

# Introduction

Epistemology has always been, at least in part, a normative discipline. It is in the business of prescription, not mere description. It characterizes certain states and practices as good, bad, or indifferent. Not content to say how the world is, it aims to say how it should be. The normative dimensions of epistemology are a consequence of the two concepts which are standardly used to define the subject, *knowledge* and *justified belief*.

Aristotle famously said that “all men by nature desire to know” (*Metaphysics*, 1.980a). He recognized that knowledge seeking is a pervasive feature of human life. Knowledge is desirable for instrumental reasons (i.e., as a means to other goods), but it also seems to be, at least sometimes, an intrinsic good as well (i.e., something that is good whether or not it leads to further goods). Accordingly, one of the central tasks of epistemology has been to investigate how we should go about acquiring this good.

The normative implications of the concept of *justified belief* are equally clear. To call a belief justified is, *at the very least*, to give a consideration in favor of that belief. I will be adopting and defending a stronger position in this book, sometimes called “the guidance conception of justification” (Pollock, 1986, p. 10), according to which to

call a belief justified is simply to say that we ought to believe it, and to say that a person is justified in believing something is simply to say that he or she ought to believe it.

Whether epistemology is understood as an investigation into how to acquire knowledge or as an investigation into what we ought to believe, it is natural to compare it to ethics, which has long been the paradigm of a normative discipline within philosophy. It is particularly natural to compare the problem of working out what we ought to believe, to the problem which Plato thought of as the central problem of ethics, that of working out how we ought to live (*Republic*, Book 1, 352).<sup>1</sup>

Analogies of this kind have recently led several philosophers to develop epistemological theories that are explicitly inspired by ethical theories. But although epistemology has been willing to turn to ethics for theoretical guidance, it has been much more reluctant to follow the lead of ethics in another way. Whereas the work of philosophers like Peter Singer and Jonathan Glover has transformed the study of ethics in recent decades, by addressing contemporary social and technological issues, the study of epistemology remains quite abstract and ahistorical. It is true that some epistemologists have applied their theorizing to contemporary issues (and a handful have even called that work “applied epistemology”), but applied epistemology, unlike applied ethics, remains an obscure and underdeveloped subject. For many people, including many professional philosophers, “applied philosophy” is virtually synonymous with “applied ethics.”

This view of the scope of applied philosophy is not inevitable, and one of the chief tasks of this book will be to argue that it is too limited. The information revolution and the knowledge economy have radically changed the way that we acquire knowledge and justify our beliefs. These changes have altered our epistemic landscape as surely as the sexual revolution and breakthroughs in reproductive technology have changed our moral landscape. The latter changes provided a good deal of the impetus for the applied turn in ethics, but the former changes have so far failed to result in a comparable turn in epistemology. Such a turn is surely inevitable, and this book aims to contribute to it. Hence this book will not only do some applied epistemology, it will also constitute an argument for the importance of applied epistemology as an emerging field of research.

It has often been observed that the applied turn in ethics of the last few decades is not a radical departure from tradition, but a rediscovery of an earlier approach to the subject.<sup>2</sup> This would be equally true of

an applied turn in epistemology. Although the expression “applied epistemology” seems to be quite new, the practice of applying epistemology to issues of topical concern is not. One classic example of applied epistemology is David Hume’s argument against belief in miracles (1966/1748, Section 10). The question of whether one could be justified in believing stories of miracles from the Bible and other religious or historical texts was of great interest to Europeans in the middle of the eighteenth century.<sup>3</sup> Epistemological considerations have also played a significant role in the history of political philosophy. For example, John Locke (1999/1689) argued for religious toleration, partly on the grounds that no government can be sure that the official religion is correct, which means that no government can be sure that it is not persecuting the true religion. Likewise, John Stuart Mill (2008/1859) argued that since no one person has infallible access to the truth, we are most likely to converge on the truth in the course of debate sustained by laws protecting free speech.

In recent decades, however, epistemology has been somewhat sidelined in political philosophy, which has increasingly come to be thought of as a branch of ethics. This is an overly narrow view. Politics *does* raise ethical issues, which political philosophy has rightly taken on. But it also raises epistemological issues, which, in recent years, political philosophy has been less willing to address.<sup>4</sup> For example, although some of the public debates about the 2003 invasion of Iraq concerned ethical issues (e.g., the principles of just war theory), most of them concerned epistemic issues (e.g., the nature of the evidence for weapons of mass destruction, or what we could know of the real intentions of the governments prosecuting the war). For the most part, philosophers, including those who emphasize the relevance of philosophy to contemporary political issues, denied that philosophy had anything to contribute to the latter debates. I hope this book will encourage them to rethink this attitude.

Before doing applied epistemology, we need a bit of theory to apply. In what follows, I will survey some of the theoretical positions to be found in contemporary normative epistemology. I will not end up endorsing any of these positions or offering any unified theoretical position of my own. There are two reasons for this. First, I am not particularly interested in making a case for any purely theoretical position. This is a work of applied philosophy. As such, I want my arguments to be persuasive to a wide range of readers with a wide range of theoretical commitments.<sup>5</sup> Second, I genuinely think that an

eclectic approach is best. Each of the theoretical positions I consider will offer some genuine insight, despite their limitations, and will serve as valuable instruments if used with care.

## **Veritism**

Alvin Goldman has recently proposed a normative epistemological theory called “veritism,” which is inspired by consequentialist ethical theory in general, and utilitarianism in particular. Knowledge occupies the role in veritism that happiness or utility occupy in utilitarianism; just as happiness or utility is the one intrinsic value in traditional utilitarian ethics, knowledge is the one intrinsic value in veritistic epistemology. In both cases, other valuable items, “such as actions, rules, or institutions, are taken to have instrumental value insofar as they tend to produce states with fundamental value” (Goldman, 1999b, p. 87).

Before evaluating veritism we need to briefly consider the nature of the value on which it focuses. What is knowledge, that is, what does the word “knowledge” mean? Attempts to define knowledge go back at least as far as Plato.<sup>6</sup> Since Plato, knowledge has standardly been characterized as a species of belief; knowledge is belief with certain (desirable) characteristics. Although there has been great controversy over precisely what those characteristics are, there has been no real controversy about one of them, truth. If someone knows *p*, then *p* must be true. But although truth is clearly a necessary condition for knowledge, most philosophers have followed Plato in denying that it is sufficient. Indeed Plato’s way of framing the problem of defining knowledge, as the problem of specifying what distinguishes *mere* true belief from knowledge, has dominated the literature.<sup>7</sup>

Despite having made some very influential contributions to this literature, Goldman has recently challenged the principal assumption on which it is based, the Platonic view that truth is not sufficient for knowledge. He argues that very often talk of knowledge is just talk of true belief:

The sentence “You don’t want to know what happened while you were gone” seems to mean: You don’t want to have the truth about what happened in your belief corpus. It does not seem to require the translation: You don’t want to have a justified belief in the truth about what

happened. So I believe there is an ordinary sense of “know” in which it means “truly believe.” (Goldman, 1999b, p. 25)

Although Goldman recognizes that this is not the only way in which the word “know” is used, he stipulates that it is knowledge in this “thin” sense that is accorded fundamental or intrinsic value in veritism.

Goldman does an excellent job countering a variety of arguments that truth is unachievable, or not a fundamental or intrinsic value. Throughout this book I will assume what I take to be the commonsense view, that he is right about this: truth (or at least a reasonable approximation to it) is quite often achievable, and truth (or approximate truth) is often valued, and valuable, for its own sake.

Veritism is, however, susceptible to criticisms which closely resemble certain common criticisms of the utilitarianism on which it is modeled. These criticisms come in two broad categories: first, those which insist that there is more than one intrinsic value that normative epistemology should be concerned to maximize or promote, and second, those which insist that there are “value side-constraints”<sup>8</sup> on *how* the value or values in question should be pursued. I will consider criticisms of both kinds.

## Error Avoidance

We have seen that veritism characterizes acquiring true belief (i.e., knowledge in the thin sense) as the only intrinsic epistemic good. Where does that leave the value of avoiding false belief (i.e., avoiding error)? Goldman considers the possibility that the acquisition of true belief and the avoidance of false belief might be two distinguishable values, but rejects the idea, claiming that there is a way to unify them by blending them “into a single *magnitude* or *quantity*” (Goldman, 2002, p. 58).

This requires him to translate the “categorical” approach to epistemology in which we talk only of belief, disbelief, and suspension of judgment (or having no opinion) into the “degrees of belief” (or subjective probability) approach in which beliefs are given values between 0 and 1, where 1 represents complete subjective certainty that the proposition in question is true and 0 represents complete certainty that it is false. Goldman says that we should identify belief *simpliciter* with having a degree of belief close to 1, while disbelief *simpliciter* is identified with having a degree of belief close to 0, and suspension of judgment (or having no opinion) is identified with having a degree of belief close to

## Introduction

0.5. Now, consider a particular true proposition,  $p$ . Having a high degree of belief in this truth is equivalent to having a low degree of belief in the falsehood  $\text{not-}p$ . So the value of true belief (i.e., of having a high degree of confidence in a particular truth) is equivalent to the value of error avoidance (i.e., of having a low degree of confidence in its negation). Thus seeking truth and shunning falsehood are, Goldman argues, simply two ways of looking at the same thing.

Goldman's argument identifies having no opinion with suspending judgment (and with having a roughly 0.5 degree of belief).<sup>9</sup> But this is clearly wrong. Someone who has never heard of Caracas will have no opinion about whether Caracas is the capital of Venezuela. But that person has not suspended judgment, nor does she have a (roughly) 0.5 degree of belief. You can neither suspend judgment about, nor have any degree of belief in, a proposition you have never considered. Goldman presupposes that everyone has some degree of confidence – high, low, or in between – in every proposition. But this is not true. For each of us, there are many propositions which we neither believe nor disbelieve, and which we do not assign any subjective probability either. And it is precisely because some strategies lead us to form more beliefs than others that there can be a conflict between the goals of truth seeking and error avoidance.

Strategies which exclusively emphasize error avoidance will inevitably lead to fewer beliefs than strategies which exclusively emphasize truth acquisition. William James recognizes this often overlooked point in the following passage:

There are two ways of looking at our duty in the matter of opinion, – ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. *We must know the truth*; and *we must avoid error*, – these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws ... Believe truth! Shun error! – these we see are two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life. We may regard the chase for truth as paramount, and the avoidance of error as secondary; or we may, on the other hand, treat the avoidance of error as more imperative, and let truth take its chance. (James, 2007/1897, Part VII, pp. 17–18)

Of course Goldman recognizes that the absence of true belief (being ignorant) is different from the presence of false belief (being

misinformed). But his unified approach to epistemic value leads him to the position that these “vices,” though distinguishable, can be “ordered” in a way which implies that the former is less epistemically or intellectually vicious than the latter:

Does this scheme of veritistic value accord with commonsense notions about intellectual attainments? I think it does. If a person regularly has a high level of belief in the true propositions she considers or takes an interest in, then she qualifies as “well-informed.” Someone with intermediate levels of belief on many such questions, amounting to “no opinion,” qualifies as uninformed, or ignorant. And someone who has very low levels of belief for true proposition – or equivalently, high levels of belief for false propositions – is seriously misinformed. Since the terms “well-informed,” “ignorant,” and “misinformed” seem to reflect a natural ordering of intellectual attainment, our scheme of veritistic value seems to be on the right track. (Goldman, 2002, p. 59)

But it is not at all clear that commonsense is on Goldman's side here. Although being well-informed<sup>10</sup> is certainly, at least *prima facie*, a higher level of intellectual attainment than being either ignorant or misinformed, it is not so clear that we should rank being ignorant ahead of being misinformed. Should we say that someone who has no opinion about whether Caracas is the capital of Venezuela has achieved more intellectually than someone who believes that Caracas is not the capital of Venezuela?<sup>11</sup> I submit that, without some contextual guidance, our intuitions simply get no traction on the question. I conclude that Goldman is wrong to suppose that ignorance and error can be put on a single scale on which ignorance is ranked higher than error. In Chapter 6 I will argue that Goldman's view that error is inherently worse than ignorance leads him to express unjustifiable concerns about the epistemic dangers of the blogosphere.

## Proceduralism

One could accept the argument up to this point while retaining the spirit, if not the letter, of veritism. Truth acquisition and error avoidance, though distinguishable, are closely related; furthermore, they are both characteristics of the outcome of inquiry (or its absence), rather than of the way in which inquiry is (or is not) conducted. In other words, the argument so far is compatible with epistemic consequentialism, though

not of the unified kind preferred by Goldman. In what follows, I will argue that epistemic consequentialism fails to do justice to the *intrinsic* value of acquiring, holding, or avoiding beliefs in some ways rather than others. In particular, it fails to do justice to the intrinsic value of doing these things in rational or justified ways, rather than irrational or unjustified ways. This *proceduralist* critique of consequentialist epistemology is reminiscent of critiques of consequentialist ethics, which insist that there is a moral value to certain kinds of actions, which is (at least partially) independent of their consequences or anticipated consequences.<sup>12</sup>

Epistemic consequentialists will not of course deny that the way in which we acquire, hold, and/or avoid belief matters. Like most people, they will insist that the process by which we do these things should (at least usually) be justified, by means of evidence, argument, and so on. Nonetheless, they will insist that the value of justification (or at least the value of “epistemic justification”<sup>13</sup>) is purely instrumental. Laurence Bonjour expresses this view in the following passage:

The basic role of justification is that of a *means* to truth ... if epistemic justification were not conducive to truth ... if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth. It is only if we have some reason for thinking that epistemic justification constitutes a path to truth that we as cognitive beings have any motive for preferring epistemically justified beliefs to epistemically unjustified ones. (Bonjour, 1985, pp. 7–8)

It is not obvious that Bonjour and Goldman are right about this. In particular, it is not obvious that we should think of justification as valuable only insofar as it helps us to believe what is true or avoid believing what is false. Consider an example to test your intuitions about this issue.

David Lewis (2000) has discussed what he takes to be a puzzling feature of the way academics are appointed to philosophy departments. From the premise that universities “exist for the sake of the advancement of knowledge” (p. 187) he draws the following conclusion:

By and large and *ceteris paribus*, we would expect the materialists in the philosophy department to vote for the materialist candidate, the dualists to vote for the dualist, and so forth ... I say this not out of cynicism. Rather, it seems to be how they *ought* to vote, and unabashedly, if they are sincere



in their opinions and serious about doing the best they can, each by his own lights, to serve the advancement of knowledge. (Lewis, 2000, p. 189)

But, of course, this is not how they typically behave. Rather than openly promoting the views which they consider right by appointing those who agree with them, “an appointing department will typically behave as if the truth or falsehood of the candidate's doctrines are weightless, not a legitimate consideration at all” (p. 190). Lewis explains this attitude by postulating a tacit treaty between academics with opposing views. According to the terms of this treaty, those with truth on their side should “ignore the advantage of being right” and not promote their own views in return for those who do not have this “advantage” agreeing not to promote their views. We ignore the truth of a particular candidate's doctrines:

Because if we, in the service of truth, decided to stop ignoring it, we know that others, in the service of error, also would stop ignoring it. We have exchanged our forbearance for theirs. If you think that a bad bargain, think well who might come out on top if we gave it up. Are you so sure that knowledge would be a winner? (Lewis, 2000, p. 200)

Despite the ingenuity of this argument, I doubt many people will be convinced by it. They are likely to object to the premise that the promotion of true belief is the fundamental value for which universities and philosophy departments exist. The natural thing to say is that it is *the process* by which beliefs are acquired and justified that is of fundamental or intrinsic value rather than, or at least as well as, the outcome of that process.

If this were not the case, it seems we should assess students' essays on the basis of the truth of the positions they argue for. But, of course, this is precisely what we struggle to avoid. Instead we try to assess them on the basis of how well they justify (or rationally defend) the positions they argue for, whether or not those positions are true. Again, it seems that when we do this, we are committed to the view that it is the process by which students arrive at and defend their conclusions which is of fundamental importance, rather than the truth of those conclusions.

There are, of course, responses that Lewis, Goldman, Bonjour, et al. could make to these objections. But I won't pursue the issue further here. It is enough to note that a purely consequentialist approach to epistemology seems counterintuitive for reasons which closely

resemble the reasons many have found purely consequentialist approaches to ethics counterintuitive.<sup>14</sup> What is more, even if epistemic consequentialism is correct, not everyone is liable to be persuaded of its correctness in the foreseeable future and I don't want nonconsequentialists to stop reading now. Hence I will not adopt a purely consequentialist approach in this book.

Nor will I adopt a purely proceduralist approach. Such an approach would, to use Goldman's words, treat justification as "sharply disconnected from truth" (Goldman, 2002, p. 55). Throughout this book I will assume that any form of inquiry which consistently fails to arrive at the truth, or which consistently leads to falsehood, cannot be justified. Like truth, justification seems to be both intrinsically and instrumentally valuable.

## **Other Values?**

I have argued that a plurality of fundamental or intrinsic values have a role to play in normative epistemology. Are there any values that I have not considered? The value of happiness (one's own or that of other people) certainly seems to be a consideration people sometimes take into account when considering what to believe. What is more, it is a value that can notoriously come into conflict with the value of truth as well as the value of justification (or at least the value of epistemic justification). This is the basis of Pascal's argument for belief in God. According to Pascal, we should have this belief, not because it's likely to be true (he claims that it is just as likely to be true as not), and not because it is justified by the evidence (he denies that it is), but because (roughly speaking) it is in our interests to believe.<sup>15</sup>

To avoid theological controversies, I will use a slightly different example, adapted from William James's "The Will to Believe." Jones is being pursued by a wild animal, when his path is blocked by a deep canyon. The only way he can hope to survive is by jumping across it, but all the available evidence suggests the canyon is too wide and he is unlikely to make it. However, he is enough of a psychologist to know that if he believes he will succeed, he is more likely to succeed. It seems to me clear that in these circumstances, Jones ought to believe the following proposition:

I will succeed in jumping over the canyon.

Objections to this claim are likely to come from two directions. Some philosophers will object on the grounds that it would be impossible for him to believe (a). Since (it is often said), ought implies can, it cannot be the case that he ought to believe something which he cannot believe. I will take up this position, as part of a broader discussion of the degree to which our beliefs are under our control, in the next section. For now, I will assume what I take to be the commonsense view, that Jones might be able to get himself to believe (a). The question is "Should he believe it?"

A position called *evidentialism*, often associated with W. K. Clifford, holds that we are always obliged to form our beliefs in accordance with the available evidence. In this case, the available evidence does not support (a), so an evidentialist will say that Jones should not believe (a), no matter what advantages belief may bring. Clifford's argument for evidentialism appeals to broadly veritistic considerations, pointing out the many advantages that truth can bring to both individuals and societies. He claims that anyone who holds a belief that is contrary to the available evidence is setting a bad example, and thereby committing a "sin against mankind" (Clifford, 1947/1877, p. 77).

I am confident that this is usually true, but the current example seems to be a clear exception to the rule. While there may be a sense in which anyone who takes the principled stance of refusing to even try to believe (a), on the grounds that it's not supported by the evidence, is being rational, it should also be clear that there is another broader, and in this case decisive, sense in which they are being irrational. A procedure that leads Jones to believe (a) would be justified in the sense that concerns us in this book. Jones would be justified in believing (a) because, all things taken into account, Jones should believe (a).

Many epistemologists would concede this, but brush aside such "pragmatic" considerations as beyond the proper scope of epistemology. On their view, what Jones should, all things considered, believe is a question for decision theorists or ethicists, not epistemologists. I see no reason to narrow the concerns of epistemology in this way. Epistemology, and especially applied epistemology, should be broad enough in its concerns to engage with any factors that may be relevant to the issue of what we should believe. Again the analogy with ethics is useful. Some people think of normative ethics as offering a specific kind of advice, which can be weighed against other potentially action-guiding considerations, for example, prudential or aesthetic considerations. The problem with this approach to ethics is that it leaves

the point of the subject obscure and open to endless contention. It is better to think of ethics, in the broad and traditional way, as a subject concerned with how we should live, all things considered. Likewise epistemology (or at least normative epistemology) should be concerned with what we should believe, all things considered.

## **Controlling Our Beliefs?**

The above discussion assumed that Jones could believe (a) if he wanted. But many philosophers would challenge this assumption. They claim that we have little or no control over our beliefs. If they are right, this is a problem, not merely for that example, but also for the analogy between ethics and normative epistemology with which we began. In fact, it seems to me that much of the resistance to the development of applied epistemology as a subject on a par with applied ethics<sup>16</sup> comes from a view, found among many working epistemologists, that whereas we can control our actions (the domain of ethics), we have little or no control over our beliefs (the domain of epistemology). I say this is wrong. I say that we often exercise voluntary control over our beliefs, just as we often exercise voluntary control over our actions. This makes me a proponent of what Alvin Goldman has called “the dubious thesis of doxastic voluntarism” (Goldman, 1999a, p. 275).<sup>17</sup> The first thing to say in defense of doxastic voluntarism is that it seems to be supported by common sense. We often talk, for example, of deciding what to believe about this or that matter or of struggling to decide what to believe about some other matter.<sup>18</sup> I say that such talk should be taken at face value. The burden of proof is on the other side.

Arguments against doxastic voluntarism go back at least as far as Hume, who claimed that belief “depends not on the will, but must arise from certain determinate causes and principles, of which we are not masters” (Hume, 1967/1740, p. 624). Hume’s position is intimately connected with the theory of ideas, which he inherited from the empiricist tradition. This tradition has a notoriously passive conception of human psychology, which few contemporary philosophers are likely to fully endorse.<sup>19</sup> Nonetheless, many contemporary philosophers do adopt a passive conception of the psychology of evidence evaluation, according to which the available evidence either leads us to form a certain belief or it does not. Choice does not come into it. As Richard Feldman puts it, “we are at the mercy of our evidence” (Feldman, 2001, p. 83).<sup>20</sup>

This claim is sometimes generalized to cover the full range of doxastic attitudes, that is, not only our beliefs, but also the degree of confidence with which we hold those beliefs. According to Neil Levy, for example, propositions have “an immediate subjective probability for us. In the light of the evidence available to us, we are simply struck by the likelihood that a given proposition is true or false” (Levy, 2007a, p. 136). I shall argue shortly that this completely passive conception of belief formation is wrong. Even if it were right, however, and we have no control over what we believe *given our evidence*, it would not follow that we have no control over what we believe. Levy himself effectively acknowledges this when he says that “though we may not be able to control what beliefs we form on the basis of our evidence we can certainly control our evidence gathering activities (if we can control anything)” (Levy, 2007a, p. 144). Although Levy characterizes himself as an opponent of doxastic voluntarism, he concedes that we do in this way exercise some “indirect” control over what we believe. We exercise some control over our evidence-gathering activities, and hence over our eventual beliefs, when we choose to investigate one subject rather than another, as well as when we choose to conduct our investigation in one way rather than another. Opponents of doxastic voluntarism tend to be dismissive of this kind of indirect control, claiming that it is of marginal significance and perhaps not really a form of control at all.<sup>21</sup> But this is a mistake. Many (indeed arguably almost all) of the things we control we control indirectly. I am certainly in control of the words that appear in this book (if I am in control of anything) but my control is indirect. I control them by controlling other things, such as the movement of my fingers over my keyboard. Even the movement of my fingers is (at least arguably) indirect, since I control it by sending signals via my nervous system to the muscles in my fingers.<sup>22</sup> So even if it were true that we only have indirect control over our beliefs, that would not imply that this control is less extensive than, or fundamentally different in kind from, the control we ordinarily take ourselves to have over most of our actions.

So far I have assumed, for the sake of argument, the correctness of the view that *given our evidence* we have no control over what we believe. Now I want to challenge it. On the face of it, our beliefs are determined, not only by our evidence, but also by the attitude we adopt towards that evidence. Hence, on the face of it, we can exercise control over our beliefs, not only by controlling our evidence-gathering activities, but also by choosing to adopt one rather than another attitude to the

strength of the evidence we end up with. We often talk, for example, of *deciding* that some evidence is compelling, or that that some evidence should be ignored. Of course, this way of speaking might be misleading. Perhaps when we talk of someone deciding that a piece of evidence is compelling, we really mean that he or she forms the belief (nonvoluntarily of course) that the evidence in question is compelling. Similarly, perhaps when we say that someone chooses to ignore a piece of evidence, we really mean that he or she forms the belief (again nonvoluntarily) that it isn't really evidence at all. Such moves are always possible, but they should be motivated by something more than a dogmatic commitment to the view that given our evidence we cannot help believing what we believe.

Although Bernard Williams (1973) has argued that it is logically impossible to directly control our beliefs, most opponents of "believing at will" accept that it is logically possible; they just think that as a matter of fact no one can do it. In the following passage William Alston presents an influential "argument" to that effect:

My argument for this, if it can be called that, simply consists in asking you whether you have any such powers. Can you, at this moment, start to believe that the United States is still a colony of Great Britain, just by deciding to do so? If you find it too incredible that you should be sufficiently motivated to try to believe this, suppose that someone offers you \$500,000,000 to believe it, and you are much more interested in the money than in believing the truth. Could you do what it takes to get that reward? (Alston, 1989, p. 122)

Many philosophers have found this persuasive. Nikolaj Nottelmann, for example, endorses Alston's view, observing that no matter how hard he might try he cannot bring himself to believe that there is a blue cat in front of him (2006, p. 560).

Speaking for myself, I cannot bring myself to believe either that the United States is still a colony of Great Britain or that there is a blue cat in front of me. But it does not follow from the fact that there are limits to what we can choose to believe that we cannot choose to believe at all, or even that we have little choice about what to believe. There is no disanalogy between beliefs and actions here. It does not follow from the fact that I cannot lick my elbow (try it) or fly to Mars that I have little or no control over my actions.<sup>23</sup>

I expect many people will be unconvinced, and continue to insist that neither they nor anyone else can ever choose what to believe.

Such people can believe whatever they like.<sup>24</sup> In this book it will be presupposed that we have significant control, not only over our evidence-gathering activities, but also over what we end up believing, and hence that when I offer advice about what to believe (or how to go about believing), readers will be able to decide for themselves whether or not to take that advice.

## **Duties to Believe and Responsibility for Belief**

William Alston's argument against doxastic voluntarism is one step in an argument against what he calls "the deontological conception of justification," according to which, to say that a belief is justified is to say that it is one's duty or obligation to believe it. Since, he claims, having a duty to believe something implies having control over whether or not to believe it, and (according to him) we do not "in general" have such control, it follows that we do not in general have duties or obligations to believe, and that justified belief should not be understood in these terms (Alston, 1989, p. 196). Needless to say, I do not accept this argument, because I do not accept the premise that we do not in general have control over what we believe. I say we have duties to believe certain things. Some of these duties may be peculiar to the circumstances in which we find ourselves, others are (almost) universal. It seems clear, for example, that we all have a duty to believe that the Holocaust occurred, because this belief is well-evidenced, true, and an important reminder to us all of what people can do.

Nonetheless, I don't accept the deontological conception of justification. On my view, to say that a person's belief is justified is simply to say that he or she ought to believe it. In other words, as I said earlier, I endorse the guidance conception of justification. But although the guidance conception of justification tends to be closely associated with the deontological conception of justification, they really should be distinguished.<sup>25</sup> Intuitively, it seems clear that we are not duty-bound to believe everything that we ought to believe, any more than we are duty-bound to do everything we ought to do.<sup>26</sup> To start with, if one has a duty to believe (or, for that matter, do) something, that implies that it is important that one believe (or do) it, and not everything that one ought to believe (or do) is particularly important. Furthermore, duties are, at least typically, duties to other people, and many of the things we ought to believe do not in any significant way concern other

## Introduction

people.<sup>27</sup> We ought to believe everything we have a duty to believe, but the converse is not true. We do not have a duty to believe everything we ought to believe.

Arguments against doxastic voluntarism have not only been used to challenge commonsense views about our duties to believe, they have also been used to challenge commonsense views about our responsibility for our beliefs. Although Neil Levy concedes that we may *sometimes* be responsible for our beliefs, he claims that such responsibility is rare, much rarer than we usually think (2007a, p. 149). He uses the example of Dr Fritz Klein, a concentration camp doctor who tried to reconcile his role in the Holocaust with his Hippocratic oath by claiming (and, we will assume, believing) that “The Jew is a gangrenous appendix in the body of mankind”<sup>28</sup> to illustrate his position. Levy is aware that most people would view Klein as blameworthy, not merely for his actions, but also for the beliefs that motivated his actions (indeed it is hard to see how he could be responsible for his actions unless he was responsible for the beliefs that led to them). Nonetheless, Levy thinks it is unlikely that Klein was responsible for this belief.

As we have seen, Levy thinks the only control we have over our beliefs is indirect, through our control of our evidence-gathering activities. Hence, according to him, any responsibility we have for our beliefs is also indirect. But, he goes on to argue, such indirect responsibility is rare, because it is inconsistent with certain common attitudes, one of which is moral certainty:

When I am morally certain that *p*, I cannot be expected to gather further evidence for or against *p*. There must be some doubt to set me on the path to checking or altering my beliefs by evidence-gathering. Certainty precludes the need (in my eyes) for me to engage in evidence-gathering, and thereby excuses me of indirect responsibility for my beliefs. (Levy, 2007a, p. 145)

Since, according to Levy, indirect responsibility is the only kind of responsibility we can have for our beliefs, it follows that holding a belief with certainty excuses us of any responsibility for it at all. This is, as Levy recognizes, extremely counterintuitive. Intuitively, it seems clear that moral certainty would not excuse Klein of responsibility for his belief. If anything, it would make him more culpable.<sup>29</sup>

I have already argued against the view that we only have indirect control over our beliefs; hence I don't accept Levy's view that we can



only ever be indirectly responsible for our beliefs. But I would also insist that even if Klein did only have indirect control over his belief, moral certainty would not excuse him of responsibility for it. Levy assumes that, although Klein may have some control over his belief (through his control of his evidence-gathering actions or omissions), he has no control over the degree of confidence with which he holds that belief. But the very evidence-gathering activities or omissions that lead him to hold the belief also lead him to whatever degree of confidence he has in it. Hence, if Klein were morally certain of his belief, he may well have been responsible, not only for having the belief, but also for being so certain of it.

I am not saying that Klein definitely was responsible for his belief, or for his certainty (if he was certain). Perhaps he was subject to brainwashing from an early age, which left him literally incapable of considering the possibility that his belief might be false. In such circumstances (arguably) he would not be responsible. But in such circumstances, it would be *the reasons* for his moral certainty that (arguably) would excuse him, not *the mere fact* that he was morally certain.<sup>30</sup>

I conclude that the commonsense view that we are typically, though not always, responsible for our beliefs, just as we are typically, though not always, responsible for our actions, is correct. Indeed the reason we are typically responsible for our actions is that they are to a large extent the result of beliefs for which we are responsible. Because people are typically responsible for their beliefs, we may be justified in praising or blaming others for their beliefs or being proud or ashamed of our own beliefs. Other responses associated with responsibility may not be justified. It is (at least usually) wrong to punish people for what they believe, though it may well be appropriate to punish them for acting on their beliefs. There are good pragmatic and moral reasons for not converting thought-sins into thought-crimes.<sup>31</sup>

## Virtue Epistemology

The discussion so far has assumed that normative epistemology is exclusively concerned with evaluating belief (or degrees of belief).<sup>32</sup> In recent decades, several philosophers have challenged this belief-based approach to epistemology. They advocate instead an approach known as *virtue epistemology*. Like veritism, virtue epistemology is a theoretical approach to normative epistemology which is explicitly modeled

on a theoretical approach to ethics. Whereas veritism is modeled on consequentialism, virtue epistemology is modeled on *virtue ethics*.

Since the late 1950s, virtue ethics has challenged act-based approaches to ethics. Virtue ethicists argue that the locus of ethical evaluation should not be the difference between right and wrong action, but the difference between good (i.e., virtuous) and bad (i.e., vicious) people. Appealing to ancient Greek, and particularly Aristotelian, traditions, they construe the principal task of ethics as the identification of moral virtues, where these are understood as dispositions to act in good ways. In an analogous way, virtue epistemologists claim that the central task of normative epistemology is not the evaluation of beliefs, but the evaluation of people and their epistemic virtues (or, to use the Aristotelian language, "intellectual virtues").

Of course, one can evaluate people and their intellectual virtues as well as their beliefs. Alvin Goldman (2002), for example, has offered an account of intellectual virtue which is derived from his veritism. But this certainly does not make him a virtue epistemologist. What distinguishes virtue epistemology from other approaches to normative epistemology is the belief that intellectual virtue is conceptually prior to knowledge and justified belief. According to virtue epistemologists, we can't understand what it is for a belief to be justified or to qualify as a piece of knowledge, until we understand what an intellectually virtuous person would be like. Hence, one prominent virtue epistemologist has defined knowledge as being (roughly) true belief grounded in intellectual virtue (Sosa, 1991, p. 277), and another has defined a justified belief as being (roughly) a belief that an intellectually virtuous person would have (Zagzebski, 1996, p. 236). Whatever one thinks of these issues of conceptual priority, it should be clear that an adequate applied epistemology should have something to say about the nature of intellectual virtue and how to go about cultivating it.

The Greek poet Archilochos famously said "The fox knows many things, but the hedgehog knows one big thing" (Gerber, 1991, Fragment 201). Linda Zagzebski (1996, p. 45) has accused contemporary belief-based epistemology of idealizing the fox at the expense of the hedgehog.<sup>33</sup> There is a great deal of truth in this. The only intellectual virtue recognized in much contemporary epistemology is the virtue of *being well-informed* (i.e., knowing a lot of things). But it should be clear that this virtue is distinct from, and can come at the expense of, a much more important intellectual virtue, *wisdom*, which involves, in part, knowing what is worth knowing. Goldman's veritism seems

a particularly clear example of an approach to epistemology which idealizes the fox at the expense of the hedgehog, since it treats maximizing truth-possession (i.e., being extremely well-informed) as the fundamental epistemic value. Although Goldman acknowledges that not all truths are equally valuable, because not all are equally interesting, he thinks this problem is easily dealt with:

We can no longer suggest that higher degrees of truth-possession are all that count in matters of inquiry. But can't we incorporate the element of interest by a slight revision in our theory? Let us just say that the core epistemic value is a high degree of truth-possession *on topics of interest*. (Goldman, 2002, p. 61)

But this is not the minor modification Goldman seems to think. Whether a topic is of interest or not is often a matter of degree. Hence we are often faced with the problem of simultaneously maximizing two values, interest and truth. How do we compare the intellectual attainments of a fox who knows many moderately interesting things, with the attainments of a hedgehog who knows just one extremely interesting thing? Much of the appeal of veritism was that it offered us the hope of treating normative epistemology as a simple matter of promoting one quantifiable value. This hope appears to be unfounded.

There is another way in which the fact that some topics are more deserving of investigation than others makes epistemology a much messier business than purely belief-based approaches to epistemology suggest. It is not as if there is a fixed set of topics in which we take (or should take) an interest, and our task is just to work out what we should believe about those topics. Preliminary investigation of a topic may reveal that it is, or is not, worth further investigation, or it may reveal some other topic that is more deserving of investigation. Hence questions about what we should believe are inseparable from questions about what we should investigate. Neither kind of question has any absolute priority. Applied epistemology cannot afford to ignore questions about what topics are interesting (or important) if it hopes to provide useful advice to people wondering what to believe. Nor can it ignore the fact that different topics are interesting or important at different times. In this book, I am concerned with what we should believe and how we should pursue knowledge now, at the beginning of the twenty-first century. Like most philosophers, I am fond of the eternal verities, but in this book I will confine my attention to issues of

current interest. There is no guarantee that they will retain that status, but that does not in any way diminish their significance.

For the most part, I will be concerned, even more specifically, with issues of interest in Western societies (particularly English-speaking ones). This is not because I think these societies are more important than others; rather it is because I know them best, and I assume that most of my readers will be similarly placed. As we shall see, some of the issues (e.g., the reliability of the blogosphere versus the conventional media) cannot be addressed without considering the social/political context in which they are raised, and some of the concepts (e.g., that of a *conspiracy theory*) may not be found in other societies. It is the nature of applied philosophy to be heavily dependent on time and place. It is not just that what one says may be interesting in one society, but not another; what one says may actually be true in one society, but not another. In several places, I will argue that certain contemporary epistemologists have an excessively *a priori* approach to some of these issues. Applied epistemology, like applied ethics, needs to engage closely with contemporary social and political realities.

## Applied Epistemology and Social Policy

Let's take the analogy between applied ethics and applied epistemology a step further. Applied ethics consists of two parts, the part concerned with questions about individuals and their behavior (e.g., How much of our income, if any, should we donate to the poor?), and the part concerned with questions of social policy (e.g., Should the tax system be used to redistribute wealth?). We can make a similar division within applied epistemology. This book is principally concerned with questions about individuals, what they should believe and how they should pursue knowledge and/or wisdom. But epistemology, like ethics, can also be applied to questions of social policy.<sup>34</sup> Such applications are an important part of the branch of epistemology known as "social epistemology."

Alvin Goldman has applied his veritism to social epistemology, calling the result "veritistic social epistemology" (from now on VSE). VSE aims to evaluate social policies in terms of their "knowledge impact" (Goldman, 1999b, p. 6). VSE, like veritism more generally, is modeled on consequentialist, and more particularly utilitarian, moral theory (Goldman, 1999b, p. 87), with knowledge in the thin sense (i.e., true belief)

**Table 1.1**

	<i>t1</i>	<i>t2</i>
S1	DB( <i>p</i> ) = 0.4	DB( <i>p</i> ) = 0.7
S2	DB( <i>p</i> ) = 0.7	DB( <i>p</i> ) = 0.9
S3	DB( <i>p</i> ) = 0.9	DB( <i>p</i> ) = 0.6
S4	DB( <i>p</i> ) = 0.2	DB( <i>p</i> ) = 0.8

playing the role in the former that happiness or utility plays in the latter. Some people reject utilitarianism as a general guide to life, but consider it to be defensible as a guide to social policy (e.g., Goodin, 1995). For similar reasons, one may reject veritism as a general theory of normative epistemology (perhaps for some of the reasons I discussed earlier), but accept it as a guide to social policy (i.e., accept VSE).

Before considering the plausibility of this position, we should get a clearer picture of the way VSE is suppose to work. Goldman (1999b, p. 93) invites us to imagine that there is a community of four agents S1–S4, and a certain true proposition, *p*, which is of interest to them. At time *t1* the agents have degrees of belief (DB) in *p* (identified with subjective probabilities) as shown in Table 1.1. A certain social policy  $\pi$  is enacted, with the result that at a later time *t2* the agents have new degrees of belief in *p*, shown in the next column of the table. In this situation, social policy  $\pi$  has positive veritistic value (V-value), because it increases the mean degree of belief in the correct answer to the question of interest (from 0.55 to 0.75).

Not surprisingly, VSE is vulnerable to criticisms which are analogous to common criticisms of utilitarian approaches to social policy. Just as utilitarian approaches to social policy have been criticized, on the grounds that they ignore justice (or more specifically justice in the distribution of happiness or utility), VSE may be criticized for ignoring a distinctively epistemic form of justice (i.e., justice in the distribution of knowledge, understood as true belief, or interesting or important true belief).<sup>35</sup>

What would a just distribution of knowledge be like? When answering this question, it is natural to consider well-known principles of distributive justice for other kinds of goods. Someone inspired by John Rawls's Difference Principle (Rawls, 1971, pp. 75–83), for example, might say that the most epistemically just social policy would be one which maximized the knowledge of the least knowledgeable members

**Table 1.2**

	$\pi$	$\psi$
S1	DB( $p$ )=0.5	DB( $p$ )=0.7
S2	DB( $p$ )=0.7	DB( $p$ )=0.9
S3	DB( $p$ )=0.4	DB( $p$ )=0.2
S4	DB( $p$ )=0.6	DB( $p$ )=0.8

of the community. So, for example, when comparing the veritistic merits of two practices,  $\pi$  and  $\psi$ , in a community of four people, we might decide that their outcomes, measured in degree of confidence in some true proposition of interest  $p$ , will be as indicated in Table 1.2. In this situation, although  $\psi$  is to be preferred to  $\pi$  by the standards of VSE, because it has greater V-value (the mean degree of true belief being 0.65, whereas the mean degree of true belief under  $p$  is 0.55),  $\pi$  is to be preferred to  $\psi$  by the standards of what we might call the Epistemic Difference Principle, because the least knowledgeable member of this community would be better informed as a result of  $\pi$  than as a result of  $\psi$ . Hence  $\pi$ , according to that principle, is more epistemically just than  $\psi$ .

Now, I am not advocating the Epistemic Difference Principle. It is open to some of the same objections that have been leveled against the original Difference Principle on which it is modeled (and perhaps others). For example, someone with libertarian sympathies, like Robert Nozick, might object to both the Epistemic Difference Principle and VSE, on the grounds that they seek to impose a pattern (Nozick, 1974, pp. 208–10) in the distribution of knowledge, and so ignore historical questions about where the knowledge in question came from and how it came to be distributed in the way it is. In other words, they both ignore questions about who had intellectual property in the knowledge in the first place, and hence they have the potential to violate the rights of holders of intellectual property who should be able to transfer their knowledge (or not to transfer it) to whomever they choose.

I will not attempt to adjudicate these disputes here. My intention has just been to argue that VSE neglects questions about justice in the distribution of knowledge, and that this constitutes an objection to VSE. Goldman could try defending VSE against this objection by pointing out that VSE was not meant to provide us with an all-things-considered evaluation of social policies. Goldman acknowledges that veritistic considerations may be outweighed by other considerations, including

considerations of justice. At one point, for example, he explicitly states that concerns for procedural justice might legitimately lead a court of law to exclude evidence which would be admissible on purely veritistic grounds (Goldman, 1999b, p. 284).

This response would be right, as far as it goes, but it would miss the point of the current objection. For the kind of justice under consideration is distinctively epistemic (indeed it is specifically veritistic). To be clear, my objection is not that general considerations of justice can rightly trump considerations of maximizing true interesting belief (though this is of course true); rather my objection is that considerations about the just allocation of such belief can rightly trump considerations about maximizing it.

Am I doing epistemology or ethics here? Although I have been talking up to this point as though epistemology and ethics are distinct, though analogous, enterprises, strictly speaking epistemology (or at least normative epistemology) is not *merely* analogous to ethics, it is a part of ethics. Belief-based epistemology is a part of ethics, because believing, as one of the things we do (albeit not always voluntarily), is a kind of acting (and not merely comparable to acting as I earlier supposed). Likewise virtue epistemology is a branch of ethics, because intellectual virtues, such as wisdom, are among the virtues required of a good life. Epistemology is the branch of ethics specifically concerned with intellectual goods, one of which is knowledge.<sup>36</sup> Hence Goldman's VSE should be understood as an ethical theory, and it is perfectly appropriate to criticize it on ethical grounds.

Suppose, by way of comparison, that an economist examined various public policies for their impact on wealth, and provisionally endorsed policies which would maximize overall (or average) wealth. We may suppose that the economist's endorsement is only provisional because he or she accepts that noneconomic values may sometimes trump the maximizing project. Nonetheless, the economist's position would certainly (and I think rightly) be criticized if it fails to address questions about the justice or injustice of wealth distributions that would result from the social policies he or she has provisionally endorsed. This criticism would not be external to the economist's project, nor should it be dismissed on the grounds that it concerns ethics *rather than* economics. Applied economics, like applied epistemology, is a branch of applied ethics.

Although epistemology is a branch of ethics, the distinction between epistemology and ethics is so well entrenched in philosophical debate, and often so useful, that it will sometimes be convenient to talk as if

they are distinct (i.e., nonoverlapping) subjects. This practice is harmless, so long as we bear in mind that when we do so we are construing ethics narrowly as a subject concerned with “outward” actions *rather than* beliefs, or with moral virtues *rather than* intellectual virtues.

## Notes

- 1 Toward the end of this chapter, I will argue for a stronger position, namely that the problem of working out what we ought to believe can be treated as part of the problem of working out how we ought to live, and that epistemology (or at least normative epistemology) can not only be compared to ethics, it can be treated as a branch of ethics. Nonetheless, it will usually be convenient to treat ethics and epistemology as separate disciplines.
- 2 See, for example, Edel, Flower, and O'Connor (1994, pp. 1–8) who point out that Aristotle, Aquinas, Hobbes, Bentham, and Kant all did what would now be called “applied ethics.”
- 3 Of course it's still an interesting question, at least for many people. However, the secularization of society since Hume's time means that it is not interesting for as many people as it would once have been.
- 4 There are some exceptions. As we shall see in Chapter 3, there is a substantial contemporary literature on the epistemology of democracy, and, as we shall see in Chapter 5, some of the current literature on the epistemology of conspiracy theories is explicitly presented as political philosophy.
- 5 In this way, my approach is comparable to Peter Singer's approach to the ethics of our treatment of nonhuman animals. Although Singer himself is a utilitarian, he tries to construct arguments which will persuade a wide range of people, including people who don't share his theoretical commitment to utilitarianism.
- 6 The topic is discussed in both the *Theaetetus* and the *Meno*.
- 7 It has usually been assumed that justification is another necessary condition for knowledge. Before Edmund Gettier (1963) it was usually assumed that truth and justification were both individually necessary and jointly sufficient.
- 8 The term comes from Nozick (1974, p. 29).
- 9 The identification of suspension of judgment with having a (roughly) 0.5 degree of belief is questionable (see Coady, 2010, p. 107). But, whatever you think about this, the real flaw with Goldman's argument is its identification of this state (or these states) with having no opinion.
- 10 I shall shortly challenge the idea that it is the preeminent intellectual attainment.
- 11 Caracas *is* the capital of Venezuela.



- 12 Such critiques of consequentialist ethics are usually called “deontological.” I will avoid that word here, however, since it is often used in epistemology to refer to a view about the nature of justification (i.e., that it is to be understood in terms of a duty or obligation), rather than a view about the value of justification (i.e., that it is intrinsically valuable). I take the term “proceduralism” from Goldman (1999b, pp. 75–9).
- 13 Epistemic justification is often contrasted with pragmatic (ethical or prudential) justification. You might be epistemically justified in believing something (i.e., the available evidence and arguments indicate that it is true); nonetheless you might (at least arguably) not be all-things-considered justified in believing it, because it is in not in your interests to believe it or because you have a moral obligation not to believe it. I will discuss this issue further in the next section.
- 14 This includes many who think that their intuitions are wrong and consequentialist ethics is right.
- 15 Strictly speaking he doesn't claim that belief is definitely in our interests, just that it has greater expected utility than disbelief.
- 16 This is certainly not the only reason. The word “epistemology” is unfamiliar and unattractive to many people, and, as a result, few people outside of professional epistemology have taken much interest in the subject.
- 17 Doxastic voluntarism is not popular. I don't know anyone else who has explicitly embraced it.
- 18 Indeed, one prominent contemporary philosopher has declared “the fundamental problem of epistemology to be that of *deciding* what to believe” (Pollock, 1986, p. 10, my emphasis).
- 19 It tends to adopt a passive view of perception and desire, as well as belief.
- 20 Similar views have been expressed by Robert Audi, who says that “belief is more like a response to external grounds than a result of internal volitive thrust” (Audi, 2001, p. 98), and John Heil, according to whom “believers are largely at the mercy of their belief-forming equipment” (Heil, 1983, p. 357).
- 21 Alston (1989 ) describes this kind of control as “weak” (p. 118). Adler just stipulates that when he speaks of controlling beliefs he is “speaking of *direct* control” (Adler, 2002, p. 57).
- 22 Of course one needn't *consciously* send signals to one's muscles in order to raise one's hand, but then one doesn't necessarily consciously control one's evidence-gathering activities either.
- 23 Not only are there things that most of us could not bring ourselves to believe, there are things that most of us cannot help believing. Again, there is no disanalogy between beliefs and actions here. Some actions are voluntary; some are not.
- 24 This is more than a facetious joke. It is a reminder of my earlier point that ordinary forms of thought and language presuppose that we have extensive control over what we believe.

## *Introduction*

- 25 Alvin Goldman, for example, critically discusses the “guidance-deontological conception of justification” (1999a, pp. 271–4).
- 26 Likewise to say that someone should do something is not necessarily to say that they are duty-bound to do it.
- 27 I do not deny that one may have duties to oneself, as Kant, for example, thought. Nonetheless, the paradigm cases of duty do seem to involve other people.
- 28 The original source for this quote is Lifton (1986, p. 16).
- 29 Intuitively it seems the more certain one is of something one shouldn't believe, the more blameworthy one is.
- 30 Moral certainty is not the only doxastic attitude Levy thinks is incompatible with responsibility for belief. He claims that Klein would be responsible for his belief that Jews are a gangrenous appendix in the body of mankind “only if he is less than morally certain that this is the case, if he thinks it matters greatly whether he is right, if he believes that further evidence is available for the proposition and that gathering the evidence is worth the trouble” (2007a, p. 147). Unless all these conditions are met, Klein cannot be expected to investigate further, and so is not responsible. We have seen that Levy just assumes that Klein has no control over the degree of confidence with which he holds his belief. He also seems to just assume that Klein's attitude to the importance of what he believes is not under his control, and that his beliefs about the availability of further evidence and the desirability of acquiring it are not under his control. But once we have accepted (as Levy has) that we have at least some control over our beliefs and over our attitudes towards our beliefs, none of this can be assumed.
- 31 I thank an anonymous reviewer for suggesting this way of putting the point.
- 32 Our concern with knowledge has also been a concern with belief, since we have treated knowledge as a kind of belief, that is, true belief, or justified true belief, or justified true belief with some further condition or conditions.
- 33 Zagzebski is discussing Isaiah Berlin's (1978) interpretation of the Archilochos fragment, rather than the fragment itself. Berlin's discussion of the fragment is largely responsible for its current fame.
- 34 Questions of social policy are not central to this book, but they do come up in several places. The Postscript is entirely devoted to a social policy issue.
- 35 Goldman does discuss the distribution of veritistic value, and acknowledges that maximizing this value may not always be desirable (see 1999b, p. 96). His discussion in this passage is not, however, about the justice or otherwise of various distributions, but about their efficiency, given a certain goal, that of promoting “the community's interest.” For example, considerations of efficiency may mean that some information should only be distributed on a “need to know” basis.
- 36 As we have seen, justified belief and wisdom are also epistemic goods.