification*, Ithaca: Cornell University Press.

Alston, W. [1993]: 'Epistemic Desiderata', *Philosophy and Phenomenological Research*, **53**, pp. 527–51.

Annis, D. [1978]: 'A Contextualist Theory of Epistemic Justification', *American Philosophical Quarterly*, **15**, pp. 213–19.

Audi, R. [1993]: *The Structure of Justification*, Cambridge: Cambridge University Press.

Axtell, G. [1997]: 'Recent Work on Virtue Epistemology', *American Philosophical Quarterly*, **34**, pp. 1–26.

Bealer, G. [1987]: 'Philosophical Limits of Scientific Essentialism', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **1**, pp. 289–365.

Bealer, G. [1999]: 'The A Priori', in J. Greco and E. Sosa (*eds*), *Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 243–70.

Bennett, J. [1990]: 'Why Is Belief Involuntary?' *Analysis*, **50**, pp. 87–107.

Boghossian, P. and Peacocke, C. [2001]: *New Essays on the A Priori*, Oxford: Oxford University Press.

Bonjour, L. [1978]: 'Can Empirical Knowledge Have a Foundation?', *American Philosophical Quarterly*, **15**, pp. 1–13.

Bonjour, L. [1985]: *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.

Bonjour, L. [1994]: 'Against Naturalized Epistemology', in P. A. French, T. E. Uehling and H. Wettstein (*eds*), *Midwest Studies in Philosophy*, **19**, pp. 283–300.

Bonjour, L. [1998]: *In Defense of Pure Reason*, Cambridge: Cambridge University Press.

Bonjour, L. [1999a]: 'The Dialectic of Foundationalism and Coherentism', in J. Greco and E. Sosa (*eds*), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 117–42.

Bonjour, L. [1999b]: 'Foundationalism and the External World', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **13**, pp. 229–50.

Burge, T. [1993]: 'Content Preservation', *Philosophical Review*, **102**, pp. 457–88.

Burge, T. [1996]: 'Our Entitlement to Self-Knowledge', *Proceedings of the Aristotelian Society*, **96**, pp. 91–116.

Casullo, A. [1988]: 'Revisability, Reliabilism and A Priori Knowledge', *Philosophy and Phenomenological Research*, **49**, pp. 187–213.

Casullo, A. [1999]: *A Priori Knowledge*, Aldershot: Ashgate/Dartmouth.

Chisholm, R. M. [1977]: *Theory of Knowledge*, 2nd edn, Englewood Cliffs, NJ: Prentice-Hall.

Chisholm, R. M. [1982]: 'A Version of Foundationalism', in *Foundations of Knowing*, Minneapolis: University of Minnesota Press, pp. 3–32.

Cohen, S. [1984]: 'Justification and Truth', *Philosophical Studies*, **46**, pp. 279–95.

Cohen, S. [1987]: 'Knowledge, Context, and Social Standards', *Synthese*, **73**, pp. 3–26.

Cohen, S. [1988]: 'How to Be a Fallibilist', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **2**, pp. 91–123.

Cohen, S. [1991]: 'Skepticism, Relevance, and Relativity', in B. McLaughlin (*ed.*), *Dretske and His Critics*, Cambridge, MA: Blackwell, pp. 17–37.

Cohen, S. [1995]: 'Is There an Issue About Justified Belief?', *Philosophical Topics*, **23**, pp. 113–27.

Cohen, S. [1998]: 'Fumerton on Metaepistemology and Skepticism', *Philosophy and Phenomenological Research*, **58**, pp. 913–8.

Cohen, S. [1999]: 'Contextualism, Skepticism, and the Structure of Reasons', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **13**, pp. 57–90.

Conee, E. [1988]: 'The Basic Nature of Epistemic Justification', *Monist*, **71**, pp. 389–404.

Cook, J. T. [1987]: 'Deciding to Believe without Self-Deception', *Journal of Philosophy*, **84**, pp. 441–46.

Davidson, D. [1981]: 'A Coherence Theory of Truth and Knowledge', in E. LePore (*ed.*) [1986], *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Oxford: Blackwell, pp. 307–19.

Davies, M. [1998]: 'Externalism, Architecturalism, and Epistemic Warrant', in C. Wright, M. Smith and C. Macdonald (*eds*), *Knowing Our Own Minds: Essays in Self-Knowledge*, Oxford: Oxford University Press, pp. 321–61.

Davies, M. [2001]: 'Externalism and Armchair Knowledge', in P. Boghossian and C. Peacocke (*eds*), *New Essays on the A Priori*, Oxford: Oxford University Press.

DeRose, K. [1992]: 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, **52**, pp. 913–29.

DeRose, K. [1995]: 'Solving the Skeptical Problem', *Philosophical Review*, **104**, pp. 1–52.

DeRose, K. [1999]: 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 187–205.

DeRose, K. [2000]: 'Ought We to Follow Our Evidence?' *Philosophy and Phenomenological Research*, **60**, pp. 697–706.

Dretske, F. [1970]: 'Epistemic Operators', *Journal of Philosophy*, **67**, pp. 1007–23.

Dretske, F. [1981]: 'The Pragmatic Dimensions of Knowledge', *Philosophical Studies*, **40**, pp. 363–78.

Dretske, F. [1991]: 'Dretske's Replies', in B. McLaughlin (ed.), *Dretske and His Critics*, Cambridge, MA: Blackwell.

Edidin, A. [1984]: 'A Priori Knowledge for Fallibilists', *Philosophical Studies*, **46**, pp. 189–97.

Fales, E. [1996]: *A Defense of the Given*, Lanham, MD: Rowman & Littlefield.

Feldman, R. [1988]: 'Epistemic Obligations', in J. Tomberlin (ed.), *Philosophical Perspectives*, **2**, pp. 235–56.

Feldman, R. [1999a]: 'Contextualism and Skepticism', in J. Tomberlin (ed.), *Philosophical Perspectives*, **13**, pp. 91–114.

Feldman, R. [1999b]: 'Methodological Naturalism in Epistemology', in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 170–86.

Feldman, R. [2000]: 'The Ethics of Belief', *Philosophy and Phenomenological Research*, **60**, pp. 667–96.

Feldman, R. [2001]: 'Voluntary Belief and Epistemic Evaluation', in M. Steup (ed.), *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Oxford: Oxford University Press.

Feldman, R. and Conee, E. [1985]: 'Evidentialism', *Philosophical Studies*, **48**, pp. 15–34.

Feldman, R. and Conee E. [forthcoming]: 'Internalism Defended', in H. Kornblith (ed.), *Epistemology: Internalism and Externalism*, URL: < http://www.ling.rochester.edu/∼feldman/papers.html >.

Foley, R. [1987]: *The Theory of Epistemic Rationality*, Cambridge, MA: Harvard University Press.

Foley, R. [1994]: 'Quine and Naturalized Epistemology', in P. A. French, T. E. Uehling and H. Wettstein (eds), *Midwest Studies in Philosophy*, **19**, pp. 243–60.

Forbes, G. [1984]: 'Nozick on Skepticism', *Philosophical Quarterly*, **34**, pp. 43–52.

Fuller, S. [1996]: 'Recent Work on Social Epistemology', *American Philosophical Quarterly*, **33**, pp. 149–66.

Fumerton, R. [1985]: *Metaphysical and Epistemological Problems of Perception*, Lincoln, NB: University of Nebraska Press.

Fumerton, R. [1988]: 'The Internalism/Externalism Controversy', in J. Tomberlin (ed.), *Philosophical Perspectives*, **2**, pp. 443–59.

Fumerton, R. [1995]: *Metaepistemology and Skepticism*, Lanham, MD: Rowman & Littlefield.

Giere, R. [1988]: *Explaining Science*, Chicago: University of Chicago Press.

Ginet, C. [1975]: *Knowledge, Perception, and Memory*, Dordrecht: Reidel.

Goggans, P. [1991]: 'Epistemic Obligations and Doxastic Voluntarism', *Analysis*, **51**, pp. 102–5.

Goldman, A. [1976]: 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, **73**, pp. 771–91.

Goldman, A. [1978]: 'Epistemics: The Regulative Theory of Cognition', *Journal of Philosophy*, **75**, pp. 509–23.

Goldman, A. [1979]: 'What Is Justified Belief?' in G. Pappas (*ed.*), *Justification and Knowledge*, Dordrecht: Reidel, pp. 1–23.

Goldman, A. [1980]: 'The Internalist Conception of Justification', in P. A. French, T. E. Uehling and H. Wettstein (*eds*), *Midwest Studies in Philosophy*, **5**, pp. 27–51.

Goldman, A. [1986]: *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.

Goldman, A. [1988]: 'Strong and Weak Justification', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **2**, pp. 51–69.

Goldman, A. [1991]: 'Epistemic Folkways and Scientific Epistemology', in *Liaisons*, Cambridge, MA: MIT Press, pp. 155–75.

Goldman, A. [1994]: 'Naturalistic Epistemology and Reliabilism', in P. A. French, T. E. Uehling and H. Wettstein (*eds*), *Midwest Studies in Philosophy*, **19**, pp. 301–20.

Goldman, A. [1999]: 'Internalism Exposed', *Journal of Philosophy*, **96**, pp. 271–93.

Greco, J. [1990]: 'Internalism and Epistemically Responsible Belief', *Synthese*, **85**, pp. 245–77.

Greco, J. [1993]: 'Virtues and Vices of Virtue Epistemology', *Canadian Journal of Philosophy*, **23**, pp. 413–32.

Greco, J. [1999]: 'Virtue Epistemology', *Stanford Encyclopedia of Philosophy*, URL: < http://plato.stanford.edu/entries/epistemology-virtue/ > .

Haack, S. [1993]: *Evidence and Inquiry*, Oxford: Blackwell.

Haack, S. [1997]: '"The Ethics of Belief" Reconsidered', in L. E. Hahn (*ed.*), *The Philosophy of Roderick Chisholm*, La Salle, IL: Open Court, pp. 129–44.

Hall, R. J. and Johnson, C. R. [1998]: 'The Epistemic Duty to Seek More Evidence', *American Philosophical Quarterly*, **35**, pp. 129–39.

Harman, G. [1968]: 'Knowledge, Inference, and Explanation', *American Philosophical Quarterly*, **5**, pp. 164–73.

Harman, G. [1973]: *Thought*, Princeton: Princeton University Press.

Harman, G. [1980]: 'Reasoning and Evidence One Does Not Possess', in P. A. French, T. E. Uehling and H. Wettstein (*eds*), *Midwest Studies in Philosophy*, **5**, pp. 163–82.

Heil, J. [1983a]: 'Believing What One Ought', *Journal of Philosophy*, **80**, pp. 752–65.

Heil, J. [1983b]: 'Doxastic Agency', *Philosophical Studies*, **43**, pp. 355–64.

Heil, J. [1984]: 'Doxastic Incontinence', *Mind*, **93**, pp. 56–70.

Heil, J. [1992]: 'Believing Reasonably', *Noûs*, **26**, pp. 47–62.

Heller, M. [1989]: 'Relevant Alternatives', *Philosophical Studies*, **55**, pp. 23–40.

Heller, M. [1999a]: 'The Proper Role for Contextualism in an Anti-Luck Epistemology', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **13**, pp. 115–30.

Heller, M. [1999b]: 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy*, **77**, pp. 196–208.

Jacobson, S. [1997]: 'Externalism and Action-Guiding Epistemic Norms', *Synthese*, **110**, pp. 381–97.

Kelly, T. [2000]: 'The Ethics of Belief', Ph.D. Thesis, Harvard University.

Kim, J. [1988]: 'What Is "Naturalized Epistemology"?' in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **2**, pp. 381–405.

Kim, K. [1993]: 'Internalism and Externalism in Epistemology', *American Philosophical Quarterly*, **30**, pp. 303–16.

Kitcher, P. [1992]: 'The Naturalists Return', *Philosophical Review*, **101**, pp. 53–114.

Korcz, K. A. [1997]: 'Recent Work on the Basing Relation', *American Philosophical Quarterly*, **34**, pp. 171–89.

Kornblith, H. [1982]: 'The Psychological Turn', *Australasian Journal of Philosophy*, **60**, pp. 238–53.

Kornblith, H. [1983]: 'Justified Belief and Epistemically Responsible Action', *Philosophical Review*, **92**, pp. 33–48.

Kornblith, H. [1993]: 'Epistemic Normativity', *Synthese*, **94**, pp. 357–76.

Kornblith, H. [1995]: 'Naturalistic Epistemology and Its Critics', *Philosophical Topics*, **23**, pp. 237–55.

Kornblith, H. [1999]: 'In Defense of a Naturalized Epistemology', in J. Greco and E. Sosa (*eds*), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 158–69.

Laudan, L. [1987]: 'Progress or Rationality? The Prospects for Normative Naturalism', *American Philosophical Quarterly*, **24**, pp. 19–31.

Laudan, L. [1996]: *Beyond Positivism and Relativism*, Boulder, CO: Westview.

Lewis, D. [1996]: 'Elusive Knowledge', *Australasian Journal of Philosophy*, **74**, pp. 549–67.

Ludlow, P. and N. Martin. [1998]: *Externalism and Self-Knowledge*, Stanford: CSLI Publications.

Luper-Foy, S. [1987]: *The Possibility of Knowledge: Nozick and His Critics*, Totowa, NJ: Rowman & Littlefield.

McGinn, C. [1984]: 'The Concept of Knowledge', in P. A. French, T. E. Uehling and H. Wettstein (*eds*), *Midwest Studies in Philosophy*, **9**, pp. 529–54.

McGrew, T. [1995]: *The Foundations of Knowledge*, Lanham, MD: Littlefield Adams.

Maffie, J. [1990a]: 'Naturalism and the Normativity of Epistemology', *Philosophical Studies*, **59**, pp. 333–49.

Maffie, J. [1990b]: 'Recent Work on Naturalized Epistemology', *American Philosophical Quarterly*, **27**, pp. 281–93.

Meiland, J. [1980]: 'What Ought We to Believe? The Ethics of Belief Reconsidered', *American Philosophical Quarterly*, **17**, pp. 15–24.

Montmarquet, J. [1993]: *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman & Littlefield.

Moser, P.K. [1985]: *Empirical Justification*, Dordrecht: Reidel.

Moser, P.K. [1987]: *A Priori Knowledge*, Oxford: Oxford University Press.

Moser, P.K. [1989]: *Knowledge and Evidence*, New York: Cambridge University Press.

Naylor, M.B. [1985]: 'Voluntary Belief', *Philosophy and Phenomenological Research*, **45**, pp. 427–36.

Nozick, R. [1981]: *Philosophical Explanations*, Cambridge, MA: Harvard University Press.

Nozick, R. [1993]: *The Nature of Rationality*, Princeton: Princeton University Press.

Plantinga, A. [1988]: 'Positive Epistemic Status and Proper Function', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **2**, pp. 1–50.

Plantinga, A. [1993a]: *Warrant: The Current Debate*, New York: Oxford University Press.

Plantinga, A. [1993b]: *Warrant and Proper Function*, New York: Oxford University Press.

Pollock, J. [1983]: 'Epistemology and Probability', *Synthese*, **55**, pp. 231–52.

Pollock, J. and Cruz, J. [1999]: *Contemporary Theories of Knowledge*, 2nd edn, Totowa, NJ: Rowman & Littlefield.

Pryor, J. [2000]: 'The Skeptic and the Dogmatist', *Noûs*, **34**, pp. 517–49.

Quine, W. V. O. [1969]: 'Epistemology Naturalized', in *Ontological Relativity and Other Essays*, New York: Columbia University Press, pp. 69–90.

Radcliffe, D. [1997]: 'Scott-Kakures on Believing at Will', *Philosophy and Phenomenological Research*, **57**, pp. 145–51.

Riggs, W. D. [1997]: 'The Weakness of Strong Justification', *Australasian Journal of Philosophy*, **75**, pp. 179–89.

Schiffer, S. [1996]: 'Contextualist Solutions to Skepticism', *Proceedings of the Aristotelian Society*, **96**, pp. 317–33.

Schmitt, F. [1999]: 'Social Epistemology', in J. Greco and E. Sosa (*eds*), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 354–82.

Scott-Kakures, D. [1994]: 'On Belief and Captivity of the Will', *Philosophy and Phenomenological Research*, **54**, pp. 77–103.

Sellars, W. [1956]: 'Empiricism and the Philosophy of Mind', in H. Feigl and M. Scriven (*eds*), *Minnesota Studies in the Philosophy of Science*, **1**, pp. 253–329.

Sosa, E. [1991]: *Knowledge in Perspective*, Cambridge: Cambridge University Press.

Steup, M. [1988]: 'The Deontic Conception of Epistemic Justification', *Philosophical Studies*, **53**, pp. 65–84.

Steup, M. [1999]: 'A Defense of Internalism', in L. Pojman (*ed.*), *Theory of Knowledge: Classical and Contemporary Readings*, Belmont, CA: Wadsworth, pp. 373–84.

Steup, M. [2001]: *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Oxford: Oxford University Press.

Stich, S. [1990]: *The Fragmentation of Reason*, Cambridge, MA: MIT Press.

Stine, G. C. [1976]: 'Skepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies*, **29**, pp. 249–61.

Swain, M. [1981]: *Reasons and Knowledge*, Ithaca: Cornell University Press.

Swain, M. [1985]: 'Justification, Reasons and Reliability', *Synthese*, **64**, pp. 69–92.

Triplett, T. [1990]: 'Recent Work on Foundationalism', *American Philosophical Quarterly*, **27**, pp. 93–116.

Unger, P. [1986]: 'The Cone Model of Knowledge', *Philosophical Topics*, **14**, pp. 125–78.

Vogel, J. [1990]: 'Are There Counter-Examples to the Closure Principle?' in M. D. Roth and G. Ross (eds), *Doubting: Contemporary Perspectives on Skepticism*, Dordrecht: Kluwer, pp. 13–27.

Vogel, J. [1999]: 'The New Relevant Alternatives Theory', in J. Tomberlin (*ed.*), *Philosophical Perspectives*, **13**, pp. 155–80.

Williams, B. [1970]: 'Deciding to Believe', in B. Williams [1973], *Problems of the Self*, Cambridge: Cambridge University Press, pp. 136–51.

Williams, M. [1996]: *Unnatural Doubts*, Princeton: Princeton University Press.

Williams, M. [1999a]: *Groundless Belief*, 2nd edn, Oxford: Blackwell.

Williams, M. [1999b]: 'Skepticism', in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 35–69.

Williamson, T. [1995]: 'Is Knowing a State of Mind?', *Mind*, **104**, pp. 533–65.

Williamson, T. [1997]: 'Knowledge as Evidence', *Mind*, **106**, pp. 717–41.

Winters, B. [1979]: 'Believing at Will', *Journal of Philosophy*, **76**, pp. 243–56.

Wolterstorff, N. [1997]: 'Obligations of Belief: Two Concepts', in L. E. Hahn (*ed.*), *The Philosophy of Roderick Chisholm*, La Salle, IL: Open Court, pp. 217–38.

Wright, C. [1985]: 'Facts and Certainty', *Proceedings of the British Academy*, **71**, pp. 429–72.

Wright, C. [1991]: 'Skepticism and Dreaming: Imploding the Demon', *Mind*, **100**, pp. 87–115.

Wright, C. [2000]: 'Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', in E. Villanuova (*ed.*), *Philosophical Issues*, **10**.

Wright, C., Smith, M., and Macdonald, C. [1998]: *Knowing Our Own Minds: Essays in Self-Knowledge*, Oxford: Oxford University Press.