

Listening to Clifford's Ghost

PETER VAN INWAGEN

The Clifford of my title is W. K. Clifford, who is perhaps best known as the exponent of a certain ethic of belief – an ethic of belief that he was probably the first to formulate explicitly and which no one has defended with greater eloquence or moral fervor. In the lecture called, appropriately enough, 'The Ethics of Belief,'¹ Clifford summarized his ethic in a single, memorable sentence: 'It is wrong always, everywhere, and for any one, to believe anything upon insufficient evidence'. It will be convenient for us to have a name for this ethical thesis. I will call it 'ethical evidentialism' – 'evidentialism' for short.

Everyone I know of who has written on 'The Ethics of Belief' has taken it for granted that Clifford propounded evidentialism with a certain target in mind, and that that target was religious belief. In the last twenty years or so, however, philosophers have come to realize that a strong case can be made for the thesis that believing things without sufficient evidence is a pervasive feature of human life, a pervasive feature of the way we hold and acquire beliefs in the ordinary business of life, in politics, in matters pertaining to literature and the arts, and in science. And they have noted that failures to observe the dictates of evidentialism in these areas are not in the main 'near misses,' cases in which these dictates might easily have been observed if only people had been a little more careful about what they believed, if only they had taken a little more trouble to collect and examine evidence relevant to their beliefs. It seems, rather, that vast numbers of people believe things (things in no way related to religion or the supernatural) for which it is impossible for them to have sufficient evidence – if not impossible in principle, impossible for those people in the circumstances in which they in fact hold those beliefs.

My concern in this essay is not with religious beliefs or political beliefs or scientific beliefs or the beliefs on the basis of which we conduct the everyday business of our lives. It is with philosophical beliefs. I shall be concerned with the question whether any important philosophical belief is, or ever could be, held by anyone (philosopher

¹ *Lectures and Essays, Vol II* (London: Macmillan, 1879). Variousy reprinted.

Peter van Inwagen

or not) otherwise than upon insufficient evidence. And I shall be concerned only with philosophical beliefs that satisfy the following two conditions.

- (1) They are positive, not negative. What it means to say that a belief (proposition, thesis, conjecture, theory, hypothesis ...) is positive or negative is hard to explain in any philosophically satisfactory way, and I will not attempt to do so. I shall have to be content to give a few examples of philosophical beliefs or propositions that are paradigmatically *not* positive: 'Formalism is not the correct philosophy of mathematics'; 'Utilitarianism is not an acceptable ethical theory'; 'Knowledge is not simply justified true belief.' And, by the same token, 'Knowledge *is* justified true belief,' although it is no doubt a false thesis, is a positive thesis, and to assent to it is to have a positive philosophical belief. Formalism and utilitarianism – assuming that these terms have been sufficiently well defined that they denote particular propositions – are positive theses, and anyone who accepts formalism or accepts utilitarianism thereby has a positive belief.
- (2) They are not held by almost all human beings. I shall not be concerned with philosophical theses that have been accepted by all sane non-philosophers and have been denied only by a few philosophers – generally practitioners of 'revisionary metaphysics.' I assume that there are such philosophical beliefs because I assume that the denial of a philosophical belief is itself a philosophical belief, and many philosophers have believed things (in, as it were, their professional capacity) that almost everyone – even most philosophers – would deny. Or so it seems at least plausible to maintain. Plausible examples of things that fall into this category would be: 'Change and motion are not real features of the world'; 'One has no reason to suppose that there are minds other than one's own'; 'There are no material objects.'² (All these

² I say *plausible* examples, because questions concerning what is uncontroversial on the Clapham omnibus can be extremely controversial in the philosophical lecture-room. Berkeley notoriously maintained that no one but a few philosophers had ever believed in the existence of matter, and my former colleague José Benardete insists that Zeno believed nothing about change and motion that contradicted the beliefs of any of the passengers on the Clapham omnibus. I'll say this: I mean to consider only those philosophical beliefs that are, so to put it, uncontroversially controversial. It will be only these beliefs that will fall within the scope of the question

Listening to Clifford's Ghost

theses, or my statements of them, contain some sort of negative construction. Nonetheless, all of them are what I would call 'positive' theses. As I said, 'positive' is a very hard term to explain.³) Thus, philosophical beliefs like 'Change and motion are real features of the world,' 'One does have reason to suppose that there are minds other than one's own,' and 'There are material objects' do not satisfy my second condition. (I'll sometimes refer to beliefs that do satisfy the second condition as 'controversial,' simply because 'not held by almost all human beings' is a clumsy phrase.)

When I speak of philosophical beliefs, then, I mean my remarks to apply only to positive philosophical beliefs that are not beliefs that are held by almost all human beings. So to restrict my topic is not *severely* to restrict it: a vast range of philosophical beliefs satisfy both the conditions by which I have narrowed my subject-matter.

Let us ask: has any philosopher ever had sufficient evidence for any (positive, controversial) philosophical belief in Clifford's sense of 'sufficient evidence'? This question immediately raises a prior

I am asking – whether any philosophical belief is, or ever could be, held by anyone otherwise than upon insufficient evidence.

³ I would say that the negation of a negative belief must be a positive belief, but that the negation of a positive belief will in some cases also be a positive belief. An analogy is perhaps provided by the concept of positive and negative geographical information. That the spy whose whereabouts we should like to know is not in London is a negative piece of geographical information, and that he *is* in London is a positive piece of geographical information. That he is in the Western Hemisphere is a positive piece of geographical information, but so is the information that he is *not* in the Western Hemisphere – at least given that he must be either in the Eastern or the Western Hemisphere –, for the latter piece of information narrows down our range of possible specific hypotheses as to his location precisely as effectively as its negation does. I might put my point this way: 'Theism is false' is a positive philosophical belief because both theism and its negation, atheism, are philosophical *theories* or at any rate philosophical *positions*. 'Utilitarianism is false' is not a positive philosophical belief because its negation, non-utilitarianism, so to call it, is not a philosophical theory or position. There are many philosophical theories – many ethical theories – that are incompatible with utilitarianism, but non-utilitarianism, or the disjunction of all ethical theories (indeed, of all propositions) incompatible with utilitarianism, is not one of them: it's incompatible with utilitarianism all right, but it's not an ethical theory – and not a theory of any sort.

Peter van Inwagen

question: what is that sense? We may well ask, for Clifford never defines the phrase ‘sufficient evidence’. Perhaps this phrase requires no definition in the cases Clifford presents as paradigms of belief upon insufficient evidence. (For example, the famous case of the ship-owner who sent his ship to sea without having her overhauled and refitted, and who, although some doubts had passed through his mind as to whether she was really fit to sail, ‘succeeded in overcoming these melancholy reflections.’) In these cases, perhaps, we can just *see* that a certain belief was held upon insufficient evidence on *any* reasonable definition of ‘insufficient evidence’. But philosophical beliefs are not much like the belief that a certain ship is seaworthy, and questions about what counts as evidence – much less, sufficient evidence – for them are more difficult to answer. We shall require some sort of understanding of ‘sufficient evidence’ if we are to answer the question I have posed, or even to say anything of interest about it.

We shall, in fact, need to have some sort of understanding of three things: of ‘(a body of) evidence,’⁴ of what it is for one to ‘have’ a certain body of evidence (so understood), and of what it is for a certain body of evidence that one ‘has’ to be ‘sufficient’ to support some belief that one has.

I will not attempt to give general definitions of these terms (or accounts of these concepts). That would be a task far beyond my abilities. I will, however, try to say something about what these terms or concepts come to when they are applied to philosophical beliefs.

One form that evidence takes in philosophy is *argument*. One might even suppose that, in philosophy, evidence and argument are so closely related that, with care, the two can be identified. After all, if one has a (good) argument for some philosophical conclusion, then, surely, when one presents that argument to an audience, one presents one’s audience with evidence for its conclusion? And if one has evidence that supports a philosophical conclusion, could that evidence not be formulated as or presented in the form of an argument?

Whatever the answers to these rhetorical questions may be, it seems that arguments for philosophical theses are at least *one* kind of evidence for them. Whatever evidence may be, what one’s evidence for a certain belief is certainly has a great deal to do with how one

⁴ In present-day English, ‘evidence’ is a mass-term: one cannot (now) speak of ‘an evidence’ or ‘evidences’. And there is no corresponding count-noun. Various idiomatic phrases like ‘a piece of evidence’ or ‘a body of evidence’ perform the function of the missing count-noun.

Listening to Clifford's Ghost

would answer the question, 'Why do you think that?'. More exactly, it has a great deal to do with how one would answer that question when the question is understood in what one might call its epistemic sense. (I oppose 'epistemic sense' to 'psychological or causal sense'. Taken in its epistemic sense, it anticipates an answer like, 'I was there. I saw him do it.' Taken in its psychological or causal sense, it anticipates an answer like, 'Alice has been saying that he did it, and I dislike him so much that I suppose I'm inclined to believe anything discreditable about him.')

When this question is understood in its epistemic sense, it seems to be indistinguishable from the question 'What's your evidence for that?' And a philosopher will typically respond to the question 'Why do you think that?' (where 'that' is a philosophical thesis) by presenting one or more arguments for the thesis in question. It is, in fact, not easy to see what other kind of answer to this question there could be. It seems plausible to say that in philosophy evidence is argument – or at least that to *present* evidence is present one or more arguments.

If that is what evidence is in philosophy, what is it to 'have' the evidence for the conclusion of a certain argument that is, or is contained in, or is constituted by, that argument? The answer is pretty clearly this: It is to grasp or understand the argument. Or, if grasping or understanding an argument is a matter of degree: It is *fully* to grasp or understand the argument.

There may be some question as to what, exactly, is involved in grasping an argument. I do not want to build too much into this notion. I take it that one may grasp an argument (even fully grasp an argument) without having considered at length the possible replies to and objections to the argument, without having considered its possible implications, and without having raised the question whether similar or parallel arguments might lead to absurd conclusions. (One might, for example, fully grasp Anselm's ontological argument without having considered the question whether a parallel argument might be used to prove the existence of an island a greater than which cannot be conceived.) I would suppose, too, that it is possible at the same time fully to understand an argument and to believe, mistakenly, that it has false premises – or even to be mistaken about whether the argument is logically valid. (That case is something like this case: You and I are both looking at a sheep in a field; I, for one of those reasons that epistemologists are so skilled at contriving, mistakenly believe I am looking at a mock sheep, artfully crafted of papier-mâché. And you are under no such misapprehension. You and I have the same evidence for there being a sheep in the field; if I have a false belief about my

Peter van Inwagen

evidence for that thesis – even if I believe that it *isn't* evidence for that thesis –, that fact doesn't prevent its being evidence for that thesis, and evidence that I *have*.)

In light of these considerations, one may want to say that although evidence in philosophy indeed consists entirely of argument, the evidence relevant to a philosophical thesis p does not consist entirely of arguments whose conclusion is p or the denial of p . If we say that such arguments comprise the *primary* evidence relevant to p , we may designate those considerations that bear on the cogency of the arguments that comprise the primary evidence as *secondary* evidence that is relevant to p . The secondary evidence, like the primary evidence, will consist of arguments, but not arguments whose conclusion is p or its denial. The conclusions of the arguments comprising the secondary evidence will rather be propositions that concern the arguments for p or its denial: that this argument depends on an equivocation, that that one has a certain suppressed premise that needs to be considered carefully, that this one does not after all depend on an equivocation. (If there is secondary evidence for philosophical theses, there is no doubt tertiary evidence, and so, in theory, *ad infinitum*. But let us not go any further down that road, which is only a byway.)

To have the piece of evidence that is relevant to a philosophical thesis and is, or is contained in, or is constituted by, an argument is, I contend, simply to understand that argument – to understand it fully. Thus, if someone says to me, ‘Why do you think that free will is incompatible with determinism?’, and if, in reply, I produce a certain argument for that thesis (incompatibilism, it's called) – perhaps I write it on a blackboard – and if that argument is a complete statement of my reasons for accepting incompatibilism (no secondary evidence in this case), then you too will have the evidence for the incompatibility of free will and determinism that is my evidence for that thesis if you inspect the argument written on the blackboard and fully understand it. (And this could be the case even if, say, you believed that the argument contained a logical fallacy when in fact it didn't.)

Now, finally, what is it for a philosophical argument to be or constitute *sufficient* evidence for the philosophical thesis that is its conclusion? I am sorry to have to say that I do not know how to answer this question. Rather than try to answer it, I am going to explore one aspect of the concept of sufficient evidence (in philosophy). My exploration will take this form: I'll present an abstract, schematic case and proceed to ask a question about it. This is the case.

Listening to Clifford's Ghost

McX believes that p (I mean 'that p ' to be a philosophical thesis of the sort I have said I should consider: a positive thesis not held by all human beings). McX has no evidence for this thesis beyond that contained in or constituted by the philosophical argument A (an argument whose conclusion is of course the proposition that p): if you asked him why (epistemic sense) he believed that p , he'd produce the argument A for your consideration, and that would be a complete answer to your question; this answer would leave out none of his grounds for believing that p . McX's colleague Wyman grasps the argument A (fully) and believes neither that p nor that not- p : if you asked Wyman whether p , she'd say (sincerely) something like, 'I don't know' or 'I haven't been able to decide what to think about that' in reply. Although Wyman grasps the argument A fully, she is not convinced by it and remains an agnostic in the matter of the truth or falsity of its conclusion.

And this is the question.

Suppose McX is aware of these facts about himself and Wyman. What, if anything, should he conclude from them? What, in particular, should he conclude about whether he believes that p 'upon sufficient evidence'?

Here is *one* chain of reasoning that might go through McX's mind when he considers these facts.

If my evidence for my belief that p were indeed sufficient evidence, it would lead any intelligent, rational person who reflected on it to believe that p . But Wyman has the same evidence for the proposition that p as I have. I say this because I recognize that my evidence is entirely contained in the argument A, and Wyman – I am convinced – fully grasps that argument. I know that she agrees with me on *this* point: the argument contains no logical fallacy. I am also convinced that she is an intelligent, rational person, and that she has carefully reflected on the argument. I must, therefore, conclude that my belief that p is not based on sufficient evidence.

This chain of reasoning, I say, might occur to McX. But if it does – and if on reflection he accepts its conclusion and proceeds to give up his belief that p – he'll be, as the history of philosophy amply demonstrates, a most unusual philosopher. What alternatives might be open to him (other than ignoring the question of what to think about the

Peter van Inwagen

implications of Wyman's failure to be convinced by the argument)? There would seem to be two alternatives:

He might conclude that there's something wrong with Wyman. In becoming acquainted with the argument A, she has acquired sufficient evidence for p – and nevertheless refuses to accept p . And that implies that she is in some way defective. She's not, after all, an intelligent, rational person. Or she lacks philosophical ability or insight – at least in the degree to which he, McX, displays these qualities. He can *see* that certain propositions (certain premises of the argument) are conceptual or necessary truths, and she can't. Or she hasn't considered the disputed premises of the argument with sufficient care – despite the fact that she said she's been thinking about nothing else for a week. Or she's intellectually lazy or dishonest: she doesn't want to accept the conclusion of the argument because it would mean tearing up most of her own philosophical work and starting over or because it contradicts philosophical or religious or political beliefs to which she's strongly emotionally attached – with the consequence that she has *managed to convince herself* that propositions that are self-evident are doubtful or even false. In short, for one reason or another, Wyman is not being rational.

He might conclude that there's nothing wrong with either Wyman or himself. He might say that he and Wyman are both being perfectly rational. They've both carefully considered argument A; he's convinced by it, and that's okay; she's unconvinced by it, and that's okay. That's just how things go in philosophy.

In any real situation, both these alternatives can seem extraordinarily unappealing. Or, if we include the first alternative that I mentioned in the range of the alternatives we are considering ('I must, therefore, conclude that my belief that p is not based on sufficient evidence'), all three of these alternatives can seem extraordinarily unappealing.

We could sum up the three alternatives that confront McX this way:

There's something wrong with me. I believe that p , and my evidence is not sufficient to warrant belief that p (and that's bad).

There's something wrong with Wyman. Her evidence is sufficient to warrant belief that p , but she does not believe that p (and that's bad).

There's nothing wrong with either of us. I believe that p and Wyman does not believe that p and the evidence that each of us

has that is relevant to the truth-value of p is identical. Therefore: either it can be all right to believe something when one's evidence is insufficient to warrant one's belief or (inclusive) it can be all right *not* to believe something when one's evidence *is* sufficient to warrant one's having that belief.

It is important to realize that the abstract story of McX and Wyman is not the mere presentation of a logical possibility. There are real situations of exactly the sort that is laid out schematically in the story. This sort of thing *happens*, and – with few if any exceptions – each of us philosophers confronts alternatives of the sort that confront McX. Asking ourselves what we make of the fact that other philosophers are not convinced by arguments we ourselves find convincing is a task we can avoid only by the ostrich method.

I will cite a concrete case of such disagreement that I have often cited, a case in which I myself figure. I believe that free will is incompatible with determinism.⁵ What evidence can I appeal to in support of this belief? The most important part of this evidence can be presented in the form of an argument, an argument I have called the Consequence Argument. To make matters as simple as possible, let us pretend for the moment that the Consequence Argument comprises *all* the evidence I have for incompatibilism. That is to say, if you asked me, 'Why do you think that free will is incompatible with determinism?', I could do no better – and no more – than to write out one or more versions of the Consequence Argument for you and try to explain to you why I thought that each of its premises was true. (I'm going to count my defenses of the premises of the Consequence Argument as parts of the argument. If that sounds incoherent to you, I'll express myself this way: the *Narrow* Consequence Argument is a certain formally valid argument with numbered premises. The conclusion of the Narrow Argument is of course the proposition that free will and determinism are incompatible. The *Wide* Consequence Argument consists of the Narrow Consequence Argument plus everything I have to say in support of the premises of the Narrow Argument. For good measure, I shall include my definitions and explanations of the philosophical terms of art that occur in the Narrow Argument in the Wide Argument.

⁵ Despite the negative form of the word 'incompatible', I regard this as a clear case of a positive philosophical belief. Any appearance to the contrary is a linguistic accident – for suppose that instead of saying ' p is incompatible with q ' we used an expression that did not have a negative form (' p denies q ', perhaps, or ' p logically excludes q ').

Peter van Inwagen

When I speak of the Consequence Argument in the sequel, I mean the Wide Consequence Argument.)

David Lewis knew all about the Consequence Argument. In fact, he wrote a characteristically wonderful paper about it called 'Are We Free to Break the Laws?'.⁶ (This paper is the best defense of compatibilism that there is. It may well be the best paper about free will that there is.) He and I studied each other's arguments about the compatibility of free will and determinism carefully. We discussed the issues connected with this question carefully and at great length throughout the 1980s. I am therefore, I think, in a position to make this judgment: Lewis fully grasped the Consequence Argument. And he was not convinced by it. (He in fact accepted the denial of its conclusion. I'll presently incorporate this fact into my discussion. But let us pretend for the moment that Lewis simply failed to be convinced by the Consequence Argument; that he considered it carefully and was thereafter an agnostic about its conclusion.) If, therefore, I have no evidence for my belief that free will and determinism are incompatible but the evidence that is contained in or is constituted by the Consequence Argument, Lewis had all the evidence I had for the proposition that free will and determinism are incompatible, and yet did not accept this proposition. What should I conclude from this?

I should, of course, like to believe that I do not, in Clifford's phrase, accept this proposition upon insufficient evidence. But if the evidence I have for this proposition is sufficient evidence, why did Lewis, who had the same evidence, not also accept it? If it is epistemically wrong or irrational to accept a proposition upon insufficient evidence, is it not likewise wrong or irrational *not* to accept a proposition upon *sufficient* evidence? If I have sufficient evidence to support my belief that, say, the earth is more than 6,000 years old, and if I present a Young Earth Creationist with this evidence – if I 'present' this evidence to him in a way that has the consequence that he 'has' this evidence in the same sense as that in which *I* have it – is he not irrational if he does not come to share my belief? As I have said, I do not know how to give an account of sufficient evidence in philosophy (or in any other area of inquiry), but it is certainly plausible to suppose that whatever 'having sufficient evidence' (in any area, philosophy, geology, what have you), may be, it should bear the following relation to rationality: if one has

⁶ *Philosophical Papers, Vol. II* (New York and Oxford: Oxford University Press, 1986), 291–98. The paper first appeared in *Theoria* 47 (1981), 113–21.

Listening to Clifford's Ghost

sufficient evidence for a proposition and does *not* accept that proposition, one is irrational. Or will someone say that one might have sufficient evidence for a proposition one does not accept and not be irrational owing to the fact that one has not carefully considered the implications of that evidence? Well, I don't mind if someone says that. If someone is inclined to, I'll simply add to my statement a clause to accommodate that person's scruple: if one has sufficient evidence for a proposition *and has carefully reflected on the implications of that evidence for the truth of that proposition* and does not accept that proposition, one is irrational. That will not affect the problem with which Lewis's failure to be convinced by the Consequence Argument confronts me, for Lewis had certainly carefully reflected on the implications of the Consequence Argument for the truth of the proposition that free will and determinism are incompatible.

So. How shall I respond to this problem? Shall I say that there's something wrong with *me*? Shall I say, that is, that I do not have sufficient evidence for my belief that free will and determinism are incompatible? Or shall I say that there was something wrong with *Lewis*? Shall I say that, although he *should* have accepted the thesis that free will and determinism were incompatible as a result of considering the Consequence Argument (since that argument constitutes sufficient evidence for its conclusion), for some reason or other he didn't accept it? Or shall I say that there was nothing wrong with either of us? – that it is epistemically permissible for me to be convinced by a certain philosophical argument *and* it was epistemically permissible for him not to be convinced by that same argument? (Remember, we are not supposing that he understood the published piece of text that contained the argument differently from the way I did. No, it was the same argument, platonically speaking, that was in my mind and in his.) That is, shall I say that either the Consequence Argument does not constitute sufficient evidence for incompatibilism and it's all right for me to accept compatibilism on the basis of that argument alone, or that it *does* constitute sufficient evidence for incompatibilism, but it was all right for Lewis not to accept incompatibilism when he was in possession of that evidence?

All these alternatives, as I have said, are remarkably unappealing. I still think that the Consequence Argument shows that free will and determinism are incompatible. I find I can't help thinking that. But why doesn't Lewis see that if it's true? Was Lewis stupid? Lacking in philosophical ability? Intellectually dishonest? I certainly can't believe any of those things. Look, it's *David Lewis* we're talking about here. I can remember a talented young philosopher saying to me in the 1970s, following his first encounter with Lewis, 'Lewis is

Peter van Inwagen

so smart it's *scary!*'; and that has been more or less the response of all philosophers who have measured themselves against that formidable mind. Nor could anyone suggest with a straight face that Lewis was lacking in philosophical ability – not unless *all* human beings are lacking in philosophical ability. And he was scrupulously honest: he may have believed one or two odd things, but he did *believe* them, and believed them because he thought that they were straight-forward objective truths.

Suppose, then, I say that there's nothing wrong with either my being convinced by the Consequence Argument or Lewis's failure to be convinced by it. Suppose I tell myself that that's just how things go in philosophy. There are arguments that some philosophers find convincing and others don't, and it's *okay* to regard the philosophical arguments that one finds convincing as having established their conclusions if one has considered them carefully and responsibly. And it's *okay* for some other philosopher not to find those same arguments convincing provided he or she has also considered them carefully and responsibly. In a word, Lewis and I were both rational – or at least may well have been.

It is now time to take account of a fact that I have been ignoring. I have conceded parenthetically that Lewis did not merely refrain from accepting incompatibilism: he accepted its denial, compatibilism. And he did not accept compatibilism simply because he had examined that thesis and discovered within himself a conviction that it was true. He accepted compatibilism on the basis of certain arguments – arguments whose essential point is as old as Hobbes's debate with John Bramhall, Bishop of Derry, about liberty and necessity. It is necessary to add that these old arguments were not the only ones that played a role in his assent to compatibilism. An argument of *his*, an argument for the conclusion that the Consequence Argument turns on an equivocation also played a role in his assent (this is a case of what I have called secondary evidence). And, having brought that argument into our discussion, I can no longer maintain the pretense that the Consequence Argument constitutes the entirety of the evidence I have that is relevant to the question of the compatibility of free will and determinism. I have some secondary evidence of my own; if nothing else, *my* argument for the conclusion that *Lewis's* argument fails to show that the Consequence Argument turns on an equivocation, is a part of the evidence I have that is relevant to that question.

There were, therefore, other arguments than the Consequence Argument 'in play' in Lewis's and my decade-long discussion of the free-will problem. But, however many arguments were involved

Listening to Clifford's Ghost

in our debate, we both knew about them all and both fully grasped every one of them. Our situation was therefore more nearly symmetrical than I have been making it out to be. I believe that I fully understood all the arguments that constituted Lewis's evidence (primary and secondary) for the proposition that free will and determinism are compatible, and that I therefore 'had' the evidence on which his belief that free will and determinism are compatible was grounded. And, of course, I was not convinced by those arguments. There was, therefore, a certain body of evidence – it comprised the Consequence Argument and all the other arguments that figured in our debate – such that Lewis and I both had this evidence and such that, on the basis of this one body of evidence, I accepted a certain proposition and he accepted its denial.

The position that, in this set of circumstances it was all right for me to accept incompatibilism *and* all right for Lewis to accept compatibilism is not one that it is easy to be entirely comfortable with. (Let's describe the position this way: it was rational for me to accept incompatibilism and rational for Lewis to accept compatibilism.) If I contend that both Lewis and I were rational, I hear Clifford's ghost whispering an indignant protest. Something along these lines (Clifford has evidently acquired, *post mortem*, a few turns of phrase not current in the nineteenth century).

If you and Lewis are both rational in accepting contradictory propositions on the basis of identical evidence, then *you* accept one of these propositions – incompatibilism – on the basis of evidence that does not direct you toward incompatibilism and away from compatibilism. (For if it did, it would have directed *him* away from compatibilism, and it would not have been rational for him to be a compatibilist.) But of all the forces in the human psyche that direct us toward and away from assent to propositions, only rational attention to relevant evidence *tracks the truth*. Both experience and reason confirm this. And if you assent to a proposition on the basis of some inner push, some 'will to believe,' if I may coin a phrase, that does not track the truth, then your propositional assent is not being guided by the nature of the things those propositions are *about*. If you could decide what to believe by tossing a coin, if that would actually be effective, then, in the matter of the likelihood of your beliefs being true, you might as well do it that way.

I am unwilling to listen to these whispers. And I find it difficult to answer them. (No doubt these two facts are connected. I am unwilling to listen to the whispers of Clifford's ghost – if I listen to them, it

Peter van Inwagen

is only because I force myself to – *because* I find them so difficult to answer.)

Could it be that the difficulty I find myself in is based on some false assumption, an assumption hidden somewhere in the various lines of reasoning I have presented? That's certainly an attractive thought. But what might this assumption be – or these assumptions, if there's more than one? Here's a candidate that in some moods I can find appealing: the assumption that all evidence for a philosophical proposition can be presented in the form of an argument. Evidence that can be presented in the form of an argument is essentially public. Any argument can be written down on a blackboard, and – so I have supposed – anyone who studies what's written on the blackboard and understands it thereby 'has' the evidence comprised in the argument. Suppose, however, that there's such a thing as interior, incommunicable evidence for certain propositions: evidence that can somehow be present to one's mind, although one is unable to articulate it, unable to put it into words, unable to present in the form of an argument.

Whether or not there is evidence of this sort for philosophical propositions, there are plausible examples of it in other areas. I sometimes know that my wife is angry when no one else does, for example, and I can't explain to anyone (even to myself) how I know this – I can't give what Plato would call an 'account' of what underlies my conviction that she is angry. It seems to me to be plausible to say that in such cases my belief that my wife is angry is grounded in certain evidence, evidence that I cannot put into words. After all, although I usually turn out to have been right about her being angry, if someone asks me, 'How did you know she was angry?', I can give no answer. Mathematics provides a very different kind of example of this phenomenon. Mathematicians are often intuitively certain that some mathematical proposition is true, although they are unable to prove it. (Gödel, I understand, was convinced that the power of the continuum was \aleph_2 , but was unable to give any statement of the ground of this conviction.) Since they often later do discover proofs of these propositions, it seems likely that, prior to their discovery of the proofs, they had some sort of evidence that those propositions were true. Now maybe the evidence they had is exactly the evidence that they would later present in the form of a proof (on those occasions on which they did later produce a proof) although for some considerable period they were unable to articulate it. It is not essential to the suggestion that I am canvassing that 'inarticulable' evidence be *essentially* or *in principle* inarticulable. The suggestion requires only that a person have at a certain time

Listening to Clifford's Ghost

evidence that he is not *then* able to articulate. Might it not be that the following two theses are both true?

- (a) I have sufficient evidence for my belief that free will and determinism are incompatible; some of this evidence is contained in the Consequence Argument, but other parts of it are either in principle interior and inarticulable or else evidence that could in principle be presented in some public form, but which, for some reason, I am at present unable to put into that form.
- (b) Lewis did not have this interior evidence that I have. I thus have more evidence that bears on the thesis that free will and determinism are incompatible than he had. His failure to accept incompatibilism was a rational response to the body of evidence he had, and mine is a rational response to the (more extensive) body of evidence that I have.

It is important to realize that thesis (b) does not imply that I am smarter than Lewis or a better philosopher. The idea is rather this. Owing to some neural accident, I have a kind of insight into the, oh, I don't know, entailment relations among various of the propositions that figure in the compatibilism/incompatibilism debate that was denied to Lewis. I *see*, perhaps, that a certain proposition *p* entails the proposition *q* (although I'm unable to formulate this insight verbally) and he was unable to see that *p* entailed *q*. And this insight really is due to a neural *quirk* (to borrow a phrase Rorty used for a different purpose). It's not that my cognitive faculties function better than Lewis's. His were as reliable as mine – no doubt more so. But his were not identical with mine, and some accidental feature of my cognitive architecture has enabled me to see an entailment that he was unable to see. (If it's open to me to say this, it would, of course, have been open to Lewis to say the same thing *mutatis mutandis*, to have contended that *he* had a body of interior, inarticulable evidence that *I* lacked and that his total evidence *vis-à-vis* the question of the compatibility of free will and determinism was more extensive than mine. It is imaginable, in fact, that we might both say this – 'this,' of course, being in each case appropriately tailored to the convictions of the speaker –, and might each regard the other as mistaken, perhaps excusably mistaken, perhaps not. Each might suppose that the other had mistaken a merely subjective conviction that some entailment held for *seeing* that that entailment held.)

I have raised the question whether (a) and (b) might not both be true. This question suggests a further question. According to (a), I have sufficient evidence to warrant my belief that free will and determinism are incompatible. According to (b), Lewis had less evidence

Peter van Inwagen

that was relevant to the question of the compatibility of free will and determinism than I and his belief that free will and determinism were compatible was rational, given the evidence that was available to him. But (b) says neither that Lewis's evidence was sufficient to warrant his belief that free will and determinism were compatible nor that it was not. The 'further question,' of course, is: Was his evidence sufficient to warrant his belief or was it not? I should like to think that it was. I find it uncomfortable to suppose that my evidence was sufficient and Lewis's was insufficient, even if in this case his believing something upon insufficient evidence was somehow excusable. But if I suppose that it was, I face this difficulty: My evidence is, if we interpret this statement very literally, not sufficient to warrant a belief in the compatibility of free will and determinism – not at least if the same body of evidence cannot be sufficient to warrant both a certain belief and its negation (since it is sufficient to warrant my belief that free will and determinism are incompatible). But Lewis's evidence was a *proper part* of my evidence. If Lewis's evidence was sufficient, it would follow that a certain body of evidence was not sufficient to warrant a certain belief, but a proper part of that evidence was sufficient to warrant that same belief. And that seems counterintuitive. It is not clear, however, that this thesis, counterintuitive though it may be, is false. Suppose, for example, that Superman and Lois Lane are looking at a field and that Lois is having visual experiences of the kind that any normal human being who was looking at a field in which there was a single sheep would have. Lois believes that there is a sheep in the field before her, and it would seem that she has sufficient evidence for this belief if any human beings ever have sufficient evidence for any of their beliefs. Superman, more than human, has the evidence that Lois has and more besides: the evidence provided by his X-ray vision, which faculty reveals to him that what appears to the mere human eye to be a sheep is one of those epistemologists' mock-ups of a sheep. It seems therefore that he has sufficient evidence for a certain proposition and that Lois has sufficient evidence for its denial and that her evidence is a proper part of his. Perhaps Lewis's belief and his evidence, on the one hand, and my belief and my evidence, on the other, are related in the same way: my evidence consists of his evidence (the evidence provided by certain philosophical arguments) together with further evidence, interior incommunicable evidence, that is mine alone; nevertheless, despite the fact that his evidence is only a proper part of mine, my evidence is sufficient to support my belief that free will and determinism are incompatible and his evidence was sufficient to support his belief that free will and determinism were compatible.

Listening to Clifford's Ghost

On reflection, however, the idea that I have evidence, incommunicable evidence, that Lewis lacked, tempting though it is, is hard to believe. At any rate, it's hard to believe that it applies in all cases in which I disagree with other philosophers about some philosophical proposition or other. After all, I accept *lots* of philosophical propositions that are denied by many able, well-trained philosophers. Am I to suppose that in every case in which I believe something many other philosophers deny (that is, in every case in which I accept some controversial philosophical thesis), I'm right and they're wrong, and that, in every such case, my evidence is superior to theirs – owing to the fact that in every such case my evidence incorporates interior, incommunicable evidence that is somehow inaccessible to those other philosophers? If I do suppose that, I must ask myself, is the neural quirk that gives me access to this evidence the same neural quirk in each case or a different one? If it's the same one, what I am postulating looks more like a case of 'my superior cognitive architecture' than a case of 'accidental feature of my cognitive architecture.' If it's a different one in each case – well, that's quite a coincidence, isn't it? All these little evidence-friendly neural quirks come together to give the right results in just one philosopher (no other philosopher agrees with me about very much), and that philosopher happens to be me.

It seems more plausible to reject the idea of interior, incommunicable evidence and to concede (to revert to the case of David Lewis and myself) that I have and Lewis had the *same* evidence in the matter of the problem of free will. But if this is so, then either at least one of us has believed something upon insufficient evidence, or else I accept incompatibilism upon sufficient evidence and Lewis accepted compatibilism upon sufficient evidence and the evidence that the two of us had that bears on the compatibility of free will and determinism is the *same* evidence.

I will not try to say which of these disjuncts is right. I will instead conclude with some remarks about what a philosopher who believes either is committed to.

Suppose a philosopher accepts the first disjunct. (I'll call this philosopher 'you'.) You believe that at least one of the two of us, Lewis and me, accepted a certain philosophical position upon insufficient evidence. Then you must conclude that you and any philosopher who disagrees with you about the truth-value of some philosophical thesis are in the same position: one of you, at least, accepts a certain thesis upon insufficient evidence. Unless you are willing to say that *you* accept the thesis in question upon insufficient evidence (presumably you are not), you must conclude

Peter van Inwagen

that you accept the thesis upon sufficient evidence and that your colleague accepts its denial upon insufficient evidence. And, surely, you will agree that there are *many* such theses – many positive, controversial philosophical theses that you accept and other philosophers deny? Let me ask you this: Do you really find it plausible to suppose that *in all or most such cases*, one of the two of you accepts a thesis upon sufficient evidence and the other upon insufficient evidence and that *you* are the one with the sufficient evidence? (The alternative is to suppose that you accept a high proportion of the philosophical theses you accept upon insufficient evidence.) We might, indeed, direct this point at Clifford himself – for the simple reason that he is one of us, a philosopher. One very good example of a philosophical thesis that Clifford accepts is, of course, the thesis we have been discussing: ethical evidentialism. Ethical evidentialism is a positive, controversial philosophical thesis. (William James rejected it, and other philosophers – Roderick Chisholm⁷ and myself, for example – have expressed doubts about it.) Clifford has, of course, presented arguments for ethical evidentialism – rather good arguments, as philosophical arguments go. But is he really in a position to contend that these arguments constitute sufficient evidence for ethical evidentialism – given that other competent philosophers fully grasp these arguments (have the evidence he has) and do not embrace ethical evidentialism?

Now the second disjunct: that Lewis and I accept contradictory propositions on the same evidence, and that this evidence is in both cases sufficient. I want to make just this point: *Clifford* cannot accept this disjunct. I concede that the second disjunct is not logically inconsistent with Clifford's thesis, with ethical evidentialism. Consider, for example, one of those religious beliefs that were the intended 'target' of ethical evidentialism. The following three propositions are certainly logically consistent:

It is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence

Professor Dawkins believes that there is no God; the total body of evidence that he has that is relevant to the existence or non-existence of God is E; E is sufficient evidence for his belief that there is no God.

⁷ See *Perceiving: A Philosophical Study* (Ithaca, N.Y.: Cornell University Press, 1957), 9. and 99–100.

Listening to Clifford's Ghost

Archbishop Williams believes that there is a God; the total body of evidence that he has that is relevant to the existence or non-existence of God is E; E is sufficient evidence for his belief that there is a God.

But if the second disjunct is consistent with ethical evidentialism, it is nevertheless inconsistent with an essential premise of the *argument* by which Clifford claims to establish ethical evidentialism.

Why does Clifford think that it is wrong to believe things upon insufficient evidence? The central nerve of Clifford's reason for supposing this is contained in some words I put into his mouth a moment ago: Of all the forces in the human psyche that direct us toward and away from assent to propositions, only rational attention to relevant evidence *tracks the truth*. Believing things *only* upon sufficient evidence is, therefore, the only device we have for minimizing the extent of our false beliefs, or at least the only such device that has any prospect of providing us with a useful set of true beliefs. (One could, of course, very effectively minimize the extent of one's false beliefs by believing nothing.) If we form our beliefs on any other basis – if we allow them to be formed by some factor that does not track the truth –, we are, in effect, believing things at random. If I form my beliefs on some basis other than rational attention to evidence, no doubt there will be a causal explanation of some sort for what I believe, but the truth of falsity of those beliefs will not figure in that explanation. Since there are a lot more ways to be wrong than there are to be right, beliefs formed by a method that does not track the truth will, to a high probability, be false. (Recall the 'electric monk' in one of the Dirk Gently books, who, owing to a malfunction in his electrical innards, had begun to believe things at random, and who, at one point in the narrative had spent the morning believing that forty-seven per cent of all tables were hermaphrodites. The example illustrates nicely the high probability of a randomly chosen proposition's being false.) A person who believes things upon insufficient evidence, therefore, is not taking care to minimize the extent of his false beliefs. And any moral person *will* take care to minimize the extent of his false beliefs. This is the moral course of action because a person with false beliefs is *ipso facto* dangerous: a driver on British roads who believes that in Britain one drives on the right-hand side of the road is dangerous indeed – as is a ship-owner who believes that his ship is seaworthy when she is not. Any moral person, obviously, will want to minimize the danger he presents to himself and others, and an essential part of realizing that end is to believe only those things for which one has sufficient evidence.

Peter van Inwagen

This argument, I contend, is the core of Clifford's defense of ethical evidentialism. There is more to his defense than this core argument, of course. Suppose, for example, that someone had asked Clifford the following rather obvious question: 'Can you really suppose that a philosopher who has false beliefs about the reality of universals or the proper analysis of causation is *ipso facto* dangerous?'. One part of Clifford's defense of ethical evidentialism is, in effect, an answer to this question (and it is a thoughtful and interesting answer). But for my present purposes, I need consider only the core argument. It is evident that anyone who accepts this argument cannot suppose that a certain body of evidence can be sufficient to support both a belief that p and a belief that not- p . For, if that were the case – and particularly if it were a common occurrence –, rational attention to evidence would not track the truth. If you believe that in Britain one drives on the left and if I believe that in Britain one drives on the right, and if the evidence that you and I have that is relevant to the question which side of the road one drives on in Britain is the same, and if this evidence is sufficient in both our cases, then rational attention to evidence does not track the truth – and making sure that one has sufficient evidence for one's beliefs therefore provides no assurance that one is not a dangerous repository of false belief. In the present case, I am a dangerous driver and you, no doubt, are not – but your basing your belief concerning the side of the road to drive on upon sufficient evidence is not what prevents you from being a dangerous driver, for I did the same thing and everyone had better steer clear of me – literally – when I'm behind the wheel. If, therefore, one accepts the second of the disjuncts on offer, one can accept ethical evidentialism, if at all, only on some basis other than the argument by which Clifford defends it.

Can the philosopher who accepts ethical evidentialism say anything in defense of his or her accepting any positive and controversial philosophical thesis? (And remember: ethical evidentialism is itself a positive and controversial philosophical thesis.) I cannot see any very plausible avenues for the ethical evidentialist to explore. I conclude that philosophers should find ethical evidentialism an unattractive thesis – as I do. But what are philosophers to say in response to Clifford's argument for ethical evidentialism? It certainly does seem clear that, for just the reason Clifford cites, *many* propositions are such that a moral person will accept them only upon sufficient evidence. Might a philosopher contend that that this stricture does not hold for all propositions, and that philosophical propositions are among those for which it does not hold? Supposing that Aristotle

Listening to Clifford's Ghost

was right to think that universals exist only *in rebus* (they might ask), did Plato's belief that universals exist *ante res* make him a dangerous man? Might this philosopher not appeal to the authority of Hume in the matter of errors in philosophy: While errors in religion (and, presumably, in politics and medicine and many other areas, including the Highway Code) are dangerous, '... errors in philosophy are only ridiculous.'? There is much that might be said in response to the thesis that false philosophical beliefs are harmless. Two of the things that might be said, and they're the only ones I will say, are that history demonstrates that wrong ideas in philosophy have done a lot of harm, and that Plato's political beliefs did not exist in isolation from his metaphysical beliefs.

If it is conceded that it is wrong to accept philosophical propositions – positive, controversial ones – otherwise than upon sufficient evidence, and if it is conceded that the same evidence cannot be sufficient for contradictory propositions, and if it is conceded that interior, incommunicable evidence plays no significant role in philosophy, there seem to be only two choices open to a philosopher who is unwilling to embrace immorality. (I will remark that I have encountered only one philosopher who has made the first of these choices and only one who has made the second.) First, the philosopher might insist that he *does* have sufficient evidence for his philosophical beliefs and that those philosophers who disagree with him on any substantive philosophical point do *not* have sufficient evidence for their beliefs; those other philosophers are irrational or lacking in philosophical ability or unintelligent or uninformed or intellectually dishonest or exhibit some other such cognitive or epistemic defect. I can only say that I regard any philosopher who embraces that option as a comic figure. Secondly, the philosopher might choose to accept *no* philosophical theses (other than negative theses and uncontroversial ones) – not even the thesis that accepting no philosophical theses is the only morally permissible course of action for a philosopher. This philosopher I regard not as a comic but as an heroic figure. I have nothing to say about such heroism other than that few of us other philosophers are likely to imitate it. I certainly am not.