

## Nicholas Sturgeon

### MORAL EXPLANATIONS

**T**HERE IS one argument for moral skepticism that I respect even though I remain unconvinced. It has sometimes been called the argument from moral diversity or relativity, but that is somewhat misleading, for the problem arises not from the diversity of moral views, but from the apparent difficulty of settling moral disagreements, or even of knowing what would be required to settle them, a difficulty thought to be noticeably greater than any found in settling disagreements that arise in, for example, the sciences. This provides an argument for moral skepticism because one obviously possible explanation of our difficulty in settling moral disagreements is that they are really un settleable, that there is no way of justifying one rather than another competing view on these issues; and a possible further explanation for the un settleability of moral disagreements, in turn, is moral nihilism, the view that on these issue there just is no fact of the matter, that the impossibility of discovering and establishing moral truths is due to there not being any.

I am, as I say, unconvinced: partly because I think this argument exaggerates the difficulty we actually find in settling moral disagreements, partly because there are alternative explanations to be considered for the difficulty we do find. Under the latter heading, for example, it certainly matters to what extent moral disagreements depend on disagreements about other questions which, however disputed they may

be, are nevertheless regarded as having objective answers, questions such as which, if any, religion is true, which account of human psychology, which theory of human society. And it also matters to what extent consideration of moral questions is in practice skewed by distorting factors such as personal interest and social ideology. These are large issues. Although it is possible to say some useful things to put them in perspective,<sup>1</sup> it appears impossible to settle them quickly or in any a priori way. Consideration of them is likely to have to be piecemeal and, in the short run at least, frustratingly indecisive.

These large issues are not my topic here. But I mention them, and the difficulty of settling them, to show why it is natural that moral skeptics have hoped to find some quicker way of establishing their thesis. I doubt that any exist, but some have of course been proposed. Verificationist attacks on ethics should no doubt be seen in this light, and J. L. Mackie's recent "argument from queerness" is a clear instance.<sup>2</sup> The quicker response on which I shall concentrate, however, is neither of these, but instead an argument by Gilbert Harman designed to bring out the "basic problem" about morality, which in his view is "its apparent immunity from observational testing" and "the seeming irrelevance of observational evidence."<sup>3</sup> The argument is that reference to moral facts appears unnecessary for the explanation of our moral observations and beliefs.

Harman's view, I should say at once, is not in the end a skeptical one, and he does not view the argument I shall discuss as a decisive defense of moral skepticism or moral nihilism. Someone else might easily so regard it, however. For Harman himself regards it as creating a strong *prima facie* case for skepticism and nihilism, strong enough to justify calling it "the problem with ethics."<sup>4</sup> And he believes it shows that the only recourse for someone who wishes to avoid moral skepticism is to find defensible reductive definitions for ethical terms; so skepticism would be the obvious conclusion to draw for anyone who doubted the possibility of such definitions. I believe, however, that Harman is mistaken on both counts. I shall show that his argument for skepticism either rests on claims that most people would find quite implausible (and so cannot be what constitutes, for them, the problem with ethics); or else becomes just the application to ethics of a familiar *general* skeptical strategy, one which, if it works for ethics, will work equally well for unobservable theoretical entities, or for other minds, or for an external world (and so, again, can hardly be what constitutes the distinctive problem with ethics). I have argued elsewhere,<sup>5</sup> moreover, that one can in any case be a moral realist, and indeed an ethical naturalist, without believing that we are now or ever will be in possession of reductive naturalistic definitions for ethical terms.

## I. The problem with ethics

Moral theories are often tested in thought experiments, against imagined examples; and, as Harman notes, trained researchers often test scientific theories in the same way. The problem, though, is that scientific theories can also be tested against the world, by observations or real experiments; and, Harman asks, "can moral principles be tested in the same way, out in the world?"

This would not be a very interesting or impressive challenge, of course, if it were merely

a resurrection of standard verificationist worries about whether moral assertions and theories have any testable empirical implications, implications storable in some relatively austere "observational" vocabulary. One problem with that form of the challenge, as Harman points out, is that there are no "pure" observations, and in consequence no purely observational vocabulary either. But there is also a deeper problem that Harman does not mention, one that remains even if we shelve worries about "pure" observations and, at least for the sake of argument, grant the verificationist his observational language, pretty much as it was usually conceived: that is, as lacking at the very least any obviously theoretical terminology from any recognized science, and of course as lacking any moral terminology. For then the difficulty is that moral principles fare just as well (or just as badly) against the verificationist challenge as do typical scientific principles. For it is by now a familiar point about scientific principles—principles such as Newton's law of universal gravitation or Darwin's theory of evolution—that they are entirely devoid of empirical implications when considered in isolation.<sup>6</sup> We do of course base observational predictions on such theories and so test them against experience, but that is because we do *not* consider them in isolation. For we can derive these predictions only by relying at the same time on a large background of additional assumptions, many of which are equally theoretical and equally incapable of being tested in isolation.

A less familiar point, because less often spelled out, is that the relation of moral principles to observation is similar in *both* these respects. Candidate moral principles—for example, that an action is wrong just in case there is something else the agent could have done that would have produced a greater balance of pleasure over pain—lack empirical implications when considered in isolation. But it is easy to derive empirical consequences from them, and thus to test them against experience, if we

allow ourselves, as we do in the scientific case, to rely on a background of other assumptions of comparable status. Thus, if we conjoin the act-utilitarian principle I just cited with the further view, also untestable in isolation, that it is always wrong deliberately to kill a human being, we can deduce from these two premises together the consequence that deliberately killing a human being always produces a lesser balance of pleasure over pain than some available alternative act; and this claim is one any positivist would have conceded we know, in principle at least, how to test. If we found it to be false, moreover, then we would be forced by this empirical test to abandon at least one of the moral claims from which we derived it.

It might be thought a worrisome feature of this example, however, and a further opening for skepticism, that there could be controversy about which moral premise to abandon, and that we have not explained how our empirical test can provide an answer to this question. And this may be a problem. It should be a familiar problem, however, because the Duhemian commentary includes a precisely corresponding point about the scientific case: that if we are at all cautious in characterizing what we observe, then the requirement that our theories merely be *consistent* with observation is a very weak one. There are always many, perhaps indefinitely many, different mutually inconsistent ways to adjust our views to meet this constraint. Of course, in practice we are often confident of how to do it: if you are a freshman chemistry student, you do not conclude from your failure to obtain the predicted value in an experiment that it is all over for the atomic theory of gases. And the decision can be equally easy, one should note, in a moral case. Consider two examples. From the surprising moral thesis that Adolf Hitler was a morally admirable person, together with a modest piece of moral theory to the effect that no morally admirable person would, for example, instigate and oversee the degradation and death of millions of persons, one can derive

the testable consequence that Hitler did not do this. But he did, so we must give up one of our premises; and the choice of which to abandon is neither difficult nor controversial.

Or, to take a less monumental example, contrived around one of Harman's own, suppose you have been thinking yourself lucky enough to live in a neighborhood in which no one would do anything wrong, at least not in public; and that the modest piece of theory you accept, this time, is that malicious cruelty, just for the hell of it, is wrong. Then, as in Harman's example, "you round a corner and see a group of young hoodlums pour gasoline on a cat and ignite it." At this point, either your confidence in the neighborhood or your principle about cruelty has got to give way. But the choice is easy, if dispiriting, so easy as hardly to require thought. As Harman says, "You do not need to conclude that what they are doing is wrong; you do not need to figure anything out; you can see that it is wrong" (p. 4). But a skeptic can still wonder whether this practical confidence, or this "seeing," rests in either sort of case on anything more than deeply ingrained conventions of thought—respect for scientific experts, say, and for certain moral traditions—as opposed to anything answerable to the facts of the matter, any reliable strategy for getting it right about the world.

Now, Harman's challenge is interesting partly because it does not rest on these verificationist doubts about whether moral beliefs have observational implications, but even more because what it does rest on is a partial answer to the kind of general skepticism to which, as we have seen, reflection on the verificationist picture can lead. Many of our beliefs are justified, in Harman's view, by their providing or helping to provide a reasonable explanation of our observing what we do. It would be consistent with your failure, as a beginning student, to obtain the experimental result predicted by the gas laws, that the laws are mistaken. But a better explanation, in light of your inexperience and the

general success experts have had in confirming and applying these laws, is that you made some mistake in running the experiment. So our scientific beliefs can be justified by their explanatory role; and so too, in Harman's view, can mathematical beliefs and many commonsense beliefs about the world.

Not so, however, moral beliefs: they appear to have no such explanatory role. That is "the problem with ethics." Harman spells out his version of this contrast:

You need to make assumptions about certain physical facts to explain the occurrence of the observations that support a scientific theory, but you do not seem to need to make assumptions about any moral facts to explain the occurrence of the so-called moral observations I have been talking about. In the moral case, it would seem that you need only make assumptions about the psychology or moral sensibility of the person making the moral observation. (p. 6)

More precisely, and applied to his own example, it might be reasonable, in order to explain your judging that the hoodlums are wrong to set the cat on fire, to assume "that the children really are pouring gasoline on a cat and you are seeing them do it." But there is no

obvious reason to assume anything about "moral facts," such as that it is really wrong to set the cat on fire. . . . Indeed, an assumption about moral facts would seem to be totally irrelevant to the explanation of your making the judgment you make. It would seem that all we need assume is that you have certain more or less well articulated moral principles that are reflected in the judgments you make, based on your moral sensibility. (p. 7)

And Harman thinks that if we accept this conclusion, suitably generalized, then, subject to one possible qualification concerning reduction that

I have discussed elsewhere,<sup>7</sup> we must conclude that moral theories cannot be tested against the world as scientific theories can, and that we have no reason to believe that moral facts are part of the order of nature or that there is any moral knowledge (pp. 23, 35).

My own view is that Harman is quite wrong, not in thinking that the explanatory role of our beliefs is important to their justification, but in thinking that moral beliefs play no such role.<sup>8</sup> I shall have to say something about the initial plausibility of Harman's thesis as applied to his own example, but part of my reason for dissenting should be apparent from the other example I just gave. We find it easy (and so does Harman [p. 108]) to conclude from the evidence not just that Hitler was not morally admirable, but that he was morally depraved. But isn't it plausible that Hitler's moral depravity—the fact of his really having been morally depraved—forms part of a reasonable explanation of why we believe he was depraved? I think so, and I shall argue concerning this and other examples that moral beliefs very commonly play the explanatory role Harman denies them. Before I can press my case, however, I need to clear up several preliminary points about just what Harman is claiming and just how his argument is intended to work.

## II. Observation and explanation

(1) For there are several ways in which Harman's argument invites misunderstanding. One results from his focusing at the start on the question of whether there can be moral observations.<sup>9</sup> But this question turns out to be a side issue, in no way central to his argument that moral principles cannot be tested against the world. There are a couple of reasons for this, of which the more important<sup>10</sup> by far is that Harman does not really require of moral facts, if belief in them is to be justified, that they figure in the explanation of moral observations. It would be enough, on the one hand, if they were needed for the

explanation of moral beliefs that are not in any interesting sense observations. For example, Harman thinks belief in moral facts would be vindicated if they were needed to explain our drawing the moral conclusions we do when we reflect on hypothetical cases, but I think there is no illumination in calling these conclusions observations.<sup>11</sup> It would also be enough, on the other hand, if moral facts were needed for the explanation of what were clearly observations, but not moral observations. Harman thinks mathematical beliefs are justified, but he does not suggest that there are mathematical observations; it is rather that appeal to mathematical truths helps to explain why we make the physical observations we do (p. 10). Moral beliefs would surely be justified, too, if they played such a role, whether or not there are any moral observations.

So the claim is that moral facts are not needed to explain our having any of the moral beliefs we do, whether or not those beliefs are observations, and are equally unneeded to explain any of the observations we make, whether or not those observations are moral. In fact, Harman's view appears to be that moral facts aren't needed to explain anything at all: though it would perhaps be question-begging for him to begin with this strong a claim, since he grants that if there were any moral facts, then appeal to other moral facts—more general ones, for example—might be needed to explain them (p. 8). But he is certainly claiming, at the very least, that moral facts aren't needed to explain any nonmoral facts we have any reason to believe in.

(2) Other possible misunderstandings concern what is meant in asking whether reference to moral facts is *needed* to explain moral beliefs. One warning about this question I have dealt with in my discussion of reduction elsewhere;<sup>12</sup> but another, about what Harman is clearly *not* asking, and about what sort of answer I can attempt to defend to the question he is asking, I can spell out here. For Harman's question is clearly not just whether there is an

explanation of our moral beliefs that does not mention moral facts. Almost surely there is. Equally surely, however, there is an explanation of our commonsense nonmoral beliefs that does not mention an external world: one which cites only our sensory experience, for example, together with whatever needs to be said about our psychology to explain why with that history of experience we would form just the beliefs we do. Harman means to be asking a question that will lead to skepticism about moral facts, but not to skepticism about the existence of material bodies or about well-established scientific theories of the world.

Harman illustrates the kind of question he is asking, and the kind of answer he is seeking, with an example from physics that it will be useful to keep in mind. A physicist sees a vapor trail in a cloud chamber and thinks, "There goes a proton." What explains his thinking this? Partly, of course, his psychological set, which largely depends on his beliefs about the apparatus and all the theory he has learned; but partly also, perhaps, the hypothesis that "there really was a proton going through the cloud chamber, causing the vapor trail, which he saw as a proton." We will not need this latter assumption, however, "if his having made that observation could have been equally well explained by his psychological set alone, without the need for any assumption about a proton" (p. 6).<sup>13</sup> So for reference to moral facts to be *needed* in the explanation of our beliefs and observations, is for this reference to be required for an explanation that is somehow *better* than competing explanations. Correspondingly, reference to moral facts will be unnecessary to an explanation, in Harman's view, not just because we can find some explanation that does not appeal to them, but because *no* explanation that appeals to them is any better than some competing explanation that does not.

Now, fine discriminations among competing explanations of almost anything are likely to be difficult, controversial, and provisional. Fortunately, however, my discussion of Harman's

argument will not require any fine discriminations. This is because Harman's thesis, as we have seen, is not that moral explanations lose out by a small margin; nor is it that moral explanations, though sometimes initially promising, always turn out on further examination to be inferior to nonmoral ones. It is, rather, that reference to moral facts always looks, right from the start, to be "completely irrelevant" to the explanation of any of our observations and beliefs. And my argument will be that this is mistaken: that many moral explanations appear to be good explanations, or components in good explanations, that are not obviously undermined by anything else that we know. My suspicion, in fact, is that moral facts are needed in the sense explained, that they will turn out to belong in our best overall explanatory picture of the world, even in the long run, but I shall not attempt to establish that here. Indeed, it should be clear why I could not pretend to do so. For I have explicitly put to one side the issue (which I regard as incapable in any case of quick resolution) of whether and to what extent actual moral disagreements can be settled satisfactorily; but I assume it would count as a defect in any sort of explanation to rely on claims about which rational agreement proved unattainable. So I concede that it *could* turn out, for anything I say here, that moral explanations are all defective and should be discarded. What I shall try to show is merely that many moral explanations look reasonable enough to be in the running; and, more specifically, that nothing Harman says provides any reason for thinking they are not. This claim is surely strong enough (and controversial enough) to be worth defending.

(3) It is implicit in this statement of my project, but worth noting separately, that I take Harman to be proposing an *independent* skeptical argument—independent not merely of the argument from the difficulty of settling disputed moral questions, but also of other standard arguments for moral skepticism. Otherwise his argument is not worth separate discussion. For

any of these more familiar skeptical arguments will of course imply that moral explanations are defective, on the reasonable assumption that it would be a defect in any explanation to rely on claims as doubtful as these arguments attempt to show all moral claims to be. But if that is why there is a problem with moral explanations, one should surely just cite the relevant skeptical argument, rather than this derivative difficulty about moral explanations, as the basic "problem with ethics," and it is that argument we should discuss. So I take Harman's interesting suggestion to be that there is a *different* difficulty that remains even if we put other arguments for moral skepticism aside and *assume*, for the sake of argument, that there are moral facts (for example, that what the children in his example are doing is really wrong): namely, that these assumed facts still seem to play no explanatory role.

This understanding of Harman's thesis crucially affects my argumentative strategy in a way to which I should alert the reader in advance. For it should be clear that assessment of this thesis not merely permits, but *requires*, that we provisionally assume the existence of moral facts. I can see no way of evaluating the claim that *even* if we assumed the existence of moral facts they would still appear explanatorily irrelevant, without assuming the existence of some, to see how they would look. So I do freely assume this in each of the examples I discuss in the next section. (I have tried to choose plausible examples, moreover, moral facts most of us would be inclined to believe in if we did believe in moral facts, since those are the easiest to think about; but the precise examples don't matter, and anyone who would prefer others should feel free to substitute her own.) I grant, furthermore, that if Harman were right about the outcome of this thought experiment—that even after we assumed these facts they still looked irrelevant to the explanation of our moral beliefs and other nonmoral facts—then we might conclude with him that there were, after all, no such facts. But I

claim he is wrong: that once we have provisionally assumed the existence of moral facts they do appear relevant, by perfectly ordinary standards, to the explanation of moral beliefs and of a good deal else besides. Does this prove that there are such facts? Well of course it helps support that view, but here I carefully make no claim to have shown so much. What I show is that any remaining reservations about the existence of moral facts must be based on those other skeptical arguments, of which Harman's argument is independent. In short, there may still be a "problem with ethics," but it has nothing special to do with moral explanations.

### III. Moral explanations

Now that I have explained how I understand Harman's thesis, I turn to my arguments against it. I shall first add to my example of Hitler's moral character several more in which it seems plausible to cite moral facts as part of an explanation of nonmoral facts, and in particular of people's forming the moral opinions they do. I shall then argue that Harman gives us no plausible reason to reject or ignore these explanations; I shall claim, in fact, that the same is true for his own example of the children igniting the cat. I shall conclude, finally, by attempting to diagnose the source of the disagreement between Harman and me on these issues.

My Hitler example suggests a whole range of extremely common cases that appear not to have occurred to Harman, cases in which we cite someone's moral character as part of an explanation of his or her deeds, and in which that whole story is then available as a plausible further explanation of someone's arriving at a correct assessment of that moral character. Take just one other example. Bernard DeVoto, in *The Year of Decision: 1846*, describes the efforts of American emigrants already in California to rescue another party of emigrants, the Donner Party, trapped by snows in the High Sierras, once their plight

became known. At a meeting in Yerba Buena (now San Francisco), the relief efforts were put under the direction of a recent arrival, Passed Midshipman Selim Woodworth, described by a previous acquaintance as "a great busybody and ambitious of taking a command among the emigrants."<sup>14</sup> But Woodworth not only failed to lead rescue parties into the mountains himself, where other rescuers were counting on him (leaving children to be picked up by him, for example), but had to be "shamed, threatened and bullied" even into organizing the efforts of others willing to take the risk; he spent time arranging comforts for himself in camp, preening himself on the importance of his position; and as a predictable result of his cowardice and his exercises in vainglory, many died who might have been saved, including four known still to be alive when he turned back for the last time in mid-March.

DeVoto concludes: "Passed Midshipman Woodworth was just no damned good" (1942, p. 442). I cite this case partly because it has so clearly the structure of an inference to a reasonable explanation. One can think of competing explanations, but the evidence points against them. It isn't, for example, that Woodworth was a basically decent person who simply proved too weak when thrust into a situation that placed heroic demands on him. He volunteered, he put no serious effort even into tasks that required no heroism, and it seems clear that concern for his own position and reputation played a much larger role in his motivation than did any concern for the people he was expected to save. If DeVoto is right about this evidence, moreover, it seems reasonable that part of the explanation of his believing that Woodworth was no damned good is just that Woodworth was no damned good.

DeVoto writes of course with more moral intensity (and with more of a flourish) than academic historians usually permit themselves, but it would be difficult to find a serious work of biography, for example, in which actions are not

explained by appeal to moral character: sometimes by appeal to specific virtues and vices, but often enough also by appeal to a more general assessment. A different question, and perhaps a more difficult one, concerns the sort of example on which Harman concentrates, the explanation of judgments of right and wrong. Here again he appears just to have overlooked explanations in terms of moral character: a judge's thinking that it would be wrong to sentence a particular offender to the maximum prison term the law allows, for example, may be due in part to her decency and fairmindedness, which I take to be moral properties if any are. But do moral features of the action or institution being judged ever play an explanatory role? Here is an example in which they appear to. An interesting historical question is why vigorous and reasonably widespread moral opposition to slavery arose for the first time in the eighteenth and nineteenth centuries, even though slavery was a very old institution; and why this opposition arose primarily in Britain, France, and in French- and English-speaking North America, even though slavery existed throughout the New World.<sup>15</sup> There is a standard answer to this question. It is that chattel slavery in British and French America, and then in the United States, was much worse than previous forms of slavery, and much worse than slavery in Latin America. This is, I should add, a controversial explanation. But as is often the case with historical explanations, its proponents do not claim it is the whole story, and many of its opponents grant that there may be some truth in these comparisons, and that they may after all form a small part of a larger explanation.<sup>16</sup> This latter concession is all I require for my example. Equally good for my purpose would be the more limited thesis which explains the growth of antislavery sentiment in the United States, between the Revolution and the Civil War, in part by saying that slavery in the United States became a more oppressive institution during that time. The appeal in these standard explanations is straightforwardly to moral facts.

What is supposed to be wrong with all these explanations? Harman says that assumptions about moral facts seem "completely irrelevant" in explaining moral observations and moral beliefs (p. 7), but on its more natural reading that claim seems pretty obviously mistaken about these examples. For it is natural to think that if a particular assumption is completely irrelevant to the explanation of a certain fact, then that fact would have obtained, and we could have explained it just as well, even if the assumption had been false.<sup>17</sup> But I do not believe that Hitler would have done all he did if he had not been morally depraved, nor, on the assumption that he was not depraved, can I think of any plausible explanation for his doing those things. Nor is it plausible that we would all have believed he was morally depraved even if he hadn't been. Granted, there is a tendency for writers who do not attach much weight to fascism as a social movement to want to blame its evils on a single maniacal leader, so perhaps some of them would have painted Hitler as a moral monster even if he had not been one. But this is only a tendency, and one for which many people know how to discount, so I doubt that our moral belief really is overdetermined in this way. Nor, similarly, do I believe that Woodworth's actions were overdetermined, so that he would have done just as he did even if he had been a more admirable person. I suppose one could have doubts about DeVoto's objectivity and reliability; it is obvious he dislikes Woodworth, so perhaps he would have thought him a moral loss and convinced his readers of this no matter what the man was really like. But it is more plausible that the dislike is mostly based on the same evidence that supports DeVoto's moral view of him, and that very different evidence, at any rate, would have produced a different verdict. If so, then Woodworth's moral character is part of the explanation of DeVoto's belief about his moral character.

It is more plausible of course that serious moral opposition to slavery would have emerged

in Britain, France, and the United States even if slavery hadn't been worse in the modern period than before, and worse in the United States than in Latin America, and that the American anti-slavery movement would have grown even if slavery had not become more oppressive as the nineteenth century progressed. But that is because these moral facts are offered as at best a partial explanation of these developments in moral opinion. And if they really *are* part of the explanation, as seems plausible, then it is also plausible that whatever effect they produced was not entirely overdetermined; that, for example, the growth of the antislavery movement in the United States would at least have been somewhat slower if slavery had been and remained less bad an institution. Here again it hardly seems "completely irrelevant" to the explanation whether or not these moral facts obtained.

It is more puzzling, I grant, to consider Harman's own example in which you see the children igniting a cat and react immediately with the thought that this is wrong. Is it true, as Harman claims, that the assumption that the children are really doing something wrong is "totally irrelevant" to any reasonable explanation of your making that judgment? Would you, for example, have reacted in just the same way, with the thought that the action is wrong, even if what they were doing *hadn't* been wrong, and could we explain your reaction equally well on this assumption? Now, there is more than one way to understand this counterfactual question, and I shall return below to a reading of it that might appear favorable to Harman's view. What I wish to point out for now is merely that there is a natural way of taking it, parallel to the way in which I have been understanding similar counterfactual questions about my own examples, on which the answer to it has to be simply: it depends. For to answer the question, I take it,<sup>18</sup> we must consider a situation in which what the children are doing is not wrong, but which is otherwise as much like the actual situation as possible, and then decide what your reaction

would be in that situation. But since what makes their action wrong, what its wrongness consists in, is presumably something like its being an act of gratuitous cruelty (or, perhaps we should add, of intense cruelty, and to a helpless victim), to imagine them not doing something wrong we are going to have to imagine their action different in this respect. More cautiously and more generally, if what they are actually doing is wrong, and if moral properties are, as many writers have held, supervenient on natural ones,<sup>19</sup> then in order to imagine them not doing something wrong we are going to have to suppose their action different from the actual one in some of its natural features as well. So our question becomes: Even if the children had been doing something else, something just different enough not to be wrong, would you have taken them even so to be doing something wrong?

Surely there is no one answer to this question. It depends on a lot about you, including your moral views and how good you are at seeing at a glance what some children are doing. It probably depends also on a debatable moral issue; namely, just *how* different the children's action would have to be in order not to be wrong. (Is unkindness to animals, for example, also wrong?) I believe we can see how, in a case in which the answer was clearly affirmative, we might be tempted to agree with Harman that the wrongness of the action was no part of the explanation of your reaction. For suppose you are like this. You hate children. What you especially hate, moreover, is the sight of children enjoying themselves; so much so that whenever you see children having fun, you immediately assume they are up to no good. The more they seem to be enjoying themselves, furthermore, the readier you are to fasten on any pretext for thinking them engaged in real wickedness. Then it is true that even if the children had been engaged in some robust but innocent fun, you would have thought they were doing something wrong; and Harman is perhaps right<sup>20</sup> about you that the actual wrongness of the action you

see is irrelevant to your thinking it wrong. This is because your reaction is due to a feature of the action that coincides only very accidentally with the ones that make it wrong.<sup>21</sup> But, of course, and fortunately, many people aren't like this (nor does Harman argue that they are). It isn't true of them, in general, that if the children had been doing something similar, although different enough not to be wrong, they would still have thought the children were doing something wrong. And it isn't true either, therefore, that the wrongness of the action is irrelevant to the explanation of why they think it wrong.

Now, one might have the sense from my discussion of all these examples, but perhaps especially from my discussion of this last one, Harman's own, that I have perversely been refusing to understand his claim about the explanatory irrelevance of moral facts in the way he intends. And perhaps I have not been understanding it as he wishes. In any case, I agree, I have certainly not been understanding the crucial counterfactual question, of whether we would have drawn the same moral conclusion even if the moral facts had been different, in the way he must intend. But I am not being perverse. I believe, as I have said, that my way of taking the question is the more natural one. And, more importantly: although there is, I grant, a reading of that question on which it will always yield the answer Harman wants—namely, that a difference in the moral facts would not have made a difference in our judgment—I do not believe this reading can support his argument. I must now explain why.

It will help if I contrast my general approach with his. I am approaching questions about the justification of belief in the spirit of what Quine has called "epistemology naturalized."<sup>22</sup> I take this to mean that we have in general no a priori way of knowing which strategies for forming and refining our beliefs are likely to take us closer to the truth. The only way we have of proceeding is to assume the approximate truth of what seems to us the best overall theory we

already have of what we are like and what the world is like, and to decide in the light of that what strategies of research and reasoning are likely to be reliable in producing a more nearly true overall theory. One result of applying these procedures, in turn, is likely to be the refinement or perhaps even the abandonment of parts of the tentative theory with which we began.

I take Harman's approach, too, to be an instance of this one. He says we are justified in believing in those facts that we need to assume to explain why we observe what we do. But he does not think that our knowledge of this principle about justification is a priori. Furthermore, as he knows, we cannot decide whether one explanation is better than another without relying on beliefs we already have about the world. Is it really a better explanation of the vapor trail the physicist sees in the cloud chamber to suppose that a proton caused it, as Harman suggests in his example, rather than some other charged particle? Would there, for example, have been no vapor trail in the absence of that proton? There is obviously no hope of answering such questions without assuming at least the approximate truth of some quite far-reaching microphysical theory, and our knowledge of such theories is not a priori.

But my approach differs from Harman's in one crucial way. For among the beliefs in which I have enough confidence to rely on in evaluating explanations, at least at the outset, are some moral beliefs. And I have been relying on them in the following way.<sup>23</sup> Harman's thesis implies that the supposed moral fact of Hitler's being morally depraved is irrelevant to the explanation of Hitler's doing what he did. (For we may suppose that if it explains his doing what he did, it also helps explain, at greater remove, Harman's belief and mine in his moral depravity.) To assess this claim, we need to conceive a situation in which Hitler was not morally depraved and consider the question whether in that situation he would still have done what he did. My answer is that he would

not, and this answer relies on a (not very controversial) moral view: that in any world at all like the actual one, only a morally depraved person could have initiated a world war, ordered the “final solution,” and done any number of other things Hitler did. That is why I believe that, if Hitler hadn’t been morally depraved, he wouldn’t have done those things, and hence that the fact of his moral depravity is relevant to an explanation of what he did.

Harman, however, cannot want us to rely on any such moral views in answering this counterfactual question. This comes out most clearly if we return to his example of the children igniting the cat. He claims that the wrongness of this act is irrelevant to an explanation of your thinking it wrong, that you would have *thought* it wrong even if it wasn’t. My reply was that in order for the action not to be wrong it would have had to lack the feature of deliberate, intense, pointless cruelty, and that if it had differed in this way you might very well *not* have thought it wrong. I also suggested a more cautious version of this reply: that since the action is in fact wrong, and since moral properties supervene on more basic natural ones, it would have had to be different in some further natural respect in order not to be wrong; and that we do not know whether if it had so differed you would still have thought it wrong. Both of these replies, again, rely on moral views, the latter merely on the view that there is *something* about the natural features of the action in Harman’s example that makes it wrong, the former on a more specific view as to which of these features do this.

But Harman, it is fairly clear, intends for us not to rely on any such moral views in evaluating his counterfactual claim. His claim is not that if the action had not been one of deliberate cruelty (or had otherwise differed in whatever way would be required to remove its wrongness), you would still have thought it wrong. It is, instead, that if the action were one of deliberate, pointless cruelty, but this *did not make it wrong*, you would still have thought it was wrong. And to

return to the example of Hitler’s moral character, the counterfactual claim that Harman will need in order to defend a comparable conclusion about that case is not that if Hitler had been, for example, humane and fairminded, free of nationalistic pride and racial hatred, he would still have done exactly as he did. It is, rather, that if Hitler’s psychology, and anything else about his situation that could strike us as morally relevant, had been exactly as it in fact was, but this had *not constituted moral depravity*, he would still have done exactly what he did.

Now the antecedents of these two conditionals are puzzling. For one thing, both are, I believe, necessarily false. I am fairly confident, for example, that Hitler really was morally depraved;<sup>24</sup> and since I also accept the view that moral features supervene on more basic natural properties,<sup>25</sup> I take this to imply that there is no possible world in which Hitler has just the personality he in fact did, in just the situation he was in, but is not morally depraved. Any attempt to describe such a situation, moreover, will surely run up against the limits of our moral concepts—what Harman calls our “moral sensibility”—and this is no accident. For what Harman is asking us to do, in general, is to consider cases in which absolutely *everything* about the nonmoral facts that could seem morally relevant to us, in light of whatever moral theory we accept and of the concepts required for understanding that theory, is held fixed, but in which the moral judgment that our theory yields about the case is nevertheless mistaken. So it is hardly surprising that, using that theory and those concepts, we should find it difficult to conceive in any detail what such a situation would be like. It is especially not surprising when the cases in question are as paradigmatic in light of the moral outlook we in fact have as is Harman’s example or is, even more so, mine of Hitler’s moral character. The only way we could be wrong about this latter case (assuming we have the nonmoral facts right) would be for our whole theory to be hopelessly wrong, so

radically mistaken that there could be no hope of straightening it out through adjustments from within.

But I do not believe we should conclude, as we might be tempted to,<sup>26</sup> that we therefore know a priori that this is not so, or that we cannot understand these conditionals that are crucial to Harman's argument. Rather, now that we have seen how we have to understand them, we should grant that they are true: that if our moral theory were somehow hopelessly mistaken, but all the nonmoral facts remained exactly as they in fact are, then, since we do accept that moral theory, we would still draw exactly the moral conclusions we in fact do. But we should deny that any skeptical conclusion follows from this. In particular, we should deny that it follows that moral facts play no role in explaining our moral judgments.

For consider what follows from the parallel claim about microphysics, in particular about Harman's example in which a physicist concludes from his observation of a vapor trail in a cloud chamber, and from the microphysical theory he accepts, that a free proton has passed through the chamber. The parallel claim, notice, is not just that if the proton had not been there the physicist would still have thought it was. This claim is implausible, for we may assume that the physicist's theory is generally correct, and it follows from that theory that if there hadn't been a proton there, then there wouldn't have been a vapor trail. But in a perfectly similar way it is implausible that if Hitler hadn't been morally depraved we would still have thought he was: for we may assume that our moral theory also is at least roughly correct, and it follows from the most central features of that theory that if Hitler hadn't been morally depraved, he wouldn't have done what he did. The parallel claim about the microphysical example is, instead, that if there hadn't been a proton there, but there *had* been a vapor trail, the physicist would still have concluded that a proton was present. More precisely, to maintain

a perfect parallel with Harman's claims about the moral cases, the antecedent must specify that although no proton is present, absolutely all the non-microphysical facts that the physicist, in light of his theory, might take to be relevant to the question of whether or not a proton is present, are exactly as in the actual case. (These macrophysical facts, as we may for convenience call them, surely include everything one would normally think of as an observable fact.) Of course, we shall be unable to imagine this without imagining that the physicist's theory is pretty badly mistaken;<sup>27</sup> but I believe we should grant that, if the physicist's theory were somehow this badly mistaken, but all the macrophysical facts (including all the observable facts) were held fixed, then the physicist, since he does accept that theory, would still draw all the same conclusions that he actually does. That is, this conditional claim, like Harman's parallel claim about the moral cases, is true.

But no skeptical conclusions follow; nor can Harman, since he does not intend to be a skeptic about physics, think that they do. It does not follow, in the first place, that we have any reason to think the physicist's theory is generally mistaken. Nor does it follow, furthermore, that the hypothesis that a proton really did pass through the cloud chamber is not part of a good explanation of the vapor trail, and hence of the physicist's thinking this has happened. This looks like a reasonable explanation, of course, only on the assumption that the physicist's theory is at least roughly true, for it is this theory that tells us, for example, what happens when charged particles pass through a supersaturated atmosphere, what other causes (if any) there might be for a similar phenomenon, and so on. But, as I say, we have not been provided with any reason for not trusting the theory to this extent.

Similarly, I conclude, we should draw no skeptical conclusions from Harman's claims about the moral cases. It is true that if our moral theory were seriously mistaken, but we still believed it, and the nonmoral facts were held

fixed, we would still make just the moral judgments we do. But this fact by itself provides us with no reason for thinking that our moral theory is generally mistaken. Nor, again, does it imply that the fact of Hitler's really having been morally depraved forms no part of a good explanation of his doing what he did and hence, at greater remove, of our thinking him depraved. This explanation will appear reasonable, of course, only on the assumption that our accepted moral theory is at least roughly correct, for it is this theory that assures us that only a depraved person could have thought, felt, and acted as Hitler did. But, as I say, Harman's argument has provided us with no reason for not trusting our moral views to this extent, and hence with no reason for doubting that it is sometimes moral facts that explain our moral judgments.

I conclude with three comments about my argument.

(1) I have tried to show that Harman's claim—that we would have held the particular moral beliefs we do even if those beliefs were untrue—admits of two readings, one of which makes it implausible, and the other of which reduces it to an application of a general skeptical strategy, a strategy which could as easily be used to produce doubt about microphysical as about moral facts. The general strategy is this. Consider any conclusion *C* we arrive at by relying both on some distinguishable "theory" *T* and on some body of evidence not being challenged, and ask whether we would have believed *C* even if it had been false. The plausible answer, if we are allowed to rely on *T*, will often be no: for if *C* had been false, then (according to *T*) the evidence would have had to be different, and in that case we wouldn't have believed *C*. (I have illustrated the plausibility of this sort of reply for all my moral examples, as well as for the microphysical one.) But the skeptic of course intends us not to rely on *T* in this way, and so rephrases the question: Would we have believed *C* even if it were false but all the evidence had been exactly as it in fact was? Now the answer has to be yes; and the

skeptic concludes that *C* is doubtful. (It should be obvious how to extend this strategy to belief in other minds, or in an external world.) I am of course not convinced: I do not think answers to the rephrased question show anything interesting about what we know or justifiably believe. But it is enough for my purposes here that no such general skeptical strategy could pretend to reveal any problems peculiar to belief in moral facts.

(2) My conclusion about Harman's argument, although it is not exactly the same as, is nevertheless similar to and very much in the spirit of the Duhemian point I invoked earlier against verificationism. There the question was whether typical moral assertions have testable implications, and the answer was that they do, so long as you include additional moral assumptions of the right sort among the background theories on which you rely in evaluating these assertions. Harman's more important question is whether we should ever regard moral facts as relevant to the explanation of nonmoral facts, and in particular of our having the moral beliefs we do. But the answer, again, is that we should, so long as we are willing to hold the right sorts of other moral assumptions fixed in answering counterfactual questions. Neither answer shows morality to be on any shakier ground than, say, physics, for typical microphysical hypotheses, too, have testable implications, and appear relevant to explanations, only if we are willing to assume at least the approximate truth of an elaborate microphysical theory and to hold this assumption fixed in answering counterfactual questions.

(3) Of course, this picture of how explanations depend on background theories, and moral explanations in particular on moral background theories, does show why someone already tempted toward moral skepticism on other grounds (such as those I mentioned at the beginning of this essay) might find Harman's claim about moral explanations plausible. To the extent that you already have pervasive doubts

about moral theories, you will also find moral facts nonexplanatory. So I grant that Harman has located a natural symptom of moral skepticism; but I am sure he has neither traced this skepticism to its roots nor provided any independent argument for it. His claim (p. 22) that we do not in fact cite moral facts in explanation of moral beliefs and observations cannot provide such an argument, for that claim is false. So, too, is the claim that assumptions about moral facts seem irrelevant to such explanations, for many do not. The claim that we should not rely on such assumptions because they are irrelevant, on the other hand, unless it is supported by some independent argument for moral skepticism, will just be question-begging: for the principal test of whether they are relevant, in any situation in which it appears they might be, is a counterfactual question about what would have happened if the moral fact had not obtained, and how we answer that question depends precisely upon whether we do rely on moral assumptions in answering it.

My own view I stated at the outset: that the only argument for moral skepticism with any independent weight is the argument from the difficulty of settling disputed moral questions. I have shown that anyone who finds Harman's claim about moral explanations plausible must already have been tempted toward skepticism by some other considerations, and I suspect that the other considerations will typically be the ones I sketched. So that is where discussion should focus. I also suggested that those considerations may provide less support for moral skepticism than is sometimes supposed, but I must reserve a thorough defense of that thesis for another occasion.<sup>28</sup>

## Notes

1 As, for example, in Alan Gewirth, "Positive 'Ethics' and Normative 'Science' ", *The Philosophical Review* 69 (1960), pp. 311–30, in which there are some useful remarks about the first of them.

- 2 J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth, England: Penguin, 1977), pp. 38–42.
- 3 Gilbert Harman, *The Nature of Morality: An Introduction to Ethics* (New York: Oxford University Press, 1977), pp. vii, viii. Parenthetical page references are to this work.
- 4 Harman's title for the entire first section of his book.
- 5 In the longer article of which this is an abridgement.
- 6 This point is generally credited to Pierre Duhem; see *The Aim and Structure of Physical Theory*, trans. Philip P. Wiener (Princeton, NJ: Princeton University Press, 1954). It is a prominent theme in the influential writings of W. V. O. Quine. For an especially clear application of it, see Hilary Putnam, "The 'Corroboration' of Theories," in *Mathematics, Matter and Method: Philosophical Papers, Volume I*, second ed. (Cambridge: Cambridge University Press, 1977), pp. 250–69.
- 7 See note 5.
- 8 Harman is careful always to say only that moral beliefs appear to play no such role, and since he eventually concludes that there are moral facts (p. 132), this caution may be more than stylistic. I shall argue that this more cautious claim, too, is mistaken (indeed, that is my central thesis). But to avoid issues about Harman's intent, I shall simply mean by "Harman's argument" the skeptical argument of his first two chapters, whether or not he means to endorse all of it. This argument surely deserves discussion in its own right in either case, especially since Harman never explains what is wrong with it.
- 9 He asks: "Can moral principles be tested in the same way [as scientific hypotheses can], out in the world? You can observe someone do something, but can you ever perceive the rightness or wrongness of what he does?" (p. 4).
- 10 The other is that Harman appears to use "observe" (and "perceive" and "see") in a surprising way. One would normally take observing (or perceiving, or seeing) something to involve knowing it was the case. But Harman apparently takes an observation to be any opinion arrived at as "a direct result of perception" (p. 5) or, at any rate (see next footnote), "immediately and without

conscious reasoning” (p. 7). This means that observations need not even be true, much less known to be true. A consequence is that the existence of moral observations, in Harman’s sense, would not be sufficient to show that there is moral knowledge, although this would be sufficient if “observe” were being used in a more standard sense. What I argue in the text is that the existence of moral observations (in either Harman’s or the standard sense) is not *necessary* for showing that there is moral knowledge, either.

- 11 This sort of case does not meet Harman’s characterization of an observation as an opinion that is “a direct result of perception” (p. 5), but he is surely right that moral facts would be as well vindicated if they were needed to explain our drawing conclusions about hypothetical cases as they would be if they were needed to explain observations in the narrower sense. To be sure, Harman is still confining his attention to cases in which we draw the moral conclusion from our thought experiment “immediately and without conscious reasoning” (p. 7), and it is no doubt the existence of such cases that gives purchase to talk of a “moral sense.” But this feature, again, can hardly matter to the argument: would belief in moral facts be less justified if they were needed only to explain the instances in which we draw the moral conclusion *slowly*? Nor can it make any difference for that matter whether the case we are reflecting on is hypothetical: so my example in which we, quickly or slowly, draw a moral conclusion about Hitler from what we know of him, is surely relevant.
- 12 In the longer paper from which this one is abridged. The salient point is that there are two very *different* reasons one might have for thinking that no reference to moral facts is needed in the explanation of moral beliefs. One—Harman’s reason, and my target in this essay—is that no moral explanations even *seem* plausible, that reference to moral facts always strikes us as “completely irrelevant” to the explanation of moral beliefs. This claim, if true, would tend to support moral skepticism. The other, which might appeal to a “reductive” naturalist in ethics, is that any moral explanations that *do* seem plausible can be paraphrased without explanatory loss in entirely nonmoral terms. I doubt this view, too, and I argue in the longer version of this paper that no ethical naturalist need hold it. But anyone tempted by it should note that it is anyway no version of moral skepticism: for what it says is that we know so much about ethics that we are always able to say, in entirely nonmoral terms, exactly which natural properties the moral terms in any plausible moral explanations refer to—that’s why the moral expressions are dispensable. These two reasons should not be confused with one another.
- 13 It is surprising that Harman does not mention the obvious intermediate possibility, which would occur to any instrumentalist: to cite the physicist’s psychological set and the vapor trail, but say nothing about protons or other unobservables. It is this explanation, as I emphasize below, that is most closely parallel to an explanation of beliefs about an external world in terms of sensory experience and psychological makeup, or of moral beliefs in terms of nonmoral facts together with our “moral sensibility.”
- 14 Bernard DeVoto, *The Year of Decision: 1846* (Boston: Houghton Mifflin, 1942), p. 426; a quotation from the notebooks of Francis Parkman. The account of the entire rescue effort is on pp. 424–44.
- 15 What is being explained, of course, is not just why people came to think slavery wrong, but why people who were not themselves slaves or in danger of being enslaved came to think it so seriously wrong as to be intolerable. There is a much larger and longer history of people who thought it wrong but tolerable and an even longer one of people who appear not to have gotten past the thought that the world would be a better place without it. See David Brion Davis, *The Problem of Slavery in Western Culture* (Ithaca, NY: Cornell University Press, 1966).
- 16 For a version of what I am calling the standard view of slavery in the Americas, see Frank Tannenbaum, *Slave and Citizen* (New York: Alfred A. Knopf, 1947). For an argument against both halves of the standard view, see Davis, *The Problem of Slavery*, esp. pp. 60–61, 223–25, 262–63.
- 17 This counterfactual test requires qualification to be exactly right, but none of the plausible qualifications matters to my examples. See the longer version of this paper.

- 18 Following, informally, Stalnaker and Lewis on counterfactuals. See Robert Stalnaker, "A Theory of Conditionals," in Nicholas Rescher, ed., *Studies in Logical Theory*, APQ Monograph No. 2 (Oxford: Basil Blackwell, 1968); and David Lewis, *Counterfactuals* (Cambridge, MA: Harvard University Press, 1973).
- 19 What would be generally granted is just that if there are moral properties they supervene on natural properties. But, remember, we are assuming for the sake of argument that there are.
- I think moral properties are natural properties; and from this view it of course follows trivially that they supervene on natural properties: that, necessarily, nothing could differ in its moral properties without differing in some natural respect. But I also accept the more interesting thesis usually intended by the claim about supervenience—that there are more basic natural features such that, necessarily, once they are fixed, so are the moral properties. (In supervening on more basic natural facts of some sort, moral facts are like most natural facts. Social facts like unemployment, for example, supervene on complex histories of many individuals and their relations; and facts about the existence and properties of macroscopic physical objects—colliding billiard balls, say—clearly supervene on the microphysical constitution of the situations that include them.)
- 20 Not *certainly* right, because there is still the possibility that your reaction is to some extent overdetermined, and is to be explained partly by your sympathy for the cat and your dislike of cruelty, as well as by your hatred for children (although this last alone would have been sufficient to produce it).
- We could of course rule out this possibility by making you an even less attractive character, indifferent to the suffering of animals and not offended by cruelty. But it may then be hard to imagine that such a person (whom I shall cease calling "you") could retain enough of a grip on moral thought for us to be willing to say he thought the action *wrong*, as opposed to saying that he merely pretended to do so. This difficulty is perhaps not insuperable, but it is revealing. Harman says that the actual wrongness of the action is "completely irrelevant" to the explanation of the observer's reaction. Notice that what is in fact true, however, is that it is *very hard* to imagine someone who reacts in the way Harman describes, but whose reaction is not due, at least in part, to the actual wrongness of the action.
- 21 Perhaps deliberate cruelty is worse the more one enjoys it (a standard counterexample to hedonism). If so, the fact that the children are enjoying themselves makes their action worse, but presumably isn't what makes it wrong to begin with.
- 22 W. V. O. Quine, "Epistemology Naturalized," in *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969), pp. 69–90. In the same volume, see also "Natural Kinds," pp. 114–38.
- 23 Harman of course allows us to assume the moral facts whose explanatory relevance is being assessed: that Hitler was depraved, or that what the children in his example are doing is wrong. But I have been assuming something more—something about what depravity is, and about what makes the children's action wrong. (At a minimum, in the more cautious version of my argument, I have been assuming that *something* about its more basic features makes it wrong, so that it could not have differed in its moral quality without differing in those other features as well.)
- 24 And anyway, remember, this is the sort of fact Harman allows us to assume in order to see whether, if we assume it, it will look explanatory.
- 25 It is about here that I have several times encountered the objection: but surely *supervenient* properties aren't needed to explain anything. It is a little hard, however, to see just what this objection is supposed to come to. If it includes endorsement of the conditional I here attribute to Harman, then I believe the remainder of my discussion is an adequate reply to it. If it is the claim that, because moral properties are supervenient, we can always exploit the insights in any moral explanations, however plausible, without resort to moral *language*, then I have already dealt with it in my discussion of reductionism (see note 10, above): the claim is probably false, but even if it is true it is no support for Harman's view, which is not that moral explanations are plausible but reducible, but that they are totally implausible. And doubts about the causal efficacy of supervenient facts seem misplaced in any case, as attention to my earlier examples (note 17) illustrates. High

unemployment causes widespread hardship, and can also bring down the rate of inflation. The masses and velocities and two colliding billiard balls causally influence the subsequent trajectories of the two balls. There is no doubt some sense in which these facts are causally efficacious in virtue of the way they supervene on—that is, are constituted out of, or causally realized by—more basic facts, but this hardly shows them inefficacious. (Nor does Harman appear to think it does: for his favored explanation of your moral belief about the burning cat, recall, appeals to psychological facts [about your moral sensibility], a biological fact [that it's a cat], and macrophysical facts [that it's on fire]—supervenient facts all, on his physicalist view and mine.) If anyone does hold to a general suspicion of causation by supervenient facts and properties, however, as Jaegwon Kim appears to (see “Causality, Identity, and Supervenience in the Mind-Body Problem,” *Midwest Studies in Philosophy* 4 [1979], pp. 31–49), it is enough to note that this suspicion cannot diagnose any special difficulty with moral explanations, any distinctive “problem with ethics.” The “problem,” arguably, will be with every discipline but fundamental physics.

- 26 And as I take it Philippa Foot, for example, is still prepared to do, at least about paradigmatic cases. See her *Moral Relativism* (Lawrence: The University of Kansas, 1978).
- 27 If we imagine the physicist regularly mistaken in this way, moreover, we will have to imagine his theory not just mistaken but hopelessly so. And we can easily reproduce the other notable feature of Harman's claims about the moral cases, that what we are imagining is necessarily false, if we suppose that one of the physicist's (or better, chemist's) conclusions is about the microstructure of some common substance, such as water. For I agree with Saul Kripke that whatever microstructure water has is essential to it, that it has this structure in every possible world in which it exists (Saul Kripke, *Naming and Necessity* [Cambridge, MA: Harvard

University Press, 1980]). If we are right (as we have every reason to suppose) in thinking that water is actually H<sub>2</sub>O, therefore, the conditional, “If water were not H<sub>2</sub>O, but all the observable, macrophysical facts were just as they actually are, chemists would still have come to think it was H<sub>2</sub>O,” has a necessarily false antecedent; just as, if we are right (as we also have good reason to suppose) in thinking that Hitler was actually morally depraved, the conditional, “If Hitler were just as he was in all natural respects, but not morally depraved, we would still have thought he was depraved,” has a necessarily false antecedent. Of course, I am not suggesting that in either case our knowledge that the antecedent is false is a priori.

These counterfactuals, because of their impossible antecedents, will have to be interpreted over worlds that are (at best) only “epistemically” possible; and, as Richard Boyd has pointed out to me, this helps to explain why anyone who accepts a causal theory of knowledge (or any theory according to which the justification of our belief depends on what explains our holding them) will find their truth irrelevant to the question of how much we know, either in chemistry or in morals. For although there certainly are counterfactuals that are relevant to questions about what causes what (and, hence, about what explains what), these have to be counterfactuals about real possibilities, not merely epistemic ones.

- 28 This essay has benefited from helpful discussion of earlier versions read at the University of Virginia, Cornell University, Franklin and Marshall College, Wayne State University, and the University of Michigan. I have been aided by a useful correspondence with Gilbert Harman; and I am grateful also for specific comments from Richard Boyd, David Brink, David Copp, Stephen Darwall, Terence Irwin, Norman Kretzmann, Ronald Nash, Peter Railton, Bruce Russell, Sydney Shoemaker, and Judith Slein.