

4

Value and Epistemic Normativity

Most writers would agree that beliefs are governed by norms, that there is a right way and a wrong way of forming and relinquishing beliefs. ‘Being governed by norms’ has at least two connotations. First there is the idea that these norms should *guide* belief’s comings and goings. Second there is the idea that where these norms are violated, it is (generally) appropriate to *criticize* the believer for this. Discussion of the form of guidance appropriate to epistemic normativity must be left for another occasion. Here I assume only that the acquisition and extinction of belief should in some sense be guided by these norms. As to criticism, those who violate epistemic norms are stigmatized as irrational, but what does that criticism involve and why is it appropriate?¹ Criticism seems like an appropriate reaction to an (illicit) threat to some human interest, either one’s own or someone else’s. So where epistemic irrationality is criticizable, we should expect belief and its rationality to be something that matters to us, something that we value, something whose presence or absence bears on our interests.

In this chapter I attempt to ground epistemic normativity in human interests, in facts about what is good for us. One might try to demonstrate that it is in our interests both to have beliefs and to conform those beliefs to epistemic norms. Yet it would frequently be in our interests to adopt quite irrational beliefs or to relinquish perfectly rational beliefs (e.g., to boost our confidence and avoid depression). Some writers have concluded from this that it is a mistake to ground epistemic normativity in value.² I instead infer that the authority of epistemic norms cannot lie in the value of actually conforming to them. Rather, it lies in the value of being *subject* to those norms, whether or not you actually conform to them.³ Before making that case I shall explore some more familiar ways of grounding epistemic norms in value.

¹ Elsewhere I’ve argued that when we stigmatize someone as irrational we are blaming them (see Introduction: Section I.5 and Owens 2000: Chapter 8), but here I remain neutral on the nature of rational criticism.

² For example (Raz 2011: 41–7) and (Shah 2013: 316).

³ This is my solution to the problem of authority raised in Chapter 2, p. XX.

4.1 The Bayesian Challenge

Let's begin with this claim:

Evidence: Believe that p only if you have sufficient evidence for the truth of p .

Evidence for a proposition is any indication that the proposition is true. Where you have no evidence at all or the evidence you have is insufficient, your belief is irrational. I am hoping something like *Evidence* will prove attractive to anyone who acknowledges the existence of norms of rationality governing the formation of belief. To keep *Evidence* as uncontroversial as possible, I leave it entirely open what it is for someone to 'have' evidence for a proposition, what form of psychic contact or availability is involved. I also leave it open what is involved in p 's being evidence for q : the indication relation may be a priori or a posteriori, necessary or contingent. Evidence may not even be distinct from the thing it is evidence for: perhaps some propositions are self-evident in that our grasp of them is a sufficient indication of their truth. *Evidence* simply requires that you believe only where you have a sufficiently strong indication of the truth of what you believe.⁴

Though I have tried to make *Evidence* as vague and inoffensive as possible, some will still balk. Bayesians will agree that anything worth calling a belief must be based on some indication of the truth of the proposition believed, but *Evidence* goes further. It postulates a dichotomy between belief and its absence and it postulates a threshold of evidential sufficiency which the believer must cross before they are entitled to be convinced of p 's truth. The difference between conviction and agnosticism so understood is like the difference between standing to someone's right or to their left, a difference marked by a (perhaps rather vague) boundary. Bayesians recognize no such boundary.

On the Bayesian model, being a good believer involves proportioning one's level of confidence in p to the strength of one's evidence for p , believing p to the appropriate degree (and not- p to a reciprocal degree). There is no need to settle on an 'all-out' belief in either proposition:

One could easily enough define a concept of acceptance which identified it with high subjective or epistemic probability (probability greater than some specified number between one-half and one), but it is not clear what the point of doing so would be. Once a subjective or epistemic probability value is assigned to a proposition, there is nothing more to be said about its epistemic status. Bayesian decision theory gives a complete account of how probability values, including high ones, ought to guide behaviour, in both the context of inquiry and the application of belief outside of this context.

⁴ There may also be norms requiring us to believe in the face of sufficient evidence.

So what could be the point of selecting an interval near the top of the probability scale and conferring on the propositions whose probability falls in that interval the honorific title ‘accepted’? (Stalnaker 1984: 91)

Stalnaker asks: what is the point of forming beliefs? What good do they do us? What function do they perform in human life? Without an answer to these questions, we have no reason to take the norms associated with belief seriously and thus have no reason to form psychological states that involve our subjection to them.⁵ Bayesians like Stalnaker do have definite views about what constitutes good evidence, views which fill out the Bayesian injunction to proportion one’s beliefs to the available evidence; they acknowledge norms of coherence, norms whose authority is grounded in the practical needs to which Stalnaker alludes (i.e., the guidance of behaviour and inquiry). But these norms do not include anything like *Evidence*.

In the next section, I’ll consider a response to Stalnaker’s challenge that derives *Evidence* from another norm, namely the requirement that our beliefs be true:

Correctness: Believe that p only if p is true.⁶

The response is a two-stage affair. First, postulate (*pace* Stalnaker) a psychological threshold between believing and not-believing such that (as *Correctness* says) we should be on the belief side of the threshold only in respect of true propositions. Second, argue that because belief is governed by *Correctness* it is also governed by *Evidence*. In Section 4.2 I’ll establish that even granted *Correctness*, we have no reason to think that belief is also governed by *Evidence*. The way to *Evidence* is not via *Correctness*.

4.2 Alethic Conceptions of Belief

Some maintain that *Correctness* captures the distinctive normative character of belief (Shah and Velleman 2005: 499). It certainly differentiates merely imagining

⁵ Some contend that certain norms are constitutive of belief, that to have a belief at all is to be subject to standards like *Evidence* (e.g. Shah and Velleman 2005: 510). On this constitutivist view, to believe that p *just is* to be required to have sufficient evidence for p and so there can be no question about whether belief ought to be governed by this norm. I doubt that answers Stalnaker’s question. He can simply ask after the point of forming beliefs so conceived. Perhaps we are compelled to believe regardless, but then we are like compulsive chess players who should do whatever it takes to stop playing.

⁶ Bayesian degrees of confidence are not governed by *Correctness*. Your confidence in p is meant to reflect the level of evidence you have for p and so, unless that level of evidence is 1, coherence requires you to have a reciprocal degree of confidence in not- p . By contrast, a rational person will believe either p or believe not- p or neither (Ross and Schroeder 2012: 17–19).

that p from believing that p : there is no requirement that we imagine only what is true. It also imposes a requirement of consistency on our beliefs: since reality is consistent, an inconsistent set of beliefs must include false beliefs. But *Correctness* and *Evidence* seem quite distinct, for you might conform to *Correctness* without conforming to *Evidence*: good luck can ensure that you get the truth by accident. And you might also conform to *Evidence* without conforming to *Correctness*: reasonable error is ubiquitous.

It is natural to suppose that *Correctness* and *Evidence* are connected as follows: the best way of conforming to *Correctness* (i.e., to believe only truths) is to conform to *Evidence* (i.e., to believe only on the basis of sufficient evidence). For the purposes of our discussion, I'll imagine that the rational believer is trying to conform to *Correctness* by consciously following *Evidence*, but nothing hangs on the nature of the guidance mechanism. Our conformity to these norms is always non-voluntary and may well be unconscious.⁷

Let's begin our discussion of this proposal by asking why we should take *Correctness* seriously. In line with our general hypothesis about value and normativity, that should involve there being value in having a true rather than a false belief. Now, if it matters to you to have a belief on a certain issue, isn't it better that the belief in question be true rather than false? In a particular case there may be various advantages to having a false belief, advantages that might outweigh the value of truth: still, isn't there something to be said in favour of getting it right? You may not care whether the number of stars in the galaxy is odd or even, which is why you have no view on the matter, but where you do care enough to form a view, mustn't you see some point in getting it right, some interest that is served by having a true rather than a false belief?

There are two problems with this account of the normative force of *Correctness*. First, once the side effects of having a true rather than a false belief are taken off the table, the value of truth becomes rather obscure. Why should you wish, for its own sake, to be right rather than wrong? Suppose I believe on whatever grounds that I'm going to die tomorrow. Is there something to be said in favour of my death, namely, that it will vindicate my belief, that it will at least prove me right? Does my belief give me (or anyone else) any reason at all to ensure that I do die?⁸ This sounds like a strange idea, but we may have misconstrued our interest in truth. Perhaps the value lies in our beliefs *tracking* the truth, in being

⁷ We have no choice about whether to form beliefs on a certain matter, at least once the evidence is in, and we also have no choice about whether to retain that belief so long as the evidence remains (Raz 2011: 44) and (Shah 2013: 316–17).

⁸ The question is pressed by (Raz 2011: 45) and (Piller 2009: 196–200).

sensitive or responsive to how things are in the world rather than in a brute match between world and content, however achieved. Let's suppose that this is so; the project of grounding *Evidence* in the value of truth now confronts a second difficulty.

Where I do see some point in getting it right, how much trouble I should take to get it right? Suppose I am trying to spell 'assessment' correctly, then I must see some value in spelling it correctly, but it remains a question exactly how hard I should try. If time is short and I must issue a written warning that you are about to be assessed, it would be silly for me to worry *too* much about precisely how many 's's the word contains. My goal of good spelling must be balanced against the other goals I am trying to achieve, and here correct spelling may count for rather little. All things considered, it may be sensible for me to write the word at a speed that makes it rather unlikely that I shall achieve my goal of spelling it correctly. This prudent haste is perfectly consistent with the fact that I was trying to spell it correctly.

This isn't how it is with belief. A believer is, we may suppose, trying to get it right, but one shouldn't weigh the value of truth against other relevant values in order to determine whether or what to believe. Take my belief that I'm going to die tomorrow and suppose that there is some value in my tracking the truth. Given that I am going to die, this means that there is some disvalue in my thinking I'm going to live, but there is also a great deal of disvalue in my believing I'm going to die: all the pain of anticipation. Now suppose that we are trying to settle what I ought to believe. Should we weigh the value of my getting this right against the disvalue of acquiring a painful belief? If so, we might conclude that I should remain agnostic in the face of overwhelming evidence of death or even become convinced that I'm going to survive given only the most tenuous grounds for hope. As a verdict on the rationality of belief, this can't be right. We can't allow the postulated value of truth-tracking to be weighed against the disvalue of future pain.⁹

This second difficulty should make us wonder how much guidance *Correctness* really gives us about when to believe.¹⁰ Whilst being guided by *Correctness* clearly involves being sensitive to the evidence at least to some extent, it need not involve

⁹ Note that this second difficulty can be generalized to threaten any attempt to ground the authority of an epistemic norm in the value of our conformity to it. Take *Evidence* and suppose there is some intrinsic value in basing one's beliefs on sufficient evidence. Whatever that value may be, we can't determine what it would be reasonable for them to believe by weighing the desirability of their doing so in accordance with *Evidence* against the desirability of their doing so in other ways.

¹⁰ As William James observed, *Correctness* does not actually require us to form any beliefs: you could conform to it by believing nothing at all (James 1956: 17–19).

being sensitive to the evidence in anything, like the way in which belief is sensitive to evidence (Owens 2003). Consider suspicion.¹¹ Suspicion, like belief, is governed by *Correctness*:¹² to suspect that O.J. murdered his wife is to be mistaken should he be innocent (however justified the suspicion and reasonable the error), but the evidence required for reasonable suspicion that p clearly differs from that required for reasonable belief in p.¹³ The same applies to hope and fear. Where p is false we speak of false hopes that p and unfounded fears that p; such hopes and fears are no less mistaken than incorrect beliefs. This is so even though false hopes and unfounded fears may, like incorrect beliefs, be perfectly reasonable, being well grounded in evidence. Yet both reasonable hopes and reasonable fears are built on much less evidence than is required for reasonable belief. So *Correctness* fails to capture the norms distinctive of belief, since *Correctness* does not tell us what constitutes evidence sufficient for the formation of belief.¹⁴

I'll conclude this section by asking whether we might avoid these problems by strengthening *Correctness* as follows:

Knowledge: Believe that p only if your belief would constitute knowledge of p.

Knowledge entails *Correctness* since knowledge entails truth. Perhaps knowledge is the value for which we have been searching, a value that can underwrite the authority of both *Correctness* and *Evidence*. Suspicions, hopes, and fears do not require knowledge. On the contrary, ignorance is a presupposition of suspicion, hope, and fear.

I agree that *Knowledge* is a norm that applies uniquely to belief (and belief-involving states). You can't think it reasonable for you to believe that p if you take yourself to be ignorant of whether p (Owens 2000: 40–1). Furthermore, this norm

¹¹ In Chapter 2 (p. XX) I used the example of a guess or conjecture to make the point. (Shah and Velleman 2005: 498 n.7) expressed doubts about whether guessing is a propositional attitude like belief. Suspicion is immune to such doubts. Note that to suspect that p is not just to hypothesize that p (*pace* (Shah and Velleman 2005: 512)). One can entertain and test the hypothesis that p whilst having a completely open mind about whether p is true, but if one suspects that p, one's mind is not entirely open (Greenspan 1988: 85–8).

¹² One might doubt that suspicions, hopes, and fears are subject to *Correctness* on the grounds that one can reasonably entertain inconsistent suspicions, hopes, and fears. The issue is complex, but on this point I see no difference between these attitudes and belief. To suspect that X did it and suspect that Y did it is usually a matter of suspecting a conspiracy or else of having a disjunctive expectation. Perhaps it can also be reasonable to maintain genuinely inconsistent beliefs and inconsistent suspicions, at least where the alternative is to make an arbitrary choice about which one to abandon. Be that as it may, inconsistent beliefs and suspicions are on a par.

¹³ Contrast the legal standard for reasonable suspicion (e.g., when arresting someone) with the legal standard for convicting them in court.

¹⁴ This conclusion should worry anyone who wants to characterize belief in terms of *Correctness* regardless of where they think the source of its authority lies.

indicates a plausible value. Isn't knowledge *pro tanto* better than ignorance, at least where it is important to have some view on the matter? This may be denied (Raz 2011: 43 n.15), but let's allow it for now. Still, given our theoretical ambitions, this is not a satisfying result, for knowledge is a value whose content can be specified only by reference to the satisfaction of the very epistemic norms whose authority we are seeking to ground in that value. Knowledge requires justified belief¹⁵ and *Evidence* specifies just how much evidence one must have in p's favour for one to be justified in believing that p. Given this, *Knowledge* must be explained in terms of *Evidence* and not vice versa.¹⁶ So, even assuming that knowledge is valuable for its own sake, we can't base the value of conforming to *Evidence* on the value conforming to *Knowledge*.

We can restate the point just made in terms of truth-tracking rather than evidence. Knowledge of p involves more than a simple match between the content of my belief and the character of the world; knowledgeable belief must track the truth, and a belief tracks the truth where it is based on some reliable indicator of its own truth. Such tracking is plausibly worth having; the problem is that one can have more or less of it. Reasonable suspicion tracks the truth to a certain extent, whilst reasonable belief hugs the truth more closely. We must explain the point of tracking truth in the way that (reasonable) belief tracks truth rather than in the way that (reasonable) suspicions, hopes, and fears do; we must say why people should want to get at the truth *in this particular way*.

We have been assuming that both truth and the value of truth can be understood without using normative notions. In that respect, *Correctness* is well placed to explain the content of other norms of belief formation and, for all I have said here, such an alethic account might well be part of the story about norms like *Evidence*. But it can't be the whole story and we can't complete that story just by reverting to *Knowledge*, for knowing that p involves having a justified belief in p, and to be justified in believing p is at least to conform to relevant epistemic norms such as *Evidence*. So the idea that knowledge has value seems to presuppose rather than support the idea that conformity to such epistemic norms has value. In Section 4.3, I'll consider one attempt to explain, rather than assume, the value of knowledge.

¹⁵ At least this is so when we are dealing with creatures for whom the question of justification arises. Perhaps cat-knowledge does not require justification. I have no view about whether human knowledge can be *analysed* in terms of justification and other factors.

¹⁶ Some epistemologists have attempted to explain the relevant notion of tracking or reliability in terms of non-normative notions such as probability, counterfactual sensitivity, and so forth. I can't examine these efforts here and only note that none have garnered general agreement. I shall assume that the notion of epistemic justification and, in particular, the notion of sufficiency invoked by *Evidence* cannot be analysed in non-normative terms.

4.3 Pragmatic Conceptions of Belief

While constructing an account of belief and knowledge, Edward Craig tells us that, ‘Fortunately there is a firmly fixed point to start from. Human beings need true beliefs about their environment, beliefs that can guide their actions to a successful outcome’ (Craig 1990: 11). For Craig, the value of true belief lies in its motivating successful agency, agency that achieves its objectives. Craig goes on to argue that true beliefs amounting to knowledge are precisely those most likely to guide action to a successful outcome, and so we can add that the value of knowledge lies in its role in motivating successful agency.

This suggestion about the source of epistemic normativity has indeed been an intellectual fixture, at least since the advent of Pragmatism. It is endorsed by American Pragmatists like Pierce, James, and C.I. Lewis; British Pragmatists like Ramsey and Braithwaite; and by many more recent writers.¹⁷ Stalnaker clearly considers the ‘guidance of behaviour’ to be a core function of belief, and the same applies within the recent literature on ‘pragmatic encroachment’. For example, Fantl and McGrath maintain that ‘the importance of the concept of knowledge’ resides in the fact that ‘it sets a meaningful lower bound on strength of epistemic position: your epistemic position regarding *p* must be strong enough to make it rational for you to act as if *p* is true’ (Fantl and McGrath 2007: 581), while Stanley and Hawthorn claim that ‘Where one’s choice is *p*-dependent, it is appropriate to treat the proposition that *p* as a reason for acting iff you know that *p*’ (Hawthorne and Stanley 2008: 578). The differences between these formulations will not matter here; I shall treat them both as characterizing knowledge by citing a norm that connects knowledge with agency. If we also assume that something like *Knowledge* is correct, we can derive a claim about belief, namely

Pragmatism: We are entitled to believe that *p* iff we are entitled to act as if *p* is true (or take *p* as a reason for action).

I want to highlight two attractions of *Pragmatism*. The first is that it connects belief to what many take to be the ultimate source of any form of normativity, namely, our ability to produce valuable states of affairs by means of our agency. On this view, false and ignorant beliefs are bad because they undermine our ability to pursue familiar goods like pleasure, beauty, and so forth. Agents are able

¹⁷ Some of these writers may have regarded pragmatic conceptions of belief as an alternative to alethic conceptions of belief. I’ll instead treat claims such as *Pragmatism* as supplementing or underwriting rather than as replacing a claim like *Correctness* or *Knowledge*. (Velleman argues that without an alethic component, we could not distinguish a belief from a fantasy and a desire from a wish (Velleman 2000: 255–77).)

to pursue those goods successfully when they can make the world satisfy the desires or fulfil the intentions on which they act. Such a match is just the converse of truth and no more requires normative notions for its specification than does truth itself.

A second attraction of *Pragmatism* is that it seems to do a good job of differentiating belief from the other propositional attitudes governed by *Correctness*. To be entitled to suspect (or hope or fear) that *p* is not to be entitled to act as if *p* is true. In certain contexts, one might be entitled to act on one's suspicions but Pragmatism states that one is always entitled to act on a (well-grounded) belief, to act as if one's belief is true, whatever the context. Thus, the norms governing belief in *p* will be those whose satisfaction always makes it sensible to act as if *p* were true. We now have an answer to our earlier question: what is the point of trying to get at the truth by forming beliefs, by coming to know the truth? Belief and knowledge provide a basis for making decisions about what to do, one on which we are entitled to rely in any context.

At this point Stalnaker would get off the train. He maintains that we don't need such 'all-out' beliefs in order to take decisions, that we could manage well enough with Bayesian degrees of confidence in the relevant propositions without seeking to divide those propositions into those that we know and those we don't. What practical purpose is served by setting an evidential threshold above which one may believe that *p* and below which one may not? Why not just proportion one's belief to the evidence, proportion one's willingness to behave as if *p* (or not-*p*) to the degree of one's belief in *p* (or not-*p*) and act accordingly? Isn't that how we determine what odds to accept in a betting situation?

The pragmatist replies that we can't keep track of the evidence for and against the innumerable propositions whose truth might matter to us.¹⁸ At some point we cease to deliberate, close the books, take a view before throwing away much of the evidence (so we don't have to store it), and act accordingly. How could it be otherwise, given the limits on our capacity to retain old information and consider the significance of new information as it arrives, while doing everything else we must do? Even where our conclusion has a probabilistic content it will still be an all-out belief, not a Bayesian level of confidence.¹⁹ In the eyes of the pragmatist we form beliefs for much the same reason that we set ourselves to do something

¹⁸ (Harman 1986: 38–42 and 46–9). For a more recent statement of the point, see (Ross and Schroeder 2012: 27–8).

¹⁹ It should be noted that 'all-out' attitudes like belief and intention are stronger or weaker, more or less firm. What makes these attitudes 'all-out' is the presence of a line between believing and not-believing, intending and not-intending. The strength of the attitude is a function of how easy it is to get the subject back across the line, to abandon the relevant belief or intention (Owens 2000: 142–5).

by forming an intention to do it: if we are to act effectively, we cannot always remain non-committal either about how things are (or are likely to be) or about what to do in the light of this.²⁰ The precise point at which we should make up our minds about how things are (at which the evidential threshold for rational belief has been reached), like the precise point at which we should settle what to do, will be fixed by the need to ensure successful agency.

One might wonder whether it is always a good thing for an agent to be able to execute her intentions successfully. Doesn't that rather depend on the quality of her intentions? I shall assume that there is always *something* to be said for having this power, even if, given the idiocy of one's intentions, one would on the whole be better off without it. What of propositions whose subject matter is of no practical relevance? For instance, some of us take enough of an interest in the number of stars in the galaxy to form beliefs about such things, though we may never have occasion to act on these beliefs. The pragmatist can respond that we rightly value the higher order capacity to determine which issues are practically relevant and which are not, a capacity that requires us to know much that is practically irrelevant at the first-order level. Let's give the pragmatist a pass on this, for the charms of *Pragmatism* are in any case illusory.

I'll start with its second attraction. Is the evidential threshold for knowledge (and so rational belief) really tied to the requirements of rational agency in the way the pragmatist suggests? It is not so clear, that it is always rational to act as if *p* simply because you know that *p*. In particular, this is not rational in some cases where (a) the costs of acting as if *p* should *p* turn out to be false are substantial or (b) the benefits of acting as if not-*p* should not-*p* turn out to be true are substantial.²¹ In these cases, though we know that *p*, we decline to act on our knowledge, we refuse to take *p* as a reason for action and so the pragmatist cannot use *Knowledge* to explain the norms of belief.

This morning I attended an examiner's meeting at which a poor degree was awarded to one of my students. I have a clear recollection of the result and I feel sorry that the student didn't do well. The next day I must write a formal letter to the student announcing the result and before writing the letter I check the list of student marks that was circulated at the meeting. Why check? Is it because

²⁰ This is perhaps less obvious when it comes to suspicions than it is with belief and intention, since we are quicker to abandon our suspicions in the light of new evidence. Still, when you suspect me of abusing your child, you don't have an open mind about whether I have abused your child and your suspicion informs your actions (and emotions). Were we never entitled to form such suspicions but had to keep all the available evidence in play until there was enough for belief, our lives would be very different.

²¹ After laying out the case for this in Chapter 7 (p. XXX), I read (Brown 2008: 175–82), from whom I borrow the exam result example.

I suddenly experience doubts, fearing that I might have mixed this student's mark up with the many others that were discussed at the meeting? That *might* be how it is; equally well it might not. Perhaps I instead reason as follows: though I know perfectly well that the student did poorly (still feel sorry about this and so forth) I also recognize that my memory is fallible, a further check can easily be made and given what hangs on the letter, I should make the check before writing it. The precaution seems eminently sensible.

I am sensible not to act on my recollection that the student did poorly because the costs of being wrong are substantial. The same point is made by cases in which I risk missing out on a considerable though unlikely benefit if I act on my belief. Suppose someone offers to pay me £10 million in return for a stake of ten pence if it turns out that I was not born male. I know that I was born male and much of the rest of what I know about myself would make little sense were I not. Nevertheless, I might reasonably accept the bet (Hawthorne 2004: 176). Can the mere fact that I have been offered this bet render one of my securest convictions unjustified? Must I confess ignorance of my original gender because it is silly to miss out on this bet?

In response, the pragmatist might weaken their position thus:

Default Pragmatism: To be entitled to believe that *p* is to be entitled to use *p* as a default assumption in one's practical reasoning.²²

In answering a question (whether practical or theoretical) we take various things for granted. The default pragmatist suggests that there are some propositions that we are entitled to take for granted in answering any given question unless we see a special reason to do otherwise, namely the propositions we are entitled to believe. And the propositions we do believe are those we actually do take for granted unless an alarm goes off, an alarm set to register situations in which it would be sensible to check or else to take a risk on it, etc. *Default Pragmatism* allows that one can fail to act as if *p* once an alarm goes off, whilst continuing to believe that *p*, provided one continues to be disposed to act as if *p* (without further consideration of whether *p*) when no alarm is sounding. Once the alarm has gone off one presumably considers what to do on the basis of other beliefs (some of which concern evidence for or against the truth of *p*), these other beliefs being used as default assumptions in the present context. That's what happens when one decides to purchase the lottery ticket or to check the list of exam results.

²² See Chapter 7 (p. XXX). Something like *Default Pragmatism* (though not under that name) is defended in (Ross and Schroeder 2012: 8–10 and 19–22).

I don't doubt that we rely on our beliefs in determining how to act unless there is a special reason to do otherwise. The problem with this claim is that it does not distinguish belief from a range of other attitudes. Some of these attitudes are like belief in that they are governed by evidential norms. If I suspect you of abusing my child then I'll refrain from placing them in your care, but I might not report you to the police: the consequences of so doing are such that they constitute a reason for not acting on a mere suspicion. Other such attitudes seem not to be governed by evidential norms. The lawyer is obliged to act as if her client is innocent, regardless of the strength of the evidence against him. Call this attitude *acceptance*. True, she will not act on this assumption outside certain professional contexts, but, equally, there are many contexts in which we are not prepared to act on our suspicions or on our beliefs.

For the default pragmatist, belief and suspicion are at bottom forms of acceptance. The difference between belief, suspicion, and the lawyer's acceptance lies in the range of contexts in which one is disposed to act on the relevant attitude. Wouldn't that suffice to explain why the evidential norms governing these attitudes differ? Indeed, but this explanation makes the difference between belief, suspicion, and acceptance out to be a matter of degree: it is just that we are both inclined and (with the relevant evidence) entitled to place greater reliance on the former. The idea that there is a categorical distinction between knowledge and ignorance—the idea that Fantl and McGrath, Stanley and Hawthorn etc. seemed to be trying to capture—has now escaped us. What remains is a continuum of propensities to assume that a certain proposition is true for the sake of action. So long as belief in *p* was taken to involve being disposed to act as if *p* in all contexts, belief occupied a very special position on this continuum. Having removed belief from that position, the default pragmatist is without an account of why the (range of) dispositions denoted by the word 'belief' should be of any special interest to us. Pragmatism starts to look like a more psychologically realistic form of Bayesianism.

In this section, I have been focused on the pragmatist's claim to be able to differentiate belief from other attitudes governed by *Correctness*. In Section 4.4, I'll present my own account of this matter and then I'll query *Pragmatism's* other attraction, the idea that belief's role in the guidance of agency is the ultimate ground of epistemic normativity because it is the source of belief's value.

4.4 Belief and the Passions

What is the point of believing things, of seeking to know them as opposed to merely suspecting them? In the rest of this chapter, I'll suggest that value of belief

and knowledge lies in its enabling a form of emotional engagement with the world (suspicion underwrites a rather different form of emotional engagement). We'll reach this conclusion by arguing that the evidential threshold at which it becomes rational to believe and so to claim knowledge of *p* is the point at which it would be appropriate to have a range of emotional attitudes towards *p* (rather than the threshold at which it would be rational to act as if *p* were true). The evidence required by *Evidence* must suffice to make it reasonable for me to feel a certain way.

Emotions come in at least two varieties. More numerous are those with the following feature: where the emotion in question has a propositional object, the subject of the emotion can feel that way only if they believe the proposition in question. It follows that they are entitled to feel that way only if they reasonably believe the proposition in question. Where I am angry that *p* I must believe that *p*, and my anger is reasonable only if *p* is a proposition in whose truth I reasonably believe. Suppose I am angry that Tom stole my bike. My anger is reasonable only if it is reasonable for me to believe that Tom stole my bike. Call such emotions *doxastic* emotions. Doxastic emotions include regret, resentment, horror, disgust, fury, sorrow, embarrassment, disappointment, shame, and on the positive side delight, gratitude, pleasure, and pride.²³

The claims just made might be further explained. Perhaps I can't even be angry that *p* unless I *know* (and thus believe) that *p*.²⁴ On this view, emotions like anger might better be called *epistemic* (rather than doxastic) emotions. Advocates of this view allow that anger can be misdirected; Tom may not have stolen my bike, and even if Tom did steal my bike, I may not know this (perhaps this news came from an informant who, though credible, is only repeating a rumour that happens to be true). But if anger is an epistemic emotion, it would be wrong to say that I am *angry that* Tom stole my bike since I cannot be angry about things of which I am ignorant. Rather, I am angry *in the belief* that Tom stole my bike (Unger 1975: 189–96).

I register no objection to claim that anger is an epistemic emotion, but nor do I wish to rely on it here because my point does not turn on the correct usage of idioms like 'angry that *p*'. Perhaps 'I am angry that Tom stole my bike' can just mean 'I am angry in the belief that Tom stole my bike'. This would be enough to

²³ Though my list of doxastic emotions is taken from Gordon, Gordon describes anger etc. as a *factive* emotion, where a factive verb is one that applies only where its propositional complement is true. That does not capture the connection to reasonable belief that interests me. And he applies the label 'epistemic emotion' to the non-doxastic emotions (Gordon 1987: Chapter 2).

²⁴ I endorsed this view in Chapter 7 (p. XXX). Here I remain agnostic.

explain why I can't reasonably be angry that Tom stole my bike unless I reasonably believe that Tom stole my bike. Furthermore, given that I can't reasonably believe that p where I think I'm ignorant of whether p , it would also explain why I can't think my anger at Tom is reasonable unless I think I know that Tom stole my bike. These implications are enough for my purposes and I shall continue to call anger, etc. doxastic emotions.

Non-doxastic emotions include hope and fear together with their variants such as worry. These emotions are like their doxastic cousins in that they may take a propositional object and where this is so, they too are governed by *Correctness*; a false hope (or fear) is a mistaken hope (or fear). But hope and fear exclude knowledge.²⁵ One is merely hopeful (rather than joyful) or fearful (rather than devastated) precisely when one does not take oneself to know the final outcome.²⁶ Once I realize that the prize is going to someone else, I can no longer hope to win it, nor can I hope to win it once it is revealed that I have won it.

Both doxastic and non-doxastic emotions come in degrees: joys and regrets, hopes and fears are all stronger or weaker and all more or less intense, but this similarity conceals an important difference. The strength of a non-doxastic emotion varies in two dimensions, the strength of a doxastic emotion in only one. Consider hope and fear. Hoping that you'll win involves suspecting that you'll win as well as wishing to win (and fearing that you'll lose involves suspecting that you'll lose as well as wishing you won't).²⁷ The more you wish to win, the more you hope that you will win. In this respect hope is like joy: the more you want the prize, the more you'll enjoy getting it. But the strength of your hopes also varies as your suspicion that you'll win waxes and wanes. As the evidence that you will win piles up, your hopes rise; as doubts intensify, your hopes fade. There is no corresponding dimension of variation in the case of joy. For joy there is an evidential threshold: either the evidence fails to convince you that you have won, in which case joy is ruled out or else it succeeds in convincing

²⁵ Wollheim criticizes Gordon's account of what I call non-doxastic emotions on this point. He argues that one can be frightened or terrified of what one *knows* will happen (e.g., death). Sometimes it is the uncertain timing, character, or consequences of one's death that one fears, but one can also be terrified of death as such. Perhaps emotions anticipating unpleasant outcomes like 'fear', 'terror', and 'dread' take both doxastic and non-doxastic forms, a difference not clearly marked in English. See (Wollheim 1999: 106–10).

²⁶ Note that one can have doxastic emotions involving propositions with a probabilistic content or a modal content, e.g., I can be pleased that I have a 90 per cent chance of winning the lottery or pleased that I might win the lottery.

²⁷ Gordon tells us that to hope that you'll win the prize you must (a) be in doubt as to whether you have won it and (b) wish to win it (Gordon 1987: Chapter 4), yet isn't hoping to win a matter of being in some degree hopeful that you will win, of having hopes of winning that might be dashed? Mere doubtfulness about winning seems less than hopeful.

you and then you should experience joy in the degree made appropriate by the desirability of its object (Adler 2002: 217).

Someone can be angry or happy at the fact that *p*, or proud of it, or grateful for it, only if they believe that *p*. This suggests the following hypothesis: often we want a view about whether *p* in order to fix our emotional bearings, to avoid having our feelings baffled by ignorance. In eliminating uncertainty we learn how to feel, not just how to act. Emotional bafflement (as well as practical uncertainty) renders doubt painful and drives inquiry. True, we might sometimes prefer to stick with those emotions—hopes and fears—that presuppose uncertainty rather than learn the truth. But very often we're anxious to discover whether we have a fatal illness and not just so we can make the appropriate arrangements. We want to know how to feel about the situation.

The pragmatist might respond that these roles are complementary: belief both sets our emotional bearings and guides our practical deliberation, but, as we saw earlier, there are cases in which feeling as if *p* would be sensible though acting on one's knowledge of *p* would not and vice versa. I felt sorry that my student had done poorly in their exams even as I checked to make sure that they had. Here I feel sad because they have done poorly, but I'm hardly checking the list because they have done poorly. Furthermore, I can feel either pride or shame that I was born a man even as I accept the bet, and I can feel this way only because I believe that I was born male. Pride and shame need not come and go in response to such offers; rather, they are part of a more permanent background, dependent on relatively stable convictions which structure our emotional lives.²⁸

I've suggested that the value of belief and knowledge lie in their enabling certain forms of emotional engagement with the world. On my hypothesis, the difference between knowledge and ignorance matters to us because the capacity to have doxastic emotions matters to us. This leaves it wide open exactly why various aspects of the world should concern us in this fashion. Some facts might engage our emotions for instrumental reasons, others because of their intrinsic interest (e.g., the number of stars in the galaxy). Either way, we are curious only where our emotions are at least potentially aroused. Yet aren't many facts rather unlikely to engage our emotions? I read the newspaper in a calm and detached frame of mind, generally believing what I read. Here, I form views on matters of apparent indifference, and they are no less governed by epistemic norms than those I form on matters close to my heart. So how, it may be asked, can epistemic

²⁸ I agree with Fantl, McGrath, Stanley, and Hawthorn, etc. that how much evidence I need for knowledge of (or reasonable belief in) a proposition depends on how much its truth matters to me, but I deny that either is a function of its transient practical significance (Chapter 7: p. XXX).

norms derive their authority from the value of my emotional engagement with the relevant subject matter?

There is such a thing as mild curiosity. Some facts interest us for their own sake and not just because of their impact on our lives or their wider moral or scientific importance. We want to know because the matter interests us, without provoking anything like anger, disgust, or even hope and fear. We are faintly pleased or slightly amused when we discover what is happening on the other side of the world. Such mild curiosity is often quite enough to explain why I both entertain suspicions on a certain matter and/or want to know the truth about it: no more intense or practical concern is required.²⁹ But suppose I am indeed totally incurious about the exact number of pencils currently on my desk. If I do form a view about this, aren't I still subject to criticism should that view be based on insufficient evidence? I doubt it. Epistemic norms apply only where deliberation is a possibility, where there is some opportunity (usually not taken) to weigh the evidence, consider defeaters, etc.³⁰ This possibility exists only where it would make sense to deliberate, and it makes sense to deliberate only where there is some basis for curiosity on the point (which might just be curiosity about whether my perceptual mechanisms are working properly). Since there is *ex hypothesi* no basis for curiosity in the pencil situation, any attitude that I acquire on the matter can only be the spontaneous output of an automatic perceptual mechanism. As noted at the outset, I'm interested in norms that can guide belief formation, thereby providing grounds for a specific form of criticism: attitudes formed by reflex will not be stigmatized as irrational.

A second question now arises. Is it always in my interests to enjoy apt rather than inapt emotions, even on topics of indisputable significance? So long as life is good, the menu of apt emotions adds relish to the dish; once our fortunes turn, an appropriate emotional engagement with what is happening seems to make life worse rather than better. My friend dies and I am devastated. Wouldn't it be better for me if I somehow didn't learn of this or even deceived myself on the point? Won't my remaining friends feel more sorry for me once I'm forced to face the truth? Of course I shouldn't visit my friend if they aren't going to be there and

²⁹ Where something matters much more to one person than another, does that entail that the relevant evidential threshold for reasonable belief varies also? If so that would disrupt the workings of testimony, for in order to determine whether some competent and sincere person is a credible source on the point, I'd have to go into how much it happens to matter to them. I agree that the threshold cannot be determined in such an individualistic fashion. What sets the threshold is the social significance of the issue where this social significance is a matter of emotional significance of this sort of issue within the relevant community (Introduction: p. XX-XX).

³⁰ For discussion of what deliberation involves, see Introduction: pp. XX-XX and Chapter 3: pp. XX-XX.

such practicalities might come to matter more than peace of mind, but then again they might not. In any case, how can this terrible knowledge *as such* be good for me? That is the topic of the Section 4.5.

4.5 Epistemic Norms and Normative Interests

The obvious way to connect epistemic norms with value would be to argue that beliefs conforming to norms such as *Evidence* have a value that beliefs violating such norms do not. As we have seen, this isn't a promising line to take. Sometimes there is no evident value or interest in conforming your beliefs to norms like *Evidence*, and where there is such an interest it is often outweighed by considerations recommending breach of the norm. I shall instead argue that what is good for us is not conformity to epistemic norms, it is rather the fact of being *subject* to those norms, of being in a position where it makes sense for us to be guided by them and are otherwise vulnerable to criticism. The validity of the norms itself has value; hence the value of belief, of the psychological state in virtue of which we are subject to them. In this final section, I will try to characterize norms like *Evidence* in a way that makes our interest in them clear and do so without instrumentalizing them, without identifying some non-normative value which conformity to them promotes.

Following the lead of Section 4.4, one might argue for this as follows: (i) we value for its own sake the capacity to engage emotionally with the world in certain ways; (ii) belief is required for that capacity; (iii) to believe is to be subject to norms like *Evidence*; therefore (iv) there is value in being subject to those norms. Here a valuable capacity is partially constituted by a belief that in turn subjects us to certain norms. Unfortunately, the argument is flawed: from the fact that A is valuable for its own sake and the fact that A involves B (even constitutively) it does not follow that B is valuable for its own sake. For example, the fact that goods like forgiveness or repentance both constitutively require a wrong does not show that there is any good in the wrong, and we should regret the wrong even as we celebrate the reaction. But, I shall urge, the present case is different: there is something good both in having the capacity to engage emotionally with the world and in being subject to the norms which possession of that capacity involves. We should not regret these things even as we regret some of the emotions they require of us.

Consider (i). What I am supposing to be of value here is not, say, the experience of an apt grief but rather the capacity to engage emotionally with my friend's fortunes. To grasp the point, let's first set aside several things that the phrase 'capacity to engage emotionally with the world' might cover but which are

not in question. First, there is our capacity to enjoy the non-doxastic emotions: hope, fear, and their variants. Second, there are those simulacra of the doxastic emotions stimulated by reading novels and watching films or plays, emotions such as ‘grief’ and ‘admiration’ for the characters portrayed. Our life is surely enriched by these psychological phenomena and a good story about my friend would move me, but that story could not satisfy my curiosity about what is now happening to them now unless it enabled me to engage emotionally with the relevant bits of the world.³¹

A specific piece of knowledge (such as knowledge of what is happening to one’s friend) enables me to engage the fact known with a range of doxastic emotions. I maintain that, at least on issues that matter to us (like the fate of our friends) we have reason to value this specific topic-focused capacity and that is the ground of the value of particular items of knowledge. This need to engage emotionally with particular facts is registered in a feeling of frustration whenever our emotions are baffled by ignorance, when we are rendered incapable of pride, shame, joy, grief, and so forth because we don’t know the truth. It makes sense to want to know whether you were admired or despised by your long-dead brother, whether you are soon going to die yourself and so forth, even if there is little you can or would wish to do about it. You want to know how to feel about the past or the future. A hedonist might respond that what you really want is to avoid the pain of the frustration caused by ignorance, but, if so, that could be achieved by popping a pill. Since the pill isn’t the thing, the desire to know must be something more than the desire to avoid the psychic costs of ignorance; rather, it is a desire for some good whose absence is registered by these costs. It is also something other than the need for good news rather than bad. Of course, you wish it to be the case that your brother admired you and that you’ll lead a long and happy life, but, quite apart from that, you have some desire simply to know where you stand (and for its own sake).

These facts may seem puzzling. To know that my friend has died is to be placed in a situation in which a feeling of devastation is apt or even required of me and my life goes much worse should I feel an apt devastation rather than a quite inapt indifference. So why should anyone wish (for my sake) that I know whether my friend has died, when this knowledge will require me to feel devastated? Why should anyone wish that I be subject to a norm that I conform to only at great

³¹ There is a familiar controversy over whether imaginative literature, films, etc. generate genuine anger, fear, and so forth, given that they do not create the relevant beliefs. I’m inclined to side with writers like (Walton 1990: 195–204) and (Velleman 2000: 270) who argue that they do not, but for present purposes I could retreat to the claim that certain valuable *forms* of anger and fear, etc. are baffled by ignorance.

cost to myself? The worry is not just that the benefits of knowledge will be outweighed by the costs, that all things considered it would be better if I didn't know at all, that I would on the whole be better off were I to remain blissfully ignorant of my friend's death. The worry is rather that, at least if the intrinsic value of knowledge (and belief) lies in their emotional significance, knowing lacks *all* intrinsic value because to know is to be required to have a feeling that there is no (non-instrumental) good in having. Since devastation as such has no advantages over peace of mind, how can knowledge derive its value from the fact that it requires this of us?

We are assuming that human interests make sense of normativity in one way only, namely, by ensuring that it is good for us to conform to the norms in question. Suppose instead that human beings have normative interests, that amongst the basic goods of human life are normative phenomena themselves: obligations and permissions but also standards of aptness and appropriateness for beliefs, feelings, and so forth. Then a psychological state could derive at least part of its value from the fact that it subjects us to certain norms, even though conformity with those norms may sometimes have nothing to be said for it. In this respect, knowledge (and belief) are not unusual. I would argue that friendship itself derives part of its value from the fact that it subjects friends to certain norms, even though conformity with those norms may do them no good in any given instance.

Many writers agree that having friends is good for you and not just because friends provide you with help and support. Friendship is good for its own sake and thereby enriches your life. One prominent feature of friendship is that it brings into play a whole set of norms. Becoming John's friend changes the normative situation between us, creating reciprocal rights and obligations and altering what feelings are appropriate between us or even required. It is once John has become a close friend that I should be devastated by his death, that I would be a bad friend if I weren't. What attitude should I adopt towards the norm of friendship that requires me to feel devastated should John die? I might regard this norm as an unfortunate aspect of an otherwise good thing, a feature perhaps inextricably bound up with the depth of our friendship but not as itself making any positive contribution to its value. Were that my attitude, I might well regret that my friendship with John would require me to feel devastated should he die. But isn't regret an inappropriate reaction to the realization that I am so required? Wouldn't it show that I failed to value the friendship and the bonds of loyalty it involves correctly? The point is *not* that I should feel glad when this norm comes into play because John dies. When this comes about I am sorry mainly for John but also for myself and others would surely pity me. What is not regrettable is the

underlying fact that, because of our friendship, I would be required to feel this way should John die. On the contrary, that normative fact is part of the good of a deep friendship, a friendship which serves our interest in being subject to such norms (Owens 2012: 111–17).

I've spoken of the norm requiring me to feel sad at John's death 'coming into play' when John dies, but that isn't quite right, for I'm subject to that norm only once I discover that John has died. In the context of friendship, knowledge establishes a normative connection between John and myself, making devastation the appropriate reaction to John's death. Friendship does not always require me to find out whether John is still alive whatever the cost, but it makes perfect sense for me to wish to know whether he is still alive, and to wish to know this because he is my friend, because his death will have a certain emotional significance for me. I may value the knowledge that requires a certain reaction of me even as I regret the feeling of devastation it demands and even though there are times when I deliberately remain in ignorance of his fate so I can focus on other things. And I value the relevant knowledge for its own sake because it partly constitutes (and does not merely cause me to acquire) a capacity for engagement that I value for its own sake.

I've said that believing involves being subject to various norms. Clearly, there is more to belief than that: believers have a certain psychology, they are disposed to think and behave in certain ways and they are subject to epistemic norms only in virtue of that psychology. A rock or an amoeba lacks the responsibilities of a believer because it lacks the required psychological sophistication. Furthermore, epistemic norms have their value (and so their authority) only in that specific psychological setting; the value and validity of these norms is dependent on the non-normative context in which they find application. For example, the argument of Section 4.4 suggests that belief and its norms might be no good to a pure Bayesian agent, one that has the internal machinery needed to weigh evidence and act accordingly but can feel nothing. In this respect, epistemic norms are once more similar to norms of friendship. The latter apply and have value only against the background of a relationship that they partially constitute, a relationship also constituted by certain habits of feeling and behaviour. Beings incapable of acquiring such habits would have no interest in these norms.

My central contention is that belief and knowledge can be valuable for their own sake because they serve certain normative interests. Some of these interests come into play in special contexts like that of friendship. Others have wider application, like my interest in being able to engage emotionally with whether I have won a certain prize. Unless I know the relevant facts, I am not subject to these norms, I am not emotionally engaged with the relevant subject matter and

the value of knowledge lies in its making this connection. By explaining where the value of belief and knowledge lies, we have thereby established the authority of the norms that govern believing (e.g., *Evidence*) and a similar story may be told about those governing suspicion. When something engages our curiosity, we'll often value the capacity to entertain hopes and fears about it even whilst we remain ignorant of the ultimate truth. Suspecting, like believing, involves being subject to norms specifying what evidence provides adequate grounds for suspicion, and the value of suspicion is, in part, the value of being subject to those norms.

I'll conclude the chapter by considering two objections to my account of the value of knowledge. First, one may wonder whether the above line of thought really establishes the value of knowledge (over ignorance) as well as that of belief (over agnosticism). The objector is correct that I have been moving freely between the two, a liberty to which I felt entitled so long as I was discussing a Bayesian who denies the value of both belief and knowledge and on much the same grounds. But once the Bayesian has been answered in respect of belief, one might wonder why *Knowledge* is true, why a belief must constitute knowledge to have the value of belief. Suppose I truly believe that I have won the prize. Whether or not this belief constitutes knowledge, it still gives me the capacity to engage emotionally with a fact about myself, a capacity quite distinct from my capacity to engage emotionally with an acknowledged fiction like a play or a film.³² To form the belief that I won is to subject myself to various norms of evidence, etc., norms to which I must subject myself if I am to be in a position to react with pride (and so forth) to my win. And if the latter capacity has a value (one that the capacity to engage with films and fiction lacks) then I have an interest in entering a psychological state that subjects me to the norms that govern belief formation. Why should we suppose that I have any further interest in knowledge?

I agree that belief is better for me than mere agnosticism whether or not that belief constitutes knowledge and that this would be enough to answer Stalnaker's challenge, vindicate *Evidence* (though not *Knowledge*) and bring my main line of argument to a successful conclusion. Whether the argument can be extended from belief to knowledge depends on a question left hanging earlier: can I be angry that Tom stole my bike only once I *know* that he stole my bike? If the answer is 'yes', merely believing this, however truly or justifiably, will not put my anger into emotional contact with the fact: it will leave me angry only in the belief that Tom stole my bike. But if the answer is 'no', then, for all I have said, knowledge has no value beyond the value of true belief.

³² (Wollheim 1999: 105–6) makes a parallel point about anger.

Second objection: isn't the value of knowledge, however understood, often outweighed by other values? Mightn't I be better off with the irrational belief that my friend is still alive, even if I do miss out on some valuable form of emotional engagement with the world? But the norms of belief formation and retention surely apply to such cases as much as to those in which the balance of advantage lies with belief in accordance with the evidence, and they alone determine whether belief is rational. Recall the objection I raised in Section 4.2 against those who sought to ground the authority of *Correctness* in the value of true belief. We imagined someone weighing the value of truth against whatever values would be served by false belief in order to determine what the subject should believe, all things considered. This procedure might endorse quite irrational belief. For the same reason, it may be thought, the value of knowledge cannot be the source of the authority of our epistemic norms. The strictness of these norms remains to be explained.

Here it is crucial to distinguish two different issues: (a) whether you should subject yourself to certain epistemic norms and (b) whether you should conform to those norms given that you are subject to them. The question raised by (a) is a practical question, a question about what one ought to do. Where (a) is the issue one should indeed weigh the value of being subject to those norms, the value of being able to engage emotionally with the relevant facts, against the disvalue involved in actually conforming with those norms. If I have an important exam tomorrow and I have reason to suspect that something terrible has happened to John (something that I can do nothing about) it may be perfectly sensible for me to wait before looking into the matter, before acquiring the capacity to engage emotionally with his fate and so placing myself in a situation where I might be required to believe the worst. That capacity is valuable for its own sake, but its value not absolute. Here ignorance may be the better option.

The question raised by (b) can also be understood as a practical question, as asking whether you ought to do something to ensure that your beliefs conform to the norms to which they are subject, perhaps by taking some training in scientific methodology or visiting your analyst to address your tendency to wishful thinking. Here the practical considerations just rehearsed again apply and it may well be that you ought not to take such steps (the analyst costs money) and should instead allow your mental processes to run their perhaps irrational course. Here the value of conforming to epistemic norms must be weighed against other relevant values to determine what one should do all things considered.

Once these practical questions have been settled, there is no further issue about whether the beliefs you have formed are subject to epistemic norms. The value relevant to the authority of epistemic norms over believers is the value of our

being subject to those very norms (rather than the value of conformity with them), and that value is not to be weighed against the disvalue of conformity in order to determine whether we are subject to those very norms. To say that a believer is ‘subject to epistemic norms’ is just to say that the process of belief formation is intelligible in a way that makes them vulnerable to criticism should they fail to believe in accordance with those norms (Introduction: pp XX–XX): the emotional costs of conformity may give the believer some excuse for wishful thinking, but they provide no justification for it. Epistemic norms are indeed strict in that any believer is subject to those norms regardless of the costs of conformity.³³

³³ I am grateful to audiences at Columbia University, the Universities of Geneva, Fribourg, Luxembourg, Southampton, and Kings College London. Many thanks also to Alex Gregory, Mark Schroeder, Nishi Shah, Matthew Silverstein, Sharon Street, David Velleman, Clayton Littlejohn, and Jose Zalabardo for comments on drafts.

