

# Williams on Ethics, Knowledge, and Reflection\*

A. W. MOORE

1. One of the principal aims of Bernard Williams's work in moral philosophy<sup>1</sup> is to provide a *critique* of ethical experience. This he does with great sensitivity and force. His work shows admirably how much moral philosophy can achieve. It is somewhat ironical, then, that one of his best known contentions concerns how little moral philosophy can achieve: moral philosophy cannot deliver the very thing which might have been expected of it, an *ethical theory* to guide moral reasoning. What it can do is to assist the self-understanding of those whose moral reasoning already has guidance from elsewhere. But as soon as it attempts to offer guidance of its own, in the shape of a theory, the question arises of where it gets its authority from to do this, sheer reasoning being in Williams's view too abstract for such practical purposes.

This conclusion is disappointing for those who hoped that we could argue somebody into morality. But Williams is adamant that this was always an ill-conceived hope.<sup>2</sup> No amount of argument can bring somebody to a sympathetic concern for the welfare of others, say. What she needs is something to expand her emotions: this will most probably be a suitable upbringing. But Williams does not see this as a defeat for rationality. It is rather a question of what kinds of forces are at work. Someone who has been suitably brought up will be incapable of doing certain sorts of things, and incapable of

\* This essay is derived from a lecture entitled 'Bernard Williams', delivered at Oxford University in 2000, in the series 'Oxford Philosophers on Oxford Philosophers', organized by Peter Hacker and David Wiggins. I am grateful to those who attended the lecture, and to Bernard Williams, for helpful comments.

<sup>1</sup> In his more recent writings Williams assigns different meanings to 'morality', 'moral', *etc.* and 'ethics', 'ethical', *etc.*: see e.g. *Ethics and the Limits of Philosophy* (London: Fontana, 1985), 6. (And see below, §2, for a brief account of the meaning that he assigns to the former.) I shall not follow this usage in this essay.

<sup>2</sup> E.g. 'The Amoralist', in his *Morality: An Introduction to Ethics* (Cambridge: Cambridge University Press, 2nd edition, 1993).

not doing others. But any *reason* she has for doing anything must ultimately be subject to such constraints rather than set them.<sup>3</sup>

More generally, Williams thinks that it is impossible for anyone to have a reason for doing anything which is not grounded in some suitably given value, desire, project, commitment, or the like—in some *motivation*, for short.<sup>4</sup> Herein is further cause for concern about the idea of an ethical theory. Theories aim at a tidiness, a systematicity, and an economy of ideas that are quite inappropriate where motivations are concerned. In particular, theories aim to eliminate conflict, whereas it is quite possible for motivations to conflict and (hence) for someone to have ineliminable reasons for doing each of two incompatible things.<sup>5</sup>

What then of the problem of *interpersonal* conflict? And what of the associated problem of moral disagreement? Williams's conception of moral reasoning seems to threaten hopes for objectivity. It seems to entail that moral reasoning is not ultimately answerable to anything; that sufficiently diverse motivations might generate moral disagreements that are irresolvable. Is this a real threat? And if it is, how disturbing are the consequences for morality?

Much of Williams's work in moral philosophy is devoted to addressing these very questions—but in an effort to clarify and refine them as much as to answer them. Blank pronouncements about objectivity conflate all sorts of issues, some logical, some epistemological, some metaphysical, and Williams is at great pains to disentangle these. Thus reconsider the idea that moral reasoning is not ultimately answerable to anything. Here we must distinguish the question of whether moral reasoning is constrained to go in a certain direction from the question of whether it is about a reality that exists independently of it. Kant, for example, believed that it had the first of these features but not the second.

Still, does not Williams's conception entail that it has neither? Does it not entail that there is a basic distinction between fact and value—between what is true irrespective of what anyone thinks about it and what is a matter of people's supplementary evaluations given

<sup>3</sup> Cf. 'Moral Incapacity', reprinted in his *Making Sense of Humanity and Other Philosophical Papers: 1982–1993* (Cambridge: Cambridge University Press).

<sup>4</sup> 'Internal and External Reasons', reprinted in his *Moral Luck: Philosophical Papers 1973–1980* (Cambridge: Cambridge University Press, 1981).

<sup>5</sup> 'Ethical Consistency', reprinted in his *Problems of the Self: Philosophical Papers 1956–1972* (Cambridge: Cambridge University Press, 1973).

## Williams on Ethics, Knowledge, and Reflection

their differing motivations—so that moral reasoning, insofar as it relates to the latter, is not answerable to anything in either sense?

There can be no doubt that Williams's conception entails *something* like that.<sup>6</sup> But it would be a grotesque oversimplification to put it in those terms. For Williams, a moral judgment can be straightforwardly true or false. Thus consider a moral judgment involving what Williams would call a 'thick' ethical concept, such as *infidelity*. By a 'thick' ethical concept Williams means a concept whose applicability is both 'action-guiding' and 'world-guided'. To apply a thick ethical concept in a given situation, for example to accuse someone of infidelity, is, in part, to evaluate the situation, which characteristically means providing reasons for doing certain things; but it is also to make a judgment which is subject to correction if the situation turns out not to be a certain way, for example if it turns out that the person who has been accused of infidelity did not in fact go back on any relevant agreement. (It is in this latter respect that thick ethical concepts differ from more general prescriptive concepts such as *rightness*.) A judgment involving a thick ethical concept can be indissolubly both moral and answerable to how things are, true or false as the case may be.<sup>7</sup>

The crucial issue, Williams thinks, is how such a judgment is to be assessed at a suitable level of reflection. It is here that the most important distinction arises. This is a distinction, not between fact and value, but between science and ethics. In the case of a true scientific judgment, Williams argues, there is some hope that people will come to accept it *because* it is true, where this is an explanatory 'because' which can be exercised at the relevant level of reflection. In the case of a true moral judgment there is no such hope. In the latter case, if people do come to accept the judgment, then the explanation for that fact, at this level of reflection, will have to be in quite other terms: historical, sociological, or anthropological, say. Indeed the explanation may not even be able to incorporate any such judgment itself. This is because continued exercise of such a judgment, again at this level of reflection, may no longer even be possible.<sup>8</sup> (I shall come back to this point later.)

The fact is that there are different ways in which human beings can live. They need to live in *some* sort of social world, but, as history amply demonstrates, there is no one sort of social world in which they need to live. Since the thick ethical concepts that people possess, and therefore the moral judgments that they accept, help to constitute the

<sup>6</sup> E.g. op. cit. note 2, 48–51.

<sup>7</sup> Op. cit. note 1, 140–5.

<sup>8</sup> Op. cit. note 1, Chapter 8.

social world in which they live, no explanation for their possessing those concepts or accepting those judgments can be metaphysically more ambitious than an explanation for their living in that world.

This suggests a kind of relativism, between societies if not between individuals. The threat of irresoluble moral disagreement has not been exorcized though it has, perhaps, been re-located. Williams acknowledges this. But what follows? Again Williams urges caution. ‘Vulgar’ relativism, according to which what follows is itself something moral—that every society should be tolerant towards the values of every other—is incoherent.<sup>9</sup> It is incoherent because the moral conclusion here is supposed to be at the same time ‘meta-moral’, that is to say it is supposed to be a conclusion not just *within* morality but *about* morality. But these conflict: as a conclusion about morality it is not itself supposed to have the relativistic attachment to a particular society that a conclusion within morality must have. In any case, no ‘meta-moral’ view can stop us finding the values of another society abhorrent and, where those values impinge on us, trying to combat them.<sup>10</sup> On the other hand, Williams’s conception must leave us dissatisfied with the blank thought, ‘We are right, and everyone else is wrong.’ So what follows, according to Williams, is what he calls ‘relativism of distance’: it is only if a society is sufficiently ‘close’ to ours, in a metaphorical sense that Williams explains, that talk of ‘right’ and ‘wrong’, with respect to their ethical outlook, need apply.<sup>11</sup> This conclusion may not look very substantial. But again there is the contrast with science. A scientific outlook, however ‘distant’ the society to which it belongs, must be either right or wrong.

Such, then, is Williams’s non-objectivism about ethics. But to what extent does he think his conclusions are disturbing? Certainly to some extent. For insofar as moral judgments seem to enjoy the objectivity of scientific judgments, his conclusions show that ethics is not all it seems. To come to terms with this we need *confidence*, that is confidence in our moral judgments. How, if at all, we achieve this will itself depend on the sort of society in which we live. Moreover, as Williams makes clear, some ways of achieving it must be resisted: confidence is a good, but it is not a supreme good.<sup>12</sup>

<sup>9</sup> Op. cit. note 2, 34–5, and op. cit. note 1, 159.

<sup>10</sup> Cf. ‘Subjectivism and Toleration’, in A. Phillips Griffiths (ed.), *A. J. Ayer: Memorial Essays* (Cambridge: Cambridge University Press).

<sup>11</sup> Op. cit. note 1, 156–67.

<sup>12</sup> Op. cit. note 1, 169–71 and 199–201; and ‘Replies’, in J.E.J. Altham and Ross Harrison (eds.), *World, Mind, and Ethics: Essays on the Ethical Philosophy of Bernard Williams*, 205–10.

## Williams on Ethics, Knowledge, and Reflection

2. There remains the question of what sort of society we actually want to live in. Here Williams has been an implacable opponent of utilitarianism.<sup>13</sup> In some of his most brilliant and most subtle writing, he has exposed the basic flaws in utilitarian thinking: not only must we acknowledge values other than happiness, the sole value that utilitarians acknowledge, we must understand happiness itself in a more realistic and textured way than utilitarians do. Since we are in any case encouraged by utilitarians to think in non-utilitarian terms (one of the many symptoms, Williams thinks, of a disquieting utilitarian indifference to social transparency), he is inclined to give utilitarianism the shortest possible shrift.

In fact he sees utilitarianism as just one manifestation of a particular style of ethical thought, pervasive in the modern world and, in Williams's view, deeply pernicious. He sometimes reserves the term 'morality' for this way of thinking.<sup>14</sup> Some of its other manifestations, notably Kantianism, are in other ways diametrically opposed to utilitarianism. But what they have in common are, very roughly: a conception of the purely voluntary, applicable to individual acts; a conception of ethics as fundamentally concerned with distinctions among purely voluntary acts—between those which ought to be performed, those which ought not to be performed, and those, if any, which fall into neither of these categories; the belief that the category into which an act falls is a matter of inescapable moral obligations, eclipsing all other considerations; and a commitment to such attendant notions as those of *guilt*, *blame*, and *responsibility*.<sup>15</sup> Williams rejects almost everything in this way of thinking. For instance he vehemently resents the pretensions to supremacy in the idea of a moral obligation. For Williams, there are many different kinds of motivation: moral motivation is just one, and it must take its place in our lives alongside all the others. But above all Williams wants to challenge the idea that moral notions are through and through *pure*. Guilt, blame, responsibility, and the rest are supposed to be uncontaminated by anything that does not relate back to the purely voluntary: there is supposed to be no such thing as 'moral luck'. This purity, Williams argues, is an illusion, except at a level which those who think in this way would themselves regard as hopelessly superficial. It is not that Williams believes there *is* such a

<sup>13</sup> 'Utilitarianism', in op. cit. note 2; and *A Critique of Utilitarianism*, in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973).

<sup>14</sup> See above, note 1.

<sup>15</sup> Op. cit. note 1, Ch. 10.

thing as moral luck, thought of in this way. His point is rather that we do better not to think in this way.<sup>16</sup>

In fact we do better, he urges, to turn to various ideas from antiquity. Among the most important of these are the ancient Greek ideas of *selfhood*, *freedom*, and *shame*. Williams argues that we have not, contrary to popular belief, advanced from these to a more refined conception. On the contrary, to the extent that we have not retained these ideas, it is often because we have adopted something much less helpful. Particularly instructive, in this respect, is the comparison of the ancient Greek idea of *shame* with the modern idea of *guilt*.<sup>17</sup>

Williams does not of course think that we can re-create the ancient Greek world: the modern world is irreversibly unlike anything that has gone before. But he does think that by comparing and contrasting ancient Greek thought with our own, we can come to understand the ancient Greeks better, we can come to understand ourselves better, and we can come to see that something of what we need they have to offer.

In particular we can learn from them about the possibility of (secular) meaning for an individual human life. Williams believes passionately in this possibility, both in the sense that he believes the possibility to be a real one (still) and in the sense that he is committed to its importance. There are few places in his writing where he explicitly expresses these beliefs.<sup>18</sup> But they inform almost all of it.

3. I now want to turn back to the idea that the exercise of a certain kind of judgment, at a certain level of reflection, may prove impossible, even though some judgments of that kind are true. Williams, as I have said, accepts this idea. But he goes a good deal further. He thinks there can be cases in which, after reflection, people find themselves unable to make judgments of a certain kind even though (a) what they are reflecting on are, precisely, certain true judgments of that kind, (b) these judgments are judgments that they themselves once made, (c) these judgments had the right sort of connection with what made them true to constitute knowledge, and (d) the people can *see* that these judgments had the right sort of connection with what made them true to constitute knowledge. This possibility is encapsulated in what Williams famously calls ‘the notably un-Socratic conclusion that... *reflection can destroy knowledge*’.<sup>19</sup>

<sup>16</sup> ‘Moral Luck’, reprinted in op. cit. note 4.

<sup>17</sup> *Shame and Necessity* (Oxford: University of California Press, 1993), *passim*.

<sup>18</sup> One exception is op. cit. note 1, 201–2.

<sup>19</sup> Op. cit. note 1, 148, his emphasis.

## Williams on Ethics, Knowledge, and Reflection

That reflection can destroy knowledge is one of Williams's most notorious doctrines. Even those who are broadly sympathetic to the metaphysical picture of ethical thought that leads Williams to espouse it are liable to say that this way of presenting the picture is needlessly provocative or needlessly paradoxical. There are many others who reject the picture altogether.<sup>20</sup> I want to devote the rest of this essay to a partial defence of Williams's doctrine.

Note first that there is a completely unremarkable sense in which reflection can destroy knowledge. A juggler, reflecting on what he is doing and on how he is doing it, may no longer know how to proceed.<sup>21</sup> The knowledge here is practical. Its exercise characteristically involves a complex and delicate physical performance. Furthermore, the destructive power of the reflection is highly contingent. There is no constitutional bar on the juggler retaining his knowledge while he is reflecting on it. But it is important that neither of these features was necessary to the example, simply *qua* example of how reflection can destroy knowledge in some unremarkable sense. There are also examples where the knowledge

<sup>20</sup> For some dissent see: Simon Blackburn, 'Making Ends Meet: A Discussion of *Ethics and the Limits of Philosophy*', in *Philosophical Books* 27 (1986), 200; John McDowell, 'Critical Notice of *Ethics and the Limits of Philosophy*', in *Mind* 95 (1986), 383; Hilary Putnam, *Renewing Philosophy* (Cambridge, Mass.: Harvard University Press, 1992), 104–5; and J. E. J. Altham, 'Reflection and Confidence', in J.E.J. Altham and Ross Harrison (eds.), *op. cit.* note 12. I myself have argued against Williams's doctrine: see A. W. Moore, 'Can Reflection Destroy Knowledge?', in *Ratio* (New Series) 4 (1991). I am now much more sympathetic to the doctrine, as this essay is an attempt to show. I now think, very roughly, that in *op. cit.* I rely on too crude a conception of how truth relates to the distinction between propositional knowledge and practical knowledge—a conception that I try to correct in A. W. Moore, *Points of View* (Oxford: Oxford University Press, 1997), Ch. 8—and, more significantly, that I rely on too crude a conception of what is involved in explaining how a judgment constitutes knowledge. This too is a conception that I try to correct in *Points of View*. I there distinguish between endorsement and indirect endorsement (12 and 15–16). In 'Can Reflection Destroy Knowledge?' I take for granted that the explanation of how a judgment constitutes knowledge will have to involve, or at least allow for, its endorsement; but in fact, it need only allow for (and involve) its indirect endorsement. Some of this, I hope, will be clarified below. Note, however, that I do voice some criticisms of Williams's doctrine in *Points of View*: see 187–8. And I am still inclined to voice them. What I go on to say in this essay, though it provides support for Williams's doctrine, does not provide unqualified support for it.

<sup>21</sup> Cf. *op. cit.* note 1, 167–8.

is propositional (the canonical evidence for someone's having it is his or her making a judgment). Thus a flustered examinee, reflecting on why she is about to write down 'Early Gothic', may no longer know that that is the correct answer. So too there are examples where the knowledge's vulnerability to reflection is constitutional. Thus there is the man who knows how to avoid facing up to some unpleasant truth about himself—until he reflects on the fact that that is what he is doing.

What sets Williams's doctrine apart from *any* of these examples, and therewith from an innocuous understanding of the claim that reflection can destroy knowledge, is that he thinks there are examples of how reflection can destroy knowledge which are simultaneously of *both* these kinds: that is to say, the knowledge in question is propositional *and* its vulnerability to reflection is, in some sense, constitutional. The sense in which its vulnerability to reflection is constitutional is not that reflection is guaranteed to destroy it, but rather that, if reflection does destroy it, then the knowledge is guaranteed to remain irrecoverable for as long as the reflection persists. What reflection does in such cases, according to Williams, is to undermine the conceptual apparatus required even to think in the relevant terms. The people engaging in the reflection can no longer make judgments of the kind that constitute the knowledge, although they can still have enough of a grasp on judgments of that kind, from without, to *see* that they constitute knowledge. They may eventually recover the knowledge: various social forces may bring this about. But such forces must also prevent them from thinking, at the relevant level, about what they are up to. They will never recover the knowledge in the full light of reflection.<sup>22</sup>

The feature of the knowledge which on Williams's view allows for this possibility—the possibility that reflection should destroy it and suppress it in the way described—is that it involves thick ethical concepts. It is because these concepts are world-guided that they are capable of expressing knowledge. It is because they are action-guiding that the knowledge has the special vulnerability to reflection that it has. Reflection can disturb or challenge people's commitment to the evaluative outlook that gives a given thick ethical concept its point. People can find that they no longer have that outlook, and thus that they can no longer exercise that concept. But they can still *understand* the outlook, and indeed the concept, just as an anthropologist or a historian can understand an evaluative out-

<sup>22</sup> Op. cit. note 1, 163–4 and 168.

## Williams on Ethics, Knowledge, and Reflection

look which he or she does not share. This explains why they can still see their erstwhile knowledge as knowledge.<sup>23</sup>

Of course, this would be altogether less mysterious if the concept simply conjoined, separably, a factual component and an evaluative component, the first of these making it world-guided and the second making it action-guiding. Some moral philosophers have thought that thick ethical concepts are of precisely this kind. They have thought that the concept of *cowardice*, for instance, has a factual component whereby it applies only to those who have a certain attitude to danger, and an evaluative component whereby whoever exercises the concept registers his or her disapproval of that attitude.<sup>24</sup> If this were right, and if what reflection did were simply to undermine the evaluative components in thick ethical concepts, then there would be a clear sense in which what had *strictly speaking* been known, in a way that had found expression through judgments involving the concepts, was still available to be known, in a way that could find expression through judgments involving purely factual counterparts of the concepts. The claim that reflection can destroy knowledge, in the sense intended by Williams, *would* then be a needlessly paradoxical way of putting something innocuous.

But Williams is adamant that this is an unworkable account of thick ethical concepts. He repudiates the idea that these concepts have, or even can have, purely factual counterparts; that they can be seen, from outside their governing outlooks, 'simply as [devices] for dividing up in a rather strange way certain neutral features of the world'.<sup>25</sup>

In this respect, one of his analogies is infelicitous. In this analogy there are conventions whereby certain items of vocabulary can be used only by certain people, or only by certain people in certain circumstances.<sup>26</sup> Granted such conventions, a situation may arise in which one person can see that assertions made by another person are expressions of knowledge even though she cannot herself make assertions of that kind: they are couched in the relevant vocabulary and she is not herself a member of the privileged class. Again, a situation may arise in which somebody can see that assertions which he himself made in the past were expressions of knowledge even though he cannot any longer make assertions of that kind: he was,

<sup>23</sup> Cf. the reference to the 'sympathetic but nonidentified observer' in op. cit. note 1, 219, note 11. Cf. also op. cit. note 12, 206–7.

<sup>24</sup> Cf. R.M. Hare, *Freedom and Reason* (Oxford: Oxford University Press, 1963), Ch. 10, §1.

<sup>25</sup> Op. cit. note 1, 141–2. (The quoted material occurs on 142.)

<sup>26</sup> Op. cit. note 1, 143–4.

but no longer is, a member of the privileged class. It is clear how the analogy is supposed to work. Indeed, on some ways of individuating ‘judgments’ and ‘kinds’ of judgment, it provides us with examples that can be described in the very same terms as the examples that Williams is interested in. Thus, in the second of the situations just envisaged, what the man is reflecting on are judgments which he himself once made, and which he can see constituted knowledge, even though he now finds himself unable to make judgments of the same kind. The analogy is nevertheless infelicitous because, in it, the message and the medium are easily separated. The knowledge is there to be shared by members and non-members of the privileged class alike. And it can be expressed, where necessary, by the use of different vocabulary. The analogy precisely fails to illustrate that indissolubility of the factual and the evaluative which Williams takes to characterize thick ethical concepts and which gives his view so much of its force.<sup>27</sup>

4. How then are we to view Williams’s doctrine that reflection can destroy knowledge? Interestingly, the materials for one powerful objection to the doctrine are to be found elsewhere in Williams’s own work, in his book on Descartes.<sup>28</sup> In that book Williams develops an argument from something which he says appears ‘basic to the notion of knowledge itself’.<sup>29</sup> ‘If knowledge is what it claims to be,’ he writes, ‘then it is knowledge of... what is there *anyway*.’<sup>30</sup> It follows that if two people both have knowledge, but their judgments differ, to the point of being uncombinable, perhaps even to the point of looking incompatible—something which can certainly happen if the judgments involve thick ethical concepts—then ‘there must be some coherent way of understanding why these [judgments] differ, and how they are related to one another.’<sup>31</sup> There is nothing in Williams’s work in ethics which is in immediate tension with this. In fact, he gives us an indication of what such an account will look like, in the case where the judgments concerned do involve thick ethical concepts. ‘An *explanation* of those local judgments,’ he writes, ‘and of the conceptual differences between societies... will

<sup>27</sup> Cf. Warren Quinn, ‘Reflection and the Loss of Moral Knowledge: Williams on Objectivity’, reprinted in his *Morality and Action* (Cambridge: Cambridge University Press, 1993), 143.

<sup>28</sup> *Descartes: The Project of Pure Enquiry* (Harmondsworth: Penguin, 1978).

<sup>29</sup> Op. cit. note 28, 65.

<sup>30</sup> Op. cit. note 28, 64, his emphasis.

<sup>31</sup> Op. cit. note 28.

## Williams on Ethics, Knowledge, and Reflection

invoke [capacities]... involved in finding [a] way around a social world.<sup>32</sup> But if the account is to satisfy the demands placed on it by the fact that what we have here is *knowledge*, then it must do more than provide a coherent story about how these judgments can be different cohabitants in the same world. It must provide a coherent story about how they can do that *and constitute knowledge*. It must represent the two people concerned as having, not just two different points of view *in* the same world, but two different points of view *on* the same world. (Otherwise there would not be any puzzle, or at any rate not the same puzzle, about how their judgments come to differ in the way that they do.) And granted that it cannot do this by distilling evaluative components from the various thick ethical concepts involved, and employing purely factual counterparts of those concepts instead, then it surely cannot do it at all—it surely cannot locate the two people together in a world which is how their knowledge severally reveals it to be—if, at the same time, it somehow precludes our fully acknowledging *what each of them knows*. But, the objection runs, this account, which, according to the argument, must be available, and which the two people themselves are not debarred from giving, involves standing back from their judgments.<sup>33</sup> In other words it involves reflection. So it must be possible (for each of the two people, though not just for them) to keep in focus what they know even when standing back and reflecting on it. In particular, their knowledge must be able to withstand the reflection that is necessary to justify its entitlement to be called knowledge in the first place. Reflection cannot destroy knowledge, in the sense intended by Williams.

To be sure, Williams could resist this objection by saying that the argument he gives in the book on Descartes is concerned with a narrower conception of knowledge than that which he deploys in his work on ethics. More specifically, he could say that ethical knowledge lies outside the ambit of the main premise of the argument he gives in the book on Descartes because it is not ‘what it claims to be’. Indeed he does write at one point that ‘ethical thought has no chance of being everything it seems... [it] will never entirely appear as what it is, and can never fully manifest the fact that it rests in human dispositions.’<sup>34</sup> But of course, if Williams did resist the objection in this way, he would have a considerable price to pay. Either this would greatly lessen the force of his claim that judgments involving thick ethical concepts can constitute knowledge or it would greatly lessen

<sup>32</sup> Op. cit. note 1, 150, his emphasis.

<sup>33</sup> Op. cit. note 28, 64–5.

<sup>34</sup> Op. cit. note 1, 199–200.

the force of the argument he gives in the book on Descartes. He gives that argument as an argument for the possibility of what he calls an 'absolute conception of reality'.<sup>35</sup> But if it were a real question whether any given item of knowledge were 'what it claims to be', then he would have his work cut out to show that, for example, knowledge about the colours of things fared any better in this respect than ethical knowledge,<sup>36</sup> or indeed, more pertinently, that any knowledge that was not already part of an absolute conception of reality fared any better in this respect than ethical knowledge, something which he would have to do if the assumption that *some* knowledge is 'what it claims to be' were not to be straightforwardly question-begging.

A far better alternative for Williams, it seems to me, would be to fasten on this crucial idea of our fully acknowledging what each of the two people knows. According to the objection, the account of how their judgments are related to one another cannot succeed if it precludes our doing this (fully acknowledging what each of them knows). But what need this involve? Certainly it must involve our seeing that their judgments constitute knowledge. But need it involve our making any judgments of the same kind? Need it involve our sharing their knowledge? We are in danger of begging the question against Williams if we say so. For if sharing their knowledge means exercising the same thick ethical concepts as they do, then precisely what Williams has given us reason to think is that we can see their knowledge as knowledge without sharing it. We can learn enough, about them, about their culture, about the history of their culture, and perhaps also about *Homo sapiens*, to be able to see how their use of various thick ethical concepts enables each of them to live, with others, in a particular social world; and we can learn enough about the particular circumstances in which they find themselves to be able to see that these circumstances warrant their making these judgments, using those very concepts. But although what we thereby learn will in some (non-trivial) sense entail the truth of their judgments,<sup>37</sup> *we* need not use any of the relevant concepts

<sup>35</sup> Op. cit. note 28, 65.

<sup>36</sup> Cf. op. cit. note 28, 237–46 and op. cit. note 1, 149–50.

<sup>37</sup> Precisely what this sense is connects with how we are to strike a balance that Clifford Geertz refers to in a passage that Williams quotes in 'Saint-Just's Illusion', in his op. cit. note 3, 143. In this passage Geertz talks about the following aim: 'to produce an interpretation of the way a people lives which is neither imprisoned within their mental horizons, an ethnography of witchcraft as written by a witch, nor systematically deaf to the distinctive tonalities of their existence, an ethnography of witchcraft as written by a geometer'.

## Williams on Ethics, Knowledge, and Reflection

either in learning it or in subsequently recounting what we have learned. It is as if there were a distinction between using concepts and mentioning them, akin to the distinction between using words and mentioning them,<sup>38</sup> and it is as if we need only ever mention any of those concepts. In *this* respect Williams's analogy concerning the conventions that restrict the use of certain items of vocabulary is a helpful one. The man who is unable to repeat assertions that he once made, because he no longer belongs to the class of people entitled to use the vocabulary in which the assertions were couched, may yet *mention* some of that vocabulary, and thereby talk *about* the assertions, as part of claiming (for instance) that they were true.

Nevertheless I think there is an analogy that is in many respects more helpful. It involves tense. Consider judgments to the effect that no-one has ever walked on the moon. Of these we can say: that true judgments of this kind have been made in the past; that these judgments have had the right sort of connection with what makes them true to constitute knowledge; that we, now, reflecting on these judgments, can *see* that they have had the right sort of connection with what makes them true to constitute knowledge; but that we are no longer in a position to make judgments of this kind ourselves. Furthermore, the reason why we are no longer in a position to make judgments of this kind ourselves is that the temporal point of view from which the original judgments were made is no longer ours. It is here, I think, that the chief virtue of the analogy lies. For what Williams is envisaging are situations in which we are no longer in a position to make judgments of a certain kind because the *ethical* point of view from which judgments of that kind were originally made is no longer ours. In both cases our change of point of view makes it impossible for us to continue to use various concepts as we did before. And this impossibility cuts deeper than the mere impossibility of our continuing to use various items of vocabulary as we did before.

The analogy is not unproblematic however. No sooner does one probe it than one sees complications. In particular a full discussion of these issues would need to include a detailed account of something that I have so far completely glided over, except for one passing reference: namely, the individuation of 'kinds' of judgment. Two assertions of the sentence, 'No-one has ever walked on the moon,' would standardly be recognized by philosophers as two

<sup>38</sup> For example, I use the word 'cats' when I say, 'Cats have four legs.' I mention it when I say, "'Cats" has four letters.'

tokens of the same type; and various questions would then arise, many of them purely terminological, about whether two different propositions had been expressed or one proposition twice, about whether the two assertions had the same content or not, and so forth. Questions about what kind or kinds of judgment had been made would clearly fall into this category. And how they were settled would dictate how the various claims that we might want to make, about what we could do and can no longer do, should be formulated. Roughly speaking, I have been treating a 'kind' of judgment as a repeatable type in this essay. But granted this, the analogy that I have offered may seem weak. For if that is what a kind of judgment is, then it seems that we *are* still in a position to make judgments of the kind that people made when they judged that no-one had ever walked on the moon. What we are not in a position to do is to make them truly. Nor is this hair-splitting. For Williams is concerned with situations in which we have lost the concepts even to be able to produce tokens of the same type as we did before. (That is, we have lost the concepts to be able to do this *in propria persona*. No doubt we can still do it as part of playing a rôle. But that is importantly different from speaking falsely.) This connects with the idea that our change of point of view has made it impossible for us to continue to use certain concepts as we did before. In the temporal case, these concepts are tensed concepts, such as the concept of what has happened, and what is impossible is for us to continue to use them to the same *effect* as we did before. In the ethical case, the concepts are thick ethical concepts, and what is impossible is for us to continue to use them at all. That there should be this difference between the two cases is not in itself so serious. But what is serious is something that the difference highlights, namely the indexicality of tense, which suggests that the new analogy shares the principal defect of Williams's own analogy concerning the restricted use of various items of vocabulary. For is not the knowledge that was once expressed by assertions of the sentence 'No-one has ever walked on the moon' still available to be expressed in different terms, say by assertions of the sentence 'Before 1969, no-one had ever walked on the moon'?

Well *is* it? Certainly this particular sentence does not fit the bill. The knowledge expressed in pre-1969 assertions of the sentence 'No-one has ever walked on the moon' did not involve any reference, explicit or implicit, to 1969. And it is a real question whether any other sentence we could assert now would fit the bill either. Some philosophers would insist not, precisely on the grounds that

## Williams on Ethics, Knowledge, and Reflection

the knowledge itself was tensed.<sup>39</sup> This raises some very large questions. But there is clearly a sense—an interesting sense, with no counterpart in Williams's own analogy—in which what people used to know, about whether anyone had walked on the moon, is no longer available to be known. Their knowledge *can no longer be shared*—though we can know enough to entail it, in some (non-trivial) sense, just as we can in the ethical case, if what I said earlier is correct. True, this is once again partly a matter of terminology. In particular, it is a matter of terminology how 'things known' are to be individuated. None the less, we can see how knowledge from different temporal points of view has some of the supposedly problematic structure that Williams takes ethical knowledge to have. So I think the new analogy does give us a better grip on that structure than Williams's own analogy does—to the extent, moreover, of being able to assuage much of the disquiet that Williams's critics have felt about the structure.

5. This is not, however, the only disquiet that his critics have felt. There remains the crucial question of how, in the ethical case, reflection precludes re-adoption of the abandoned point of view. That is, there remains the question of how reflection achieves what is achieved in the temporal case by the sheer passage of time. It is beyond the scope of this essay to mount a full discussion of that question. This is one reason why I said earlier that my defence of Williams's doctrine would be partial. (I am in any case much more sceptical about this aspect of his doctrine.) Even so, I should like to finish by deflecting one or two of the worries that arise here too.

Williams himself, of course, has much to say on this matter.<sup>40</sup> At one point he indicates that, when reflection destroys ethical knowledge, what it does is to lead those who have the knowledge 'to give it up, lose hold on it, or simply drift away from it, as modern societies in the past two centuries or less have, for instance, done one or more of those things in relation to the concept of *chastity*'.<sup>41</sup> The

<sup>39</sup> Cf. John Perry, 'The Problem of the Essential Indexical', reprinted in his *The Problem of the Essential Indexical and Other Essays* (New York: Oxford University Press, 1993), 45–6. The issues that arise here are connected with how much is demanded by the 'tantamount' in a principle which Williams discusses in *op. cit.* note 1, 143, namely that 'A cannot correctly say that B speaks truly in uttering *S* unless A could also say something tantamount to *S*'.

<sup>40</sup> See *esp.* *op. cit.* note 1, 167–71. See also *op. cit.* note 37 and *op. cit.* note 12, 205–10.

<sup>41</sup> *Op. cit.* note 12, 207, his emphasis.

destruction of ethical knowledge by reflection is a social phenomenon. It occurs when various advances in understanding, for instance advances in the understanding of human nature or of history, undermine some of the beliefs that give a thick ethical concept its point, thereby causing the concept to be abandoned—and, for as long as those advances in understanding are operative, preventing it from being recovered. The way in which the beliefs are undermined is simply by being shown to be false. But the concept itself is not shown to be false. ('True' and 'false' do not apply to concepts.) Rather, the concept is shown not to play one or more of the rôles that it was thought to play. The effect is that people no longer want to think in those terms. But the fact remains that, when people did want to think in those terms, they were able to put the concept to use in making certain true judgments, judgments having the right sort of connection with what made them true to constitute knowledge.

Is this coherent? Well, here is an analogy. Consider the concepts of Euclidean geometry. These can be put to use in making certain true judgments about Euclidean space (for instance, that the sum of the angles in a triangle is equal to two right angles). And these judgments have the right sort of connection with what makes them true to constitute knowledge (they can be rigorously derived from axioms that determine their subject matter). Even so, we now know something that for centuries we did not know, namely that physical space is not Euclidean: the concepts of Euclidean geometry are not suitable for describing it. It is possible to imagine—is it not?—that, in due course, those concepts will simply be abandoned.

Someone might reply, 'It is. But to be abandoned is one thing. To be irrecoverable is another. Even if the concepts of Euclidean geometry *are* eventually abandoned, on the grounds that they lack the application which for centuries they were thought to have, it will still be possible to bring them back into use, if only into purely theoretical use. The knowledge that the sum of the angles in a (Euclidean) triangle is equal to two right angles may become otiose, but it will never (thereby) be *destroyed*.'

True. But all this shows is that the analogy is limited. Precisely what distinguishes the concepts of Euclidean geometry from thick ethical concepts, in this respect, is that the former can indeed be put to 'purely theoretical use'. Thick ethical concepts cannot be used without incurring certain practical commitments. They are *action-guiding*. It remains entirely possible that, if reflection causes one of them to be abandoned, then it will also prevent it from being recovered. After all, we cannot make use of a thick ethical concept unless

## Williams on Ethics, Knowledge, and Reflection

we have a certain confidence in it. (I talked at the end of §1 about the confidence that we need to have in our moral judgments.) But how can we have confidence in a concept that reflection has once caused us to abandon? Suppose, for instance, that we originally use the concept to draw a distinction that we think is correlated with some genetic distinction. And suppose that what leads us to abandon it is the discovery that there is no such genetic distinction to be drawn. How can we then regain the confidence necessary to start applying the concept again? What are we to make of the distinction that we originally drew by means of it? How can we see the drawing of that distinction as anything other than completely unprincipled?

To be sure, even if these rhetorical questions show that we cannot recover the concept while we are still reflectively aware of what led us to abandon it, it remains to be shown that this 'cannot' is sufficiently demanding to justify our regarding the vulnerability to reflection of the ethical knowledge in question as 'constitutional'. (Does a psychological impossibility suffice? Does a sociological impossibility? Or does nothing less than inconceivability suffice?) This in turn may degenerate into a tiresome terminological quibble. But it need not. There are substantive issues, for instance, about how far the vulnerability of ethical knowledge in general to reflection is itself something that we can (and ought to) try to lessen.

There is another way in which reflection can destroy knowledge, very different in kind from anything that we have considered hitherto, which raises similar questions about the sort of power that reflection exerts. I have in mind the capacity of reflection to give rise to Cartesian doubts.<sup>42</sup> Thus consider someone who knows perfectly well that there is a table in front of him. He may, after reflection, come to have no more than a shaken belief that there is a table in front of him, a belief which no longer counts as knowledge nor, for as long as he reflects, can be converted back into knowledge. Here, as in cases of the sort that Williams is envisaging, there is a loss of confidence. And here, as in cases of the sort that Williams is envisaging, we can ask what sort of impossibility it is for the person concerned to recapture his confidence, in the full light of reflection. Is it merely a psychological impossibility? Or has the reflection created a demand for justification that cannot be met, with the result that no reflective state that he can now get into is properly to *count* as a state of confidence? If it would not be absurd to accede to the

<sup>42</sup> See 'Can Reflection Destroy Knowledge?', *op. cit.* note 20, Appendix, and *Points of View*, *op. cit.* note 20, 235–6.

**A. W. Moore**

second of these—and I think it would not be—then neither would it be absurd to say that the impossibility in cases of the sort that Williams is envisaging is similarly constitutive. And this would certainly make that impossibility demanding enough for his purposes.

As I have already indicated, Williams himself has much more to say about these issues. Suffice to conclude that, like all else in Williams, it withstands a great deal of reflection.

*St. Hugh's College Oxford*