

7. The Epistemic Significance of Disagreement

Thomas Kelly

Looking back on it, it seems almost incredible that so many equally educated, equally sincere compatriots and contemporaries, all drawing from the same limited stock of evidence, should have reached so many totally different conclusions—and always with complete certainty.

(John Michell, *Who Wrote Shakespeare?*)

1. INTRODUCTION

Consider the following issues, each of which is the object of considerable controversy:

- (1) the extent to which a desire to intimidate the Soviet Union played a role in Harry Truman's decision to drop the atomic bomb on Japan in 1945
- (2) whether Truman's decision to do so was morally justified
- (3) whether there are in fact any truths of the kind that Immanuel Kant called "synthetic a priori"

I have a belief about each of these issues, a belief that I hold with some degree of conviction. Moreover, I ordinarily take my beliefs about each of these matters to be rational—I think of myself as having good reasons for holding them, if pressed to defend my position I would

For helpful discussion and correspondence, I am grateful to Adam Elga, David Chalmers, David Christensen, Richard Feldman, Anil Gupta, Peter van Inwagen, Derek Parfit, Jim Pryor, Pamela Hieronymi, Michael Rescorla, Kerian Setiya, Jonathan Vogel, Ralph Wedgwood, and Roger White. Earlier versions of this paper were read at the University of Notre Dame, Harvard University, the University of Pittsburgh and at a Pacific Division meeting of the American Philosophical Association; I am grateful to the audiences present on those occasions.

cite those reasons, and so on.¹ On the other hand, I am very much aware of the fact that, with respect to each issue, there are many others who not only do not share my belief, but in fact, take a diametrically opposed position. Of course, the mere fact of disagreement need not be problematic: if, for example, I was convinced that all of those who disagreed with me were simply being foolish, or hadn't bothered to think about the matter carefully enough, or were unfamiliar with evidence that I happen to possess (evidence which, if presented to them, would result in a change in their views), then I might simply shrug off this disagreement. But in fact, I believe no such thing: I acknowledge that on many controversial issues with respect to which I have a firmly held belief, there are some who disagree with me whose judgement cannot be simply discounted by appeal to considerations of intelligence, thoughtfulness, or ignorance of the relevant data.

Can one rationally hold a belief while knowing that that belief is not shared (and indeed, is explicitly rejected) by individuals over whom one possesses no discernible epistemic advantage? If so, what assumptions must one be making about oneself and about those with whom one disagrees? In deciding what to believe about some question, how (if at all) should one take into account the considered views of one's **epistemic peers**?²

My aim in this paper is to explore the *epistemic significance of disagreement*. A central concern is whether the practice of retaining beliefs that are rejected by individuals over whom one claims no epistemic advantage is a defensible one. It is, of course, far from clear that the relevant practice *is* defensible. For it is natural to suppose that

¹ Of course, not all rational beliefs are rationalized by supporting reasons: my belief that $2 + 2 = 4$ is (I assume) a rational belief, but it is not rationalized in virtue of standing in a certain relation to supporting considerations, in the way that my rational belief that *communist economies tend to be inefficient* is. In this paper, however, I will ignore the case of beliefs whose rationality consists in their status as 'properly basic' (to borrow a phrase from Alvin Plantinga). Indeed, I suspect that beliefs of this kind would require a very different treatment than the one offered here.

² I owe the term 'epistemic peer' to Gutting (1982). Gutting uses the term to refer to those who are alike with respect to 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (p. 83). I will use the term in a somewhat extended sense. As I will use the term, the class of epistemic peers with respect to a given question are equals, not only with respect to their possession of the sort of general epistemic virtues enumerated by Gutting, but also with respect to their exposure to evidence and arguments which bear on the question at issue. I discuss this notion further in § 2.3 below.

persistent disagreement among epistemic peers should undermine the confidence of each of the parties in his or her own view. This natural intuition was voiced by Henry Sidgwick in a memorable passage in *The Methods of Ethics*:

the denial by another of a proposition that I have affirmed has a tendency to impair my confidence in its validity . . . And it will be easily seen that the absence of such disagreement must remain an indispensable negative condition of the certainty of our beliefs. For if I find any of my judgements, intuitive or inferential, in direct conflict with a judgement of some other mind, there must be error somewhere: and if I have no more reason to suspect error in the other mind than in my own, reflective comparison between the two judgements necessarily reduces me . . . to a state of neutrality. (p. 342)

Sidgwick's idea—that reflection on the relevant sort of disagreement should reduce one to 'a state of neutrality'—has been endorsed by thinkers both early and late. The idea played a prominent role in ancient skepticism as one of the 'modes of Pyrrhonism' designed to rationally induce suspension of judgement. Here is the characterization offered by Sextus Empiricus:

According to the mode deriving from dispute, we find that undecidable dissension about the matter proposed has come about both in ordinary life and among the philosophers. Because of this we are not able either to choose or to rule out anything, and we are driven to suspend judgement. (I. 165)³

Indeed, in his own presentation of the case for skepticism, Sextus seems to indicate that the existence of such disagreement is ultimately the most fundamental consideration of all.⁴ The same idea is a recurrent theme in Montaigne's case for skepticism as presented in his *Essays*. More recently, Keith Lehrer (1976) has claimed that there is simply no room for rational disagreement among those who share the same information and have even a minimal level of respect for each other as judges: in such circumstances, each party to the dispute should revise his or her own judgement until consensus is achieved. In economics, a substantial body of literature similarly seems to suggest that the

³ As reported by Sextus, the argument from disagreement was one of the Ten Modes of Aenesidemus as well of one of the Five Modes of Agrippa; it thus played a part in both early and late Pyrrhonism.

⁴ See I. 178–9, where the standard modes seem to ultimately depend upon the existence of 'interminable controversy among the philosophers'.

uniquely rational response to known disagreement is to revise one's original beliefs so as to bring about consensus.⁵

Despite its attractiveness, this line of thought is, I believe, mistaken. Disagreement does not provide a good reason for skepticism or to change one's original view. In what follows, I will argue for the following thesis: once I have thoroughly scrutinized the available evidence and arguments that bear on some question, the mere fact that an epistemic peer strongly disagrees with me about how that question should be answered does not itself tend to undermine the rationality of my continuing to believe as I do. Even if I confidently retain my original view in the face of such disagreement, my doing so need not constitute a failure of rationality. Indeed, confidently retaining my original belief might very well be the *uniquely* reasonable response in such circumstances.

According to the view that I will defend then, disagreement does not have the kind of significance that has often been claimed for it. However, it would be a mistake to conclude that disagreement is therefore without epistemic significance. I will thus also attempt to clarify the nature of the significance that disagreement does have in those cases in which it *is* of significance.

The discussion which follows is as much exploratory as it is polemical. A primary concern is to make fully explicit the substantive commitments and assumptions about rationality of one who defends the views that I defend. I do not pretend that the relevant commitments are costless. I myself do not find the costs unacceptably high. But this, of course, is itself something about which others might very well disagree.

2. SOME PRELIMINARY DISTINCTIONS

2.1.

I begin by locating the question that I want to pursue relative to certain other, closely related questions. Here, the most straightforward distinction to be drawn is that between descriptive and normative questions. There is a considerable amount of empirical evidence which suggests that an awareness of disagreement tends to lead us to significantly

⁵ Here I have in mind the tradition of research which dates from Aumann's seminal paper 'Agreeing to Disagree' (1976). I discuss the import of this literature in § 3 below.

moderate our opinions. That is, within isolated groups, there are strong psychological pressures that tend to lead to the formation of consensus, or at least, to the formation of a dissensus that is less polarized than the one which would otherwise have obtained. Questions about the pervasiveness and scope of such phenomena have been fruitfully explored by social psychologists.⁶

In contrast to descriptive questions about how an awareness of disagreement in fact affects our beliefs, the question that I want to pursue belongs to the class of normative questions—questions about how an awareness of disagreement *should* affect our beliefs. Answering these normative questions could, in principle, lead us to revise our actual practice, to alter our characteristic responses to disagreement. Alternatively, if it is beyond our power to revise our actual practice—say, because our actual responses to disagreement are psychologically fixed⁷—how we answer these normative questions might affect our attitudes towards our unalterable responses. Thus, suppose that, as a matter of fact, an awareness of disagreement tends to more or less inevitably lead us to revise our views in the direction of greater consensus. If we conclude that it is epistemically appropriate to give a great deal of weight to the judgements of others in revising our own beliefs, then we might view this unavoidable psychological tendency with relative equanimity, or even with pride, as symptomatic of our natural and reflexive rationality. If, on the other hand, we conclude that doing so is not the epistemically appropriate response, then we might view our inevitable tendency to respond in this way in a less favorable light: perhaps as symptomatic of a somewhat craven desire to adhere to orthodoxy for orthodoxy's sake.

2.2.

As I have emphasized, it is at least somewhat natural to suppose that when one discovers that others explicitly reject some view that one holds, this discovery ought to make one more skeptical of that view. It is important, however, to distinguish carefully between two quite

⁶ The classic studies in this tradition were conducted by Solomon Asch (1952, 1956).

⁷ A prominent theme in recent epistemology is that much of the epistemological tradition seems to presuppose that we possess a degree of control over our beliefs that we do not in fact possess. See e.g. Alston (1988) and Plantinga (1993: esp. ch. 1).

distinct kinds of skepticism that such a discovery might be thought to warrant. The first kind of skepticism is skepticism about whether there is, after all, a fact of the matter about the disputed question. That is, it might be thought that persistent disagreement with respect to a given domain warrants some kind of non-factualism or error theory about that domain. Thus, in moral philosophy the existence of disagreement with respect to fundamental ethical questions is often claimed to strengthen the case for non-factualism or some variety of error theory on the grounds that there being no fact of the matter is the best explanation of our inability to reach agreement.⁸ Similarly, the phenomenon of persistent disagreement among theorists concerning the correct solution to various decision problems is sometimes thought to bolster the case for expressivist accounts of discourse about practical rationality. Although in contemporary philosophy this move is most often made with respect to normative domains, it has in the past often been made with respect to non-normative domains as well. Thus, the logical positivists frequently insisted that the seemingly interminable controversies among theologians and metaphysicians are due to the fact that the relevant bodies of discourse are not truth-apt but rather ‘cognitively meaningless’. Here again, the driving idea is that the best explanation of why we cannot agree about what the facts are is simply that there are no facts upon which we might agree.

Questions about the circumstances in which disagreement warrants some variety of non-factualism or error theory about a given domain are interesting ones, but they will not be pursued here. Instead, I want to examine cases in which we are confident that there *is* a genuine fact of the matter—*despite* the existence of disagreement—in order to inquire as to how an awareness of that disagreement should affect our beliefs in such cases. I assume that there are some domains with respect to which we occupy this position. Consider, for example, history. There is, I assume, a fact of the matter about whether a desire to intimidate the Soviet Union played a role in Harry Truman’s decision to drop the atomic bomb—however much knowledgeable and highly qualified historians might disagree about what that fact of the matter is. Of course, certain postmodernists and anti-realists about the past might question this. But here it is fair to say, I think, that our commitment to a robust

⁸ Mackie (1977) is a classic attempt to motivate an error theory by appeal to facts about ethical disagreement.

factualism about historical discourse is stronger than any argument that such thinkers have yet provided.

Compared to questions about whether disagreement should undermine our commitment to factualism about various domains, questions about the extent to which disagreement poses a distinctively *epistemic* challenge have been relatively underexplored. In fact, much of what little discussion this question has received has taken place within the philosophy of religion: philosophers of religion have debated the extent to which an awareness of the great diversity of (sometimes) incompatible religious traditions ought to make a theist more skeptical about the distinctive claims of her own tradition.⁹ It is unclear, however, whether there is any *special* problem about religious claims in particular. For, as Peter van Inwagen (1996) has emphasized, everyone, or almost everyone, would seem to be in the position of the theist with respect to at least some questions. That is, virtually everyone has at least some beliefs that are explicitly rejected by individuals over whom he or she possesses no discernible epistemic advantage. This phenomenon, while no doubt familiar enough from everyday life, is perhaps especially salient for philosophers. For philosophy is notable for the extent to which disagreements with respect to even the most basic questions persist among its most able practitioners, despite the fact that the arguments thought relevant to the disputed questions are typically well-known to all parties to the dispute. (It is not, after all, as though Compatibilists about free will think themselves privy to some secret master argument, such that if this argument were presented to the Incompatibilists, the Incompatibilists would see fit to abandon their view.)

2.3.

It is uncontroversial that there are some circumstances in which one should give considerable weight to the judgements of another party in deciding what to believe about a given question. Paradigmatic examples consist of cases in which it is clear that the other party enjoys some epistemic advantage with respect to the question at issue. The list of possible advantages which one party might enjoy over another seems to divide naturally into two general classes. First, there are advantages that

⁹ See, for example, Gutting (1982), Plantinga (2000), and the essays collected in Quinn and Meeker (2000).

involve a superior familiarity with or exposure to evidence and arguments that bear on the question at issue. Thus, suppose that I know that you possess not only all of the evidence which I possess but also some relevant evidence which I lack. (That is, my total evidence is a proper subset of your total evidence.) In these circumstances, it makes sense for me to treat your beliefs as indicators of the actual state of the evidence since I have no independent access to the character of that evidence. More subtly: it might be that although we have both been exposed to the same body of evidence, you have carefully scrutinized that evidence while I have considered it only hastily or in a cursory manner. Here again, it is your superior familiarity with the evidence which makes a certain measure of deference on my part the appropriate course.

A second class of epistemic advantages which one might potentially enjoy consists in superiority with respect to general epistemic virtues such as intelligence, thoughtfulness, freedom from bias, and so on. Thus, if I know that I have great difficulty being objective when it comes to assessing the quality of my work but that you labor under no such handicap, then I have a reason to defer to your judgements about my work, all else being equal.¹⁰

Any plausible view, I take it, will allow for the fact that I should give considerable weight to your judgements when I have reason to believe that your epistemic position is superior to my own in either of these ways (at least, provided that I do not claim some compensating advantage). Because some measure of deference seems clearly appropriate in such circumstances, the question that I want to pursue concerns the normative significance of disagreement in cases in which *neither* of the parties enjoys such an advantage.

Let us say that two individuals are **epistemic peers** with respect to some question if and only if they satisfy the following two conditions:

¹⁰ As this last example makes clear, it is no doubt overly simple to attribute to an individual some particular level of (e.g.) objectivity or thoughtfulness irrespective of a particular subject matter: the extent to which an individual possesses such qualities might very well (and in the usual case, will) vary significantly from domain to domain. Attributions of a given level of objectivity or thoughtfulness should thus be relativized to particular domains. (It is an empirical question, I take it, how the relevant domains should be demarcated.) In what follows, the need for such relativization should be taken as understood; for expository purposes, I will avoid repeated mentions of this need, and write simply of an individual's objectivity (etc.) rather than her objectivity-with-respect-to-domain-A.

- (i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and
- (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias.¹¹

The question at issue, then, is whether known disagreement with those who are one's epistemic peers in this sense must inevitably undermine the rationality of maintaining one's own views.

3. NO AGREEING TO DISAGREE?

Why might one think that it is unreasonable to steadfastly maintain one's views in the face of such disagreement? In economics, there is a substantial body of literature which purports to show the irrationality of 'agreeing to disagree' in various circumstances. The first to develop general results along these lines was Robert Aumann (1976). In a classic paper, Aumann showed that if two or more individuals (i) update their beliefs by Bayesian conditionalization, (ii) have common prior probabilities, and (iii) have common knowledge of each other's opinions, then (iv) those individuals will not knowingly disagree on the answer to any question: rather, they will continuously revise their beliefs until consensus is reached. Subsequent work has shown that Aumann's 'no

¹¹ It is a familiar fact that, outside of a purely mathematical context, the standards which must be met in order for two things to count as *equal* along some dimension are highly context-sensitive. Thus, inasmuch as classes of epistemic peers with respect to a given issue consist of individuals who are 'epistemic equals' with respect to that issue, whether two individuals count as epistemic peers will depend on how liberal the standards for epistemic peerhood are within a given context. That is, whether two individuals count as epistemic peers will depend on *how much* of a difference something must be in order to count as a *genuine* difference, according to the operative standards. In the same way, whether two individuals count as 'the same height' will depend on the precision of the standards of measurement that are in play. (Lewis 1979 is a classic discussion of the relevant kind of context-sensitivity.) Of course, given sufficiently demanding standards for epistemic peerhood, it might be that no two individuals *ever* qualify as epistemic peers with respect to any question. (Perhaps there is always at least some *slight* difference in intelligence, or thoughtfulness, or familiarity with a relevant argument.) Similarly, it might be that no two individuals count as the same height given sufficiently demanding standards of equality. My sense is that, often enough, the standards that we employ in assessing intelligence or thoughtfulness (like the standards that we employ when measuring height) are sufficiently liberal to allow individuals to qualify as equal along the relevant dimensions.

agreeing to disagree' result survives various weakenings of his original assumptions.¹²

Contrary to what one might naturally assume, however, this tradition of research does not in fact support the conclusion that known disagreement among epistemic peers provides each of the peers with a good reason to revise his or her view. Indeed, close examination reveals that the technical results which have been established thus far do not bear on the case of disagreement among epistemic peers at all. As noted, Aumann's original proof depends on the assumption of common prior probabilities. This assumption is tantamount to assuming that there is a prior agreement as to the normative import of any piece of evidence which might be encountered. In effect, Aumann's 'no agreeing to disagree' result holds only for individuals *who would hold identical views given the same evidence*. And although subsequent work in this tradition has shown that Aumann's result can survive certain weakenings in his original assumptions, the assumption of common prior probabilities has not proven dispensable. Now, by definition, individuals who are epistemic peers with respect to a given question have been exposed to the same evidence which bears on that question. Disagreement among epistemic peers then, is disagreement among those who disagree *despite* having been exposed to the same evidence. Thus, our question concerns a case which stands outside the range of cases for which Aumann's result holds.

The guiding idea behind the 'no agreeing to disagree' literature is that, in many circumstances, the discovery that another person holds a view that one is inclined to reject constitutes evidence that the other person has access to relevant evidence which one does not possess. By giving some weight to the view of the other person, one is able to take into account the import of that evidence to which one would otherwise lack access. Thus, one does not have to possess the evidence for oneself in order to take its epistemic import into account. This guiding idea represents a genuine insight.¹³ Indeed, as emphasized above (§ 2.3), *any* plausible epistemological view will allow for the fact that I should give considerable weight to your beliefs when I have reason to think

¹² Geanakoplos (1994) provides a basic exposition of Aumann-like results through 1994.

¹³ For further exploration of this theme, as well as an attempt to specify normative principles which should guide our attempts to take account of evidence that we do not possess, see Kelly (forthcoming *b*).

that the fact that you believe as you do is attributable to your possession of some relevant evidence which I do not possess. However, our present question is not how one should respond to the beliefs of others when one lacks access to the evidence on which those beliefs are based. The question, rather, is how one should respond when one *does* have access to the relevant evidence.

The technical results of the 'no agreeing to disagree' literature then, do not bear directly on the question at issue. It might be thought, however, that one who appreciates the guiding idea which underlies these technical results (viz. that one takes into account evidence which one does not possess by taking into account the views of those with whom one disagrees) will naturally embrace the view that disagreement provides a good reason for skepticism. For suppose that I know that you are significantly better informed than I am with respect to some question. In these circumstances, it makes sense for me to defer to your better informed judgement in deciding what to believe about that question. In reasoning in this way, I presuppose that you are a competent evaluator of that evidence which is available to you but not to me. (If I knew that you were *not* a competent evaluator of this evidence, then it would be illegitimate for me to draw inferences about the character of your evidence from the content of your beliefs.) Suppose that at some later time, our epistemic positions are equalized: I gain access to that evidence which was previously available only to you. I am now in a position to make my own judgement about the probative force of the evidence. Still, it might be thought that consistency requires that I continue to give considerable weight to your judgement about what our (now common) total evidence supports. After all, even if I'm strongly inclined to disagree with you as to the overall import of the evidence in a given case, shouldn't I give considerable weight to your judgement given my readiness to defer to you when I am otherwise ignorant of that evidence? Indeed, unless I have some positive reason to think that one of us is more likely to do a better job with respect to assessing the relevant evidence than the other, shouldn't I give *equal* weight to our considered judgements? Recall Sidgwick's remark: 'if I find any of my judgements, intuitive or inferential, in direct conflict with a judgement of some other mind, there must be error somewhere: and if I have no more reason to suspect error in the other mind than in my own, reflective comparison between the two judgements necessarily reduces me . . . to a state of neutrality'.

Ultimately, the argument that it is unreasonable to maintain one's views in the face of such disagreement depends on considerations of symmetry. According to this line of thought, the only thing that would justify one in maintaining views that are rejected by one's epistemic peers would be if one had some positive reason to privilege one's own views over the views of those with whom one disagrees. But *ex hypothesi*, no such reason is available in such cases. In the next section, I take up this argument and attempt to show how it might be resisted.

4. THE APPEAL TO SYMMETRY

Suppose that two epistemic peers—let's call them 'you' and 'I'—are each deliberating about what attitude to take towards a given hypothesis H in the light of the available evidence. Suppose further that, as a result of my assessment of the evidence, I come to believe H, while as a result of your assessment of the evidence, you come to believe not-H. If we subsequently become aware of our disagreement, how if at all should we revise our respective views? Again, in these circumstances, considerations of symmetry would seem to dictate that suspension of judgement is the uniquely reasonable response on both of our parts: that is, each of us should abandon his or her prior conviction and retreat towards an attitude of agnosticism with respect to H. For how could either of us defend doing otherwise? Consider how the situation appears from my perspective. *Ex hypothesi*, I admit that there are no objective criteria that make it antecedently more probable that I am more likely than you are to be correct on this particular occasion—I do not claim to be any smarter, a better reasoner, or to possess some relevant evidence which you lack. Given the acknowledged, perfect symmetry of our positions, how can I possibly justify *not* giving equal weight to your considered judgement? After all, wouldn't this be the most rational course for some objective, on-looking third party who knew nothing about our dispute other than the fact that it is two judges of equal competence and qualifications who disagree? Given this, wouldn't my failure to give equal weight to your judgement amount to a kind of epistemic arbitrariness on my part, an indefensible privileging of my own position for no other reason than the fact that it is my own?

However, the claim that things are perfectly symmetrical in such cases deserves further scrutiny. Indeed, to uncritically assume that

things are perfectly symmetrical with respect to all of the epistemically relevant considerations in such cases is, I think, to subtly beg the question in favor of the skeptical view. For consider: I am no smarter than you are, no better at reasoning, no better informed, and (hence) no more fit to pronounce upon the issue at hand. So far, it is uncontroversial that things are perfectly symmetrical between us. Then a body of evidence is introduced, and we are asked to make a judgement about how strongly that body of evidence confirms or disconfirms a certain hypothesis. Suppose that, as it turns out, you and I disagree. From my perspective, of course, this means that you have misjudged the probative force of the evidence. The question then is this: why shouldn't I take *this* difference between us as a relevant difference, one which effectively breaks the otherwise perfect symmetry?

After all, the question of how well someone has evaluated the evidence with respect to a given question is certainly the kind of consideration that is relevant to deciding whether his or her judgement ought to be credited with respect to that question. That is, it is exactly the sort of consideration that is capable of producing the kind of asymmetry that would justify privileging one of the two parties to the dispute over the other party. And from my vantage point—as one of the parties *within* the dispute, as opposed to some on-looking third party—it is just this undeniably relevant difference that divides us on this particular occasion.¹⁴

One might wonder: is my assessment that you have misjudged the probative force of the evidence consistent with my continuing to regard you as a genuine epistemic peer? Yes, it is. Of course, if I came to believe that I am, in general, a better evaluator of evidence than you are, then this would be a good reason for me to demote you from the ranks of those to whom I accord the status of epistemic peer. But a revision in my assessment of our relative levels of competence is in no way mandated by the judgement that one of us has proven superior with respect to the exercise of that competence on a given occasion. Two chess players of equal skill do not always play to a draw; sometimes one or the other wins, perhaps even decisively.

¹⁴ In cautioning against the tendency to think that the correct way to view such disputes is from a purely external, third-person point of view, I echo Richard Foley. As he puts the point: 'It is deeply misleading to think about such conflicts in terms of a model of neutral arbitration between conflicting parties' (2001: 79). Cf. Foley (1994: 65–6) and Plantinga (1995: 182).

At the outset of the paper, I asked what I must be assuming about myself and about others who have been exposed to the same evidence when I continue to hold a belief that they reject. My answer to this question is: perhaps not very much. In particular, I need not assume that I was better qualified to pass judgement on the question than they were, or that they are likely to make similar mistakes in the future, or even more likely to make such mistakes than I am. All I need to assume is that *on this particular occasion* I have done a better job with respect to weighing the evidence and competing considerations than they have.

Of course, there is still the question of whether I am *correct* in thinking that I have done a better job with respect to evaluating the evidence and arguments than those with whom I find myself in disagreement. Suppose that they reason in a parallel way and conclude that *I'm* the one who has misjudged the evidence. On the present view, the rationality of the parties engaged in such a dispute will typically depend on who has *in fact* correctly evaluated the available evidence and who has not. If you and I have access to the same body of evidence but draw different conclusions, which one of us is being more reasonable (if either) will typically depend on which of the different conclusions (if either) is in fact better supported by that body of evidence. No doubt, especially in the kinds of cases at issue, it will often be a non-trivial, substantive intellectual task to determine what the totality of relevant evidence supports. Therefore, the question of which one of us is doing a better job evaluating the evidence will often be a non-trivial, substantive intellectual question. But here as elsewhere, life is difficult. On any plausible conception of evidence, we will be extremely fallible with respect to questions about what our evidence supports.¹⁵ The amount of disagreement that we find among well-informed individuals simply makes this fact more salient than would otherwise be the case.

On the present view, the rationality of one's believing as one does is not threatened by the fact that there are those who believe otherwise. Rather, any threat to the rationality of one's believing as one does depends on whether those who believe otherwise have good reasons for believing as they do—reasons that one has failed to accurately

¹⁵ Indeed, if Williamson (2000) is correct, then our ability to fully appreciate our evidence is subject to *in principle* limitations. However, even if one finds Williamson unconvincing on this point, one should still admit that we are in fact extremely fallible when it comes to evaluating large and diverse bodies of evidence. I discuss this fallibility further in § 6 below.

appreciate in arriving at one's own view. I explore this theme further in the following section.

5. RATIONALITY AND MERELY POSSIBLE DISAGREEMENT

Consider the circumstances in which we are apt to find disagreement intellectually threatening. Here, the first point to appreciate is the following: it is extremely implausible that *actual* disagreement is always more epistemically significant than certain kinds of *merely possible* disagreement. After all, whether there is any actual disagreement with respect to some question as opposed to merely possible disagreement might, in a particular case, be an extremely contingent and fragile matter. In particular, whether there is any actual disagreement might very well depend on factors that everyone will immediately recognize as irrelevant to the truth of the question at issue. (Suppose, for example, that there would be considerable disagreement with respect to some issue, but that all of the would-be dissenters have been put to death by an evil and intolerant tyrant.)

The existence of actual disagreement, then, need be no more intellectually threatening than certain kinds of merely possible disagreement. However, not *every* kind of merely possible disagreement will be intellectually threatening: the possibility that individuals who are insane or who are otherwise clearly irrational might disagree with some view that we hold would presumably *not* provide us with a good reason to doubt that view. The question, then, is this: under what circumstances should we find the possibility of disagreement intellectually threatening? Whether we find the possibility of disagreement intellectually threatening, I suggest, will and should ultimately depend on our considered judgements about *how rational* the merely possible dissenters might be in so dissenting. And our assessment of whether rational dissent is possible with respect to some question (or our assessment of the extent to which such dissent might be rational) will depend in turn on our assessment of the strength of the evidence and arguments that might be put forward on behalf of such dissent. But if this is correct, then the extent to which merely possible dissent should be seen as intellectually threatening effectively reduces to questions about the strength of the reasons that might be put forward on behalf of such

dissent. Now, there might be cases in which we judge that the arguments and evidence that could be brought forth on behalf of a hypothetical dissent are truly formidable, and this might justifiably make us doubt our own beliefs. But in that case, the reasons that we have for skepticism are provided by the state of the evidence itself, and our own judgements about the probative force of that evidence. The role of disagreement, whether possible or actual, ultimately proves superfluous or inessential with respect to the case for such skepticism.

Suppose that those members of the philosophical community who have both (i) thought seriously about Newcomb's Problem and (ii) are familiar with the main arguments on both sides are approximately evenly divided between One-Boxers and Two-Boxers.¹⁶ We can imagine various ways in which this state of disagreement gives way to a consensus. Here is one way: someone thinks of an ingenious argument that convinces all of the One-Boxers (or, alternatively, all of the Two-Boxers) that they have been in error up until now. Here is a second way: an evil and intolerant tyrant, bent on eliminating the scourge of One-Boxing once and for all, seizes power and initiates a systematic and ultimately wholly successful campaign of persecution against the One-Boxers. (Again, in these circumstances, I assume that the mere absence of disagreement is of no epistemic significance at all.) These cases, clearly, lie at opposite ends of a certain spectrum. Consider finally a third possible world in which disagreement about Newcomb's Problem is absent. In this possible world, there is no evil tyrant, nor is there any ingenious argument that inspires rational conviction in all of those who consider it. The only known arguments that are thought relevant to Newcomb's Problem are exactly those arguments that *we* presently possess. The only difference between this possible world and our own world is the following. In this possible world, everyone who has studied Newcomb's Problem happens to be a One-Boxer, because everyone who

¹⁶ In his original presentation of Newcomb's Problem, Robert Nozick wrote: 'I have put this problem to a large number of people . . . To almost everyone it is perfectly clear and obvious what should be done. The difficulty is that these people seem to divide almost evenly on the problem, with large numbers thinking that the opposing half is just being silly' (1969: 48). My sense is that the by-now over three decades worth of sustained debate on Newcomb's Problem has resulted in a significant shift in the original distribution of opinion in favor of a policy of Two-Boxing. But I will abstract away from this fact in what follows: what is crucial for my purposes is simply that there are some actual defenders of One-Boxing as well as some actual defenders of Two-Boxing. (In what follows, one might consider the actual world as it stood *circa* 1969, as reported by Nozick.)

has studied Newcomb's Problem has been convinced by the very arguments that convince One-Boxers in our world. There just are, as a matter of contingent fact, no actual defenders of Two-Boxing—although the known arguments for Two-Boxing are just as strong as the known arguments for Two-Boxing in our world. (Indeed, they are the same arguments.) Moreover, there is no *deep* explanation of *why* this is so—it is not as though the brain chemistry of the inhabitants of this world differs from ours, in a way that makes them peculiarly susceptible to the allures of One-Boxing. It is just that in this possible world, everyone who has thought about it up until now finds the case for One-Boxing more compelling, and there is thus complete consensus that One-Boxing is the uniquely rational strategy.

Do these empirical and contingent facts about the state of opinion make a difference about what it is rational to believe about Newcomb's Problem? Imagine an intelligent student who sets out to study Newcomb's Problem. She scrupulously exposes herself to all of the arguments and intuition pumps that favor One-Boxing and to all of the arguments and intuition pumps that favor Two-Boxing. In the process of thinking about the problem, she increasingly comes into contact with others who have thought about the problem, and she naturally begins to take note of their views. In our world, the student finds that roughly half of those she meets are One-Boxers and half Two-Boxers. In the other possible world, she finds that everyone she meets is a One-Boxer. Having thoroughly investigated the issue, she thus resolves to make up her own mind about Newcomb's Problem. Should she take a different view about Newcomb's Problem in the other, unanimous world than she does in our fragmented and divided world? *Despite* the fact that she has access to exactly the same arguments in both worlds? This seems extremely dubious—after all, can't the student in the unanimous possible world simply look over at our own fragmented world, and realize that *here* she has epistemic peers who extol Two-Boxing? But to judge that there are close possible worlds in which individuals can rationally take certain considerations as warranting a given belief is just to make a judgement about the probative force of those considerations themselves.

Interestingly, there are philosophical questions with respect to which the state of informed philosophical opinion in our world seems to be unanimous, or very close to unanimous, in much the way that the state of informed philosophical opinion about Newcomb's Problem is unanimous in my imaginary world. Consider the way in which radical forms

of skepticism—about the external world, or about other minds, for example—have traditionally been treated within epistemology. In view of how seriously skepticism has been taken throughout the history of philosophy, it is a striking fact that there have been relatively few genuine skeptics. I am unaware of a single contemporary philosopher, for example, who genuinely believes that she does not know whether anyone besides herself has a mind.¹⁷

There are, of course, various possibilities here. One possibility is that I am just factually wrong—there are, in fact, a significant number of philosophers who believe that they don't know whether anyone else has a mind, but, for understandable reasons, they don't bother announcing this belief to the rest of the world. (As Bertrand Russell once observed, there is no doubt something extremely odd about a genuine skeptic about other minds who makes a point of *professing* this belief to others.) Another, perhaps more important possibility is the following. It might be that there are a considerable number of individuals who would be genuine skeptics, but that it is simply *psychologically impossible* (or very nearly so) to believe the conclusion of a skeptical argument except at the moment when one is attending to the argument, if then. (Here I have in mind the kind of epistemic weakness of the will that Hume so famously made vivid in book 1 of the *Treatise*.) However, the fact that there have been few if any genuine skeptics about other minds is not, I think, primarily due to the fact that individuals find themselves simply psychologically unable to believe the conclusions of skeptical arguments. Rather, there have been, I think, very few individuals who have believed that there is some sound argument for skepticism about other minds. Of course, many philosophers have defended skeptical arguments by attempting to show that particular objections to their soundness are misguided, or even that all extant objections are misguided. Some philosophers no doubt believe that we have yet to produce good objections to skeptical arguments, or even that we can reasonably hope to find good objections to skeptical arguments in the future. But all of these broadly sympathetic stances *vis-à-vis* skepticism are much weaker than genuine skepticism, in the sense of believing that there is some sound argument that has as its conclusion 'I cannot know that

¹⁷ Peter Unger seems to have been an exception at the time of his (1975) but later writings reveal that his attitudes towards skepticism have evolved considerably; see e.g. his (1984: ch. 3).

there is any mind other than my own'. And it is in this sense, I think, that there have been few if any genuine skeptics about other minds.¹⁸

Nevertheless, the relative absence of genuine skeptics has not been taken to be a significant fact in the assessment of skepticism itself. That is, in assessing the case for skepticism, the discussion has been about the probative force of skeptical arguments. The contingent fact (assuming that it is a fact) that there are few if any philosophers who actually endorse some skeptical argument as a sound argument has not been taken to be relevant. To put it in another way: it has *not* been considered a good *objection* to skepticism to simply note that there are few if any genuine skeptics. We can, of course, easily imagine that things are otherwise—that is, we can imagine that philosophical opinion about the truth of skepticism about other minds is more or less evenly divided in our world between genuine skeptics and non-skeptics, in much the way that philosophical opinion is genuinely divided between One-Boxers and Two-Boxers. Would the case for skepticism about other minds be any stronger if it were so? Given that the best arguments offered by the genuine skeptics are simply *our* best arguments? In general, it has been thought—correctly, I believe—that the case for skepticism stands or falls with the probative force of skeptical arguments and does not depend on contingent and empirical facts concerning the actual existence or nonexistence of skeptics. As Laurence Bonjour has written: 'the need to consider skepticism does not depend in any crucial way . . . on whether or not serious proponents of skepticism are actually to be found; if skeptics did not exist, one might reasonably say, the serious epistemologist would have to invent them' (1985: 14–15).

6. THE VIEWS OF ONE'S PEERS AS HIGHER-ORDER EVIDENCE

It is a presupposition of the issue under discussion that we are fallible with respect to our ability to correctly appreciate our evidence. Of course, reasonable individuals are disposed to respond correctly to their evidence. But even generally reasonable individuals are susceptible

¹⁸ The relatively recent advent of skeptic-friendly varieties of contextualism (e.g. Lewis 1996) might cause some difficulties for this (admittedly rough) construal of what counts as 'genuine skepticism'. But not, I think, in a way that materially affects the point at issue.

to making mistakes on particular occasions. The possibility of error makes the following question salient: how do we know what our evidence supports? Could one have evidence which is relevant to the question of what one's evidence supports? If so, what would such evidence consist of?

Given that reasonable individuals are disposed to respond correctly to their evidence, the fact that a reasonable individual responds to her evidence in one way rather than another is itself evidence: it is evidence about her evidence. That is, the fact that a (generally) reasonable individual believes hypothesis H on the basis of evidence E is some evidence that it is reasonable to believe H on the basis of E. The beliefs of a reasonable individual will thus constitute *higher-order* evidence, evidence about the character of her first-order evidence. Of course, such higher-order evidence, like most other evidence, is not conclusive evidence: it does not follow from the fact that a generally reasonable individual believes H on the basis of E that it is reasonable to believe H on the basis of E. In a case in which E does not adequately support H but a generally reasonable individual mistakenly believes H on the basis of E, the fact that the individual believes as she does constitutes *misleading* evidence about the character of the evidence E. But misleading evidence is evidence nonetheless. In general, then, the fact that a reasonable person believes H on the basis of E constitutes evidence about the character of E.

Given the general reasonableness of one's epistemic peers, what they believe on the basis of one's shared evidence will thus constitute evidence about what it is reasonable to believe on the basis of that evidence. There are subtle questions, I think, about how one should integrate such higher-order considerations into one's own deliberations and what difference such considerations make to what it is reasonable for one to believe. At the very least, evidence about one's evidence will make a difference to what it is reasonable for one to believe *about one's evidence*. Will such higher-order evidence also make a difference to what it is reasonable for one to believe about propositions that are *not* about one's evidence? Let E represent our shared total evidence with respect to H. Consider the epistemic proposition that

- (1) E is good evidence that H is true

On the present view, if I discover that you believe that H on the basis of E, I should treat this discovery as confirming evidence for (1). Should

I also treat it as confirming evidence for H itself? If I discover instead that you believe that not-H on the basis of E, this discovery would constitute disconfirming evidence for (1). Would it also constitute evidence against H?

Here is a reason for thinking that I should *not* treat the evidence for or against (1) that is afforded by your believing as you do as evidence for or against H itself. Imagine that I have yet to make up my mind about H: that is, I am in the process of actively deliberating about what attitude I should take up towards the hypothesis. Suppose further that I find that you believe H on the basis of our shared first-order evidence E. If I treat the fact that you believe as you do as an additional piece of evidence which bears on the truth of H, then, when I enumerate the considerations which tend to confirm H, I will list not only the various first-order considerations that speak in favor of H, but also the fact that you believe that H is true. That I treat your belief in this way might seem to involve a certain admirable modesty or humility on my part. But notice that, when *you* enumerate the reasons why *you* believe that H is true, you will list the various first-order considerations that speak in favor of H—but presumably, *not* the fact that you yourself believe that H is true. From your perspective, the fact that you believe as you do is the *result* of your assessment of the probative force of the first-order evidence: it is not one more piece of evidence to be placed alongside the rest. That is, you do not treat the fact that you believe H as a further reason to believe that H, above and beyond the first-order considerations that you take to rationalize your belief. (If you subsequently changed your mind and came to doubt that the first-order evidence was sufficient to rationalize your believing H, you would not treat the fact that you believe that H as a reason to continue believing it. Similarly, when you first came to believe that H on the basis of your initial consideration of the first-order evidence E, you did not then proceed to treat the fact that *I believe that H is true* as a reason to *increase* your confidence that H is true. Rather, you arrived at that level of confidence which you thought appropriate given the first-order evidence E.) I am thus in the awkward position of treating your belief that H as a reason to believe that H, despite the fact that you do not treat this as an epistemically relevant consideration. Again, it might make sense for me to treat your belief in this way if I lacked access to your first-order evidence: in that case, your belief stands in as a sort of proxy for the evidence on which it is based (cf. § 3 above). But when I do have access to your

first-order evidence for H, and I continue to treat the belief that you have formed in response to that evidence as a further reason to believe that H, aren't I essentially engaged in a kind of double-counting with respect to the relevant evidence?¹⁹

Perhaps the relevance of my knowing that you believe as you do with respect to a given question is much like the relevance of an insurance company's knowledge that some particular driver happens to be a teenager. Because teenage drivers are, taken as a group, more reckless than other drivers, it makes sense for an insurance company to give a considerable amount of weight to this fact in particular cases. But if the insurance company had direct access to the underlying facts about the actual recklessness of a particular teenager, then this person's age would be rendered an irrelevant piece of information, and continuing to give weight to it would be to engage in a kind of illegitimate double-counting. In the language of the statisticians: access to the underlying facts about the actual recklessness of the driver 'screens off' knowledge of the driver's age, rendering the latter probabilistically irrelevant. Similarly, it might be that my having access to all of the evidence on which you base your belief screens off the fact that you believe as you do.²⁰

¹⁹ I have assumed that, when you enumerate the considerations that you take to bear on the truth of the hypothesis H, you will not include your own belief that H is true among those considerations. Consider, however, the view known as *epistemic conservatism*. According to epistemic conservatism, the mere fact that one presently believes that H makes it normatively appropriate to go on believing H, in the absence of positive reasons for abandoning that belief. Suppose that epistemic conservatism is, in fact, a correct view about belief revision. In that case—it might be argued—you *ought* to treat the fact that you believe that H as a reason to believe that H.

But this suggestion misunderstands the nature of epistemic conservatism. Adherents of epistemic conservatism typically do not present their view as implying that one possesses a reason to believe a proposition in virtue of believing that proposition. Rather, the view is that one does not *need* a reason for it to be normatively appropriate to continue believing a proposition that one already believes. (Beliefs are 'innocent until proven guilty', as opposed to the more traditional view that they are 'guilty until proven innocent'.)

For endorsements of epistemic conservatism, see Sklar (1975), Harman (1986), and Quine and Ullian (1978). For criticism, see Foley (1987), Vogel (1992), and Christensen (1994).

²⁰ Compare also the legal norm of 'Best Evidence'. If an original document is unavailable, a transcription of the original might be admitted as evidence of its author's intentions. But if the original document *is* available, then the transcription is considered inadmissible. The underlying thought, of course, is that while the transcription might have significant evidential value in the absence of the original, it is rendered irrelevant by the original's presence, since whatever evidential value it does have is exhausted by its (perhaps imperfect) reflection of the contents of the original. Similarly, one might think that, since the evidential value of the belief of some other party ultimately depends on the

At the very least then, there seems to be a certain awkwardness in my giving additional weight to your belief that H is true when I have already taken into account all of those considerations on which your belief is based, considerations which you take to exhaust the case for H. I do not think that this line of thought is decisive, however. Issues about how one's higher-order evidence does or does not interact with one's first-order evidence when that first-order evidence is itself available are, I think, extremely complicated. I will not attempt to resolve these issues here.²¹ Rather, in the remainder of this section, I will argue that even if we *do* treat the higher-order evidence that is provided by the views of our epistemic peers as further evidence that bears on the disputed questions themselves, it does not follow that skepticism or agnosticism is the reasonable response to disagreements of the relevant kind.

Again, let E represent our total evidence with respect to H at time t0. In order to avoid premature complications, let's suppose that each of us is ignorant of the other's existence at this point.²² Let's further stipulate that E is such as to rationalize the belief that H. Recognizing this fact, you form the reasonable belief that H at time t1, an instant later. Unfortunately, however, I badly misjudge the probative force of the evidence E at time t1 and take up the unreasonable belief that not-H.

At time t1 then, prior to our learning about the other person, the situation stands as follows. You hold the reasonable belief that H on the basis of your total evidence E while I hold the unreasonable belief that not-H on the basis of my total evidence E. The asymmetry in the epistemic statuses of our respective beliefs is due simply to the fact that E *really does* support H and does not support not-H.

fact that is a (perhaps imperfect) reflection of some more fundamental piece of evidence on which it is based, the belief is rendered irrelevant by the presence of the more fundamental piece of evidence (even if the same belief would be highly relevant in the absence of the more fundamental piece of evidence).

²¹ For further discussion, see Kelly (forthcoming *a*).

²² Of course, it might be that the most typical way for two individuals to have the same evidence is for them to have shared their evidence with one another—or at least, for both of them to be members of some community which shares its evidence (think of the Compatibilists and the Incompatibilists here). But it is, I assume, at least logically possible for two individuals to have arrived at the same evidence independently of one another. I want to begin, then, by considering what's true in a case in which you and I have the same evidence, but where both of us are ignorant of the fact that there is someone else who has exactly that evidence.

Suppose that we become aware of our disagreement at time t_2 . According to the view in question, our total evidence with respect to H has now changed. Let's call our new total evidence at time t_2 E' . What does E' include? E' will include the following:

- $E' =$ (i) the original, first-order evidence E ,
 (ii) the fact that you believe H on the basis of E , and
 (iii) the fact that I believe not- H on the basis of E .

The crucial fact here is the following: there is no reason to think that the new evidence E' will invariably mandate an attitude of abstention or agnosticism with respect to the hypothesis H . In particular, there is no reason to think that your continuing to believe H is unreasonable on evidence E' given that it was reasonable when your total evidence consisted of E . For in the usual case, the character of the new evidence E' will depend a great deal on the character of the original evidence E . Indeed, if we give equal weight to (ii) and (iii), then H will be more probable than not- H on the new evidence E' , given that it was more probable on the original evidence E . Our original evidence E does not simply vanish or become irrelevant once we learn what the other person believes on the basis of that evidence: rather, it continues to play a role as an important subset of the new total evidence E' . In general, what one is and is not justified in believing on the basis of E' will depend a great deal on the character of the first-order evidence E .

Thus, even if one treats the higher-order evidence which is provided by the beliefs of one's epistemic peers as evidence which bears on the disputed theses, it does not follow that agnosticism or suspension of judgement is the correct response to such disputes.

7. ACTUAL DISAGREEMENT RECONSIDERED

I have argued that disagreement does not have the kind of epistemic significance that has sometimes been claimed for it. Still, it would be a mistake to think that disagreement is therefore epistemically *insignificant*. What epistemic role or roles are left for disagreement, on the view that I have defended? Of course, an awareness of disagreement can serve to call one's attention to arguments that one might never have considered or to evidence of which one might have been unaware. However, even when all parties to a dispute have access to the same evidence and

arguments, disagreement can still play an epistemically salutary role. In the last section, I noted that the views of one's epistemic peers provide higher-order evidence. In this section, I want to highlight two other important roles that disagreement can play in cases of shared evidence.

First, it might be that the presence of disagreement with respect to some question at earlier times tends to produce a better pool of evidence bearing on that question at later times. That is, over time, the goals of inquiry might be best promoted when there is a diversity of opinions among inquirers. This theme has been endorsed and developed by a distinguished tradition of thinkers, a tradition which includes John Stuart Mill, Frederick Hayek, and Paul Feyerabend.²³

In addition, there is, I want to suggest, a more subtle way in which disagreement can prove epistemically beneficial. My suggestion is that the role of actual disagreement among epistemic peers is analogous to the role that actuality sometimes plays in falsifying modal claims that are mistakenly thought to be justified a priori.

Taken as a class, philosophers are somewhat notorious for making claims, ostensibly justified a priori, about *what must be the case*, or *what could not be otherwise*, that are subsequently falsified by empirical discoveries.²⁴ Not only does a putatively a priori necessary truth fail to hold in all possible worlds, it does not even hold in our own, actual world. (The logical positivists often accused Kant of making this mistake.) Presumably, there is a sense in which these empirical discoveries were not essential to falsifying the modal claim in question: someone with sufficient imaginative abilities would not need actual, empirically discovered counterexamples in order to see that the modal claim is false. But because human beings not infrequently suffer from persistent blindspots or failures of imagination, actuality occasionally plays a key role in falsifying such modal claims. (Once the modal claim is seen to be false, it can then come to seem *obviously* false; additional counterexamples are easy to come by, and it can seem almost embarrassing that we needed an empirical discovery in order to perceive its falsity.)

I suggest that something analogous is true of the role of actual dissenters. In principle, we ought to be able to give due weight to the

²³ Mill; Hayek (1960); Feyerabend (1975). A contemporary philosopher who has further developed this general theme is Philip Kitcher; see especially his (1993).

²⁴ For a recent excoriation of philosophers on this score, see Nozick (2001: esp. ch. 3).

available reasons that support a given view, even in the absence of actual defenders of the view who take those reasons as compelling. But in practice, the case for a view is apt to get short shrift in the absence of any actual defenders. The existence of actual defenders can serve to overcome our blindspots by forcefully reminding us of just how formidable the case is for the thesis that they defend, just as actual counterexamples are sometimes needed to overcome our blindspots concerning modality. But the case for a given view itself is no stronger in virtue of the fact that that view has actual defenders—just as a genuine counterexample to a modal claim is no stronger in virtue of being an actual, empirically discovered counterexample.

8. CONCLUSION: EPISTEMIC EGOISM WITHOUT APOLOGY

I have argued that disagreements of the relevant kind do not provide a compelling basis for skepticism. The mere fact that others whom I acknowledge to be my equal with respect to intelligence, thoughtfulness, and acquaintance with the relevant data disagree with me about some issue does not undermine the rationality of my maintaining my own view. I admit to finding this conclusion somewhat unsettling. Among my reasons for finding it unsettling is the following: many of those whom I take to be my epistemic peers disagree with me about this issue. Disappointingly, even some of those whom I would expect to be most sympathetic to my view given their own practice tend to argue against it as a matter of theory.

That I find it unsettling that many people I know and respect disagree with me about the epistemic significance of disagreement is perhaps unsurprising. There are, after all, psychological studies that suggest that we are highly disposed to being greatly influenced by the views of others, and I have no reason to think that I am exceptional with respect to this particular issue. It is, of course, a different question whether the fact that many others disagree with my thesis provides a good reason for me to doubt that thesis. And my answer to this question, as might be expected, is ‘No’: because I accept the general thesis that known disagreement is not a good reason for skepticism, I do not, in particular, regard the fact that people disagree with me about this general thesis as

a reason for being skeptical of *it*. Although I tend to find it somewhat unsettling that many disagree with my view, I am inclined to regard this psychological tendency as one that I would lack if I were more rational than I in fact am. In contrast to my psychological ambivalence, my considered, reflective judgement is that the fact that many people disagree with me about the thesis that disagreement is not a good reason for skepticism is not itself a good reason to be skeptical of the thesis that disagreement is not a good reason for skepticism.

The fact that I both endorse this thesis and refuse to take the fact that others disagree with me as a compelling reason for doubting its truth means that my views have a certain kind of internal coherence. This kind of internal coherence is not trivial: all combinations of views do not have it. However, I am not inclined to put too much weight on this kind of internal coherence, for this particular virtue proves surprisingly robust. Suppose, for example, that despite my considered judgement I one day do give in to the psychological pressure occasioned by the fact that so many of those who I know and respect disagree with me, and I abandon my thesis. (In the question-and-answer session following a talk at which I present these ideas, all of the questioners make it clear that they think that my thesis is clearly false. It is not that anyone provides a rationally compelling argument for this conclusion; rather, I am simply overwhelmed by my ever-increasing ideological isolation.) From my present vantage point, the envisaged change in my beliefs seems to be a craven (if understandable and all too predictable) capitulation to brute psychological pressure. *After* I have changed my mind about the epistemic significance of disagreement, however, it is of course open to me to look upon my recent conversion in a much more charitable light. I have changed my mind, after all, because I am influenced by the fact that others disagree with me, and this—according to the view that I will *then* hold—is the epistemically rational response to an epistemically relevant consideration. My later self might then say: my fundamental epistemic rationality—that is, the responsiveness of my beliefs to considerations that are in fact epistemically relevant—won out, in the end, over my misguided adherence to a mistaken philosophical thesis that would have permitted me to treat these epistemically relevant considerations as though they were irrelevant. So it looks as though, either way, a certain amount of self-congratulation will seem to be in order in the future.

REFERENCES

- Alston, William (1988) 'The Deontological Conception of Epistemic Justification', in James Tomberlin (eds.), *Philosophical Perspectives 20: Epistemology* (Atascadero, Calif.: Ridgeview), 257–99.
- Asch, Solomon (1952) *Social Psychology* (Englewood Cliffs, NJ: Prentice-Hall).
- (1956) *Studies of Independence and Conformity: A Minority of One against a Unanimous Majority*, *Psychological Monographs* 70(9).
- Aumann, Robert J. (1976) 'Agreeing to Disagree', *Annals of Statistics*, 4(6): 1236–9.
- BonJour, Laurence (1985) *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press).
- Christensen, David (1994) 'Conservatism in Epistemology', *Nous*, 28(1): 69–89.
- Feyerabend, Paul (1975) *Against Method* (London: Verso).
- Field, Hartry (2000) 'Apriority as an Evaluative Notion', in Boghossian and Peacocke (eds.), *New Essays on the A Priori* (Oxford: Oxford University Press), 117–49.
- Foley, Richard (1987) *The Theory of Epistemic Rationality* (Cambridge, Mass.: Harvard University Press).
- (1994) 'Egoism in Epistemology', in F. Schmitt (ed.), *Socializing Epistemology* (New York: Rowman & Littlefield), 53–73.
- (2001) *Intellectual Trust in Oneself and Others* (Cambridge: Cambridge University Press).
- Geanakoplos, John (1994) 'Common Knowledge', in *Handbook of Game Theory*, ii, ed. R. J. Aumann and S. Hart (Amsterdam: Elsevier Science), 1437–96.
- Gutting, Gary (1982) *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press).
- Harman, Gilbert (1986) *Change in View* (Cambridge, Mass.: MIT Press).
- Hayek, Friederich (1960) *The Constitution of Liberty* (Chicago: University of Chicago Press).
- Kelly, Thomas (forthcoming a). 'Confidence and Belief Revision', in John Hawthorne and Dean Zimmerman (eds.), *Philosophical Perspectives 20: Epistemology* (Oxford: Blackwell).
- (forthcoming b) 'Reasoning about Evidence one does Not Possess'.
- Kitcher, Philip (1983) *The Advancement of Science* (New York: Oxford University Press).

- Leher, Keith (1976) 'When Rational Disagreement is Impossible', *Noûs*, 10: 327–32.
- Lewis, David (1979) 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8: 339–59. Reprinted in his *Philosophical Papers*, i (Oxford: Oxford University Press, 1983).
- (1996) 'Elusive Knowledge'. Reprinted in his *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999).
- Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong* (New York: Penguin).
- Mill, John Stuart (1859) *On Liberty* (many editions).
- Nozick, Robert (1969) 'Newcomb's Problem and Two Principles of Choice'. First appeared in N. Rescher *et al.* (eds), *Essays in Honor of Carl G. Hempel*. Reprinted in Nozick's collected papers, *Socratic Puzzles* (Cambridge, Mass.: Harvard University Press, 1997), 44–73. Page references are to the reprinted version.
- (2001) *Invariances: The Structure of the Objective World* (Cambridge, Mass.: Harvard University Press).
- Plantinga, Alvin (1993) *Warrant: The Current Debate* (Oxford: Oxford University Press).
- (1995) 'Pluralism: A Defense of Religious Exclusivism', in Thomas Senor (ed.), *The Rationality of Belief and the Plurality of Faith* (Ithaca, NY: Cornell University Press).
- (2000) *Warranted Christian Belief* (Oxford: Oxford University Press).
- Quine, W.V., and J. Ullian (1978) *The Web of Belief* (New York: McGraw-Hill).
- Quinn, P., and K. Meeker (eds.) (2000) *The Philosophical Challenge of Religious Diversity* (New York: Oxford University Press).
- Sextus Empiricus (2000) *Outlines of Skepticism*, ed. Julia Annas and Jonathan Barnes (Cambridge: Cambridge University Press).
- Sidgwick, Henry (1981) *The Methods of Ethics* (Cambridge: Hackett).
- Sklar, Lawrence (1975) 'Methodological Conservatism', *Philosophical Review*, 74: 374–401.
- Unger, Peter (1975) *Ignorance: A Case for Skepticism* (Oxford: Oxford University Press).
- (1984) *Philosophical Relativity* (Minnesota: University of Minnesota Press).
- van Inwagen, Peter (1996) "'It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence'", in J. Jordan and D. Howard-Snyder (eds.) *Faith, Freedom, and Rationality* (Lanham, Md.: Rowman & Littlefield), 137–54.

Vogel, Jonathan (1992) 'Sklar on Methodological Conservatism', *Philosophy and Phenomenological Research*, 52: 125–31.

Williamson, Timothy (2000) *Knowledge and its Limits* (Oxford: Oxford University Press).