

CHAPTER 4

IN DEFENCE OF MORAL REALISM

1. INTRODUCTION

Moral realism as here conceived has two tenets, one semantic thesis to the effect that when we pass moral judgements, we make truth-claims of an irreducible kind, and one ontological thesis to the effect that there are objective normative facts (truths) out there, *sui generis*, to be discovered. Some moral propositions are true. I have defended moral realism elsewhere (in my book *Moral Realism*) and recently Russ Shafer-Landau has also published a book with the same title. The main difference between my argument here, and in my earlier book, is that I now defend normative realism — I earlier focused rather on intrinsic value. The difference is perhaps more verbal than real, since I did hold then, and still hold, that there is only *one* problem to be pondered in ethics. I now find it more natural to speak of it as the normative problem, however. My main disagreement with Shafer-Landau concerns how we should answer Mackie's argument from queerness, in its part concerning supervenience, but we also disagree on matters to do with moral epistemology. I have not tried to relate my new argument in detail either to my own previous view or to Shafer-Landau's more recent argument; my argument here stands on its own feet. And the question I ponder is why we should believe in moral realism.

Now, the semantic aspect of this question seems to me simple. It is obvious, that when we pass moral judgements to the effect that certain actions are right, other actions wrong, and so forth, then we intend to make objective judgements capable of being true or false. It is also clear, I think, that we make judgements that are not merely descriptive of empirical realities. For example, moral judgements are not elliptical, saying things like, *according to the norms existing in my society, this action is right or wrong* or, *this action is liked, or disliked, by me*. It is noteworthy that even a philosopher like John Mackie, who thought that there are no objective moral facts, thought that, in issuing moral judgements, we imply (wrongly) that such facts exist.

Why should we draw the conclusion that our moral judgements are objective (in intent)? Well, it is sufficient to observe carefully how we make these judgements, and how we react to moral phenomena such as moral disagreement, and (putative) moral mistakes.

When we run upon conflicting moral judgements we believe that both parties cannot be right. Why don't we believe that they can? The best explanation is that we take them to make objective and contradictory judgements. But, if we run upon inconsistent judgements, at least one judgement must be false. And we often feel that we come to the conclusion that what we once believed was a proper moral judgement was, upon closer inspection, not a proper one. We then tend to believe, not only that we judged the case differently before, but that, before, we judged it *wrongly*. This means that we assume that there is a fact of the matter to be right or wrong about.

Or, upon reflection, this claim may be too strong. It is not certain that *everyone, all the time*, is using a moral vocabulary in the objective sense here defined. Perhaps this is even a rare, and fairly recent phenomenon. It may well be true that some people, who have been indoctrinated into a particular moral creed, perhaps with the argument that it has been dictated to us human beings by God, are merely intending to report what is required by this conventional moral creed, with which they sympathise, when they pass moral judgements. We may suspect that this is a case if there is little or no moral disagreement in a society. In that case, a moral semantic (indexical) relativism of the kind defended by Wong¹ and Harman² may be true of the moral utterances made by people in this society. Other people, living in other societies, may have revolted against the morality into which they have been socialised, but not because they have started to make claims about a moral reality they believe exists independently of their conceptualisation; instead they may be just evincing their strong emotions (even their gut feelings) when they pass moral judgements. Hence emotivism or expressivism of the kind defended by Stevenson³ and Gibbard⁴ may be true of *their* moral utterances. However, even if there are such cases, and even if they are what we often meet with, when we observe how people pass moral judgements, there are also cases

¹ See for example his *Moral Relativity*.

² See for example his contribution to Gilbert Harman and Judith Jarvis Thomson *Moral Relativism and Moral Objectivity*.

³ See his *Ethics and Language and Facts and Values*.

⁴ See his *Wise Choices, Apt Feelings*.

where people do intend their moral judgements to capture a moral reality that is seen by them as *sui generis*. This is certainly how Sidgwick and Moore conceived of their own moral judgements, and this is how Moore conceived of his disagreement with Sidgwick, and this is how contemporary moral realists such as Parfit and myself conceive of our own moral judgments and disagreements. And even if it is *possible* that we delude ourselves, even if it is *possible* that we are mistaken in our understanding of our own words, this possibility is far-fetched, considered the way at least we moral realists conceive of the subject-matter of morality. Once again, we pursue the truth, we want to avoid inconsistency (since we believe that two inconsistent beliefs cannot both be true), we sometimes feel that we were previously wrong when we held opinions different from the ones we now hold. Hence, it is reasonable to believe that Mackie is right when he claims that an objectivist moral semantics is true at least of *our* moral utterances.

However, John Mackie is not only famous because he thought that our moral judgements are objective, he also, as was alluded to above, thought that they were all false. He argued that 'although most people in making moral judgements implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false.'⁵ I will speak of this ontological view as 'moral nihilism', even if this was not the term Mackie himself used.

Mackie's semantic view, to the effect that we refer to objective moral properties when we make moral statements, together with his ontological nihilism, to the effect that nothing possesses these properties — Mackie's error theory — has by some been considered problematic. If the claim that torture is wrong is false, does this not mean that torture is not wrong? But does not this entail that torture is all right? And is this not a strange moral implication, which appears to follow from Mackie's error theory? Moreover, if it is also false that torture is all right, does not this imply that torture is wrong? But then torture is both wrong and right (not wrong), which seems to be not only morally speaking strange but incoherent.

Appearances are deceptive. If Mackie's error theory is correct, then both the claim that torture is right and the claim that torture is wrong are false. We realise this when we realise that only together with the premise that all actions have normative status (they are right *or* wrong) are we allowed to conclude that, if an action is not wrong, then it is right. A moral realist accepts the first premise, I submit. All actions have normative status, the moral realist is likely to claim. And our common sense moral metaphysics is realist as well. But Mackie denies the first premise, of course. According to his theory, all actions lack moral status. According to his theory, all positive moral judgements to the effect that an act is right, wrong, or obligatory, are hence false.⁶ I suppose he is also making the weaker claim, now made popular by so called particularists, that all moral *principles* are false. All this means that his theory is consistent and it has no ethical implications.

Yet, it is a strange theory, or a strange *combination* of views. But what is strange about the combination is perhaps not the semantic objectivist part, but the corresponding nihilist ontological assumption. What lead Mackie to his moral nihilism?

His most important arguments to this effect are his arguments from 'queerness' and from 'relativity'. I will examine them in order.

2. THE ARGUMENT FROM QUEERNESS

Mackie assumed that, when we make moral judgements, not only do we ascribe an objective property to the actions we are passing moral judgements upon, but we also attribute to it a property which is inherently prescriptive. He thought that we think, when we say of a certain action that it is right, not only that it has a certain objective property (rightness) but also that this property has the property in its turn to motivate us to act, at least when we think it is present. Moral properties, then, have 'intrinsic to-be-pursuedness' built into them. But such a property would be very queer indeed. This is how he states his argument:

If there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else.⁷

I think Mackie is right when he insists that this would be a very strange (queer) property. However, I think he is wrong when he believes that we attribute such a property to an action when we claim that it is right. We do attribute an objective property, I concede. However, we do not attribute any inherently motivating property. And yet, for all that, Mackie's observation is close to the mark. As a matter of fact, his mistake, when he believes

⁵ *Ethics*, p. 35.

⁶ This is a very natural and plausible interpretation of Mackie's error theory; for an elaboration on it, see Sinnott-Armstrong, *Moral Skepticisms* pp. 34-36, even though Sinnott-Armstrong does not generalise it to permissions.

⁷ *Ibid.* p. 38.

that, when we pass moral judgements we want to attribute an inherently motivating property is highly instructive, so let us try to see how he goes wrong.

It is certainly true that most people when they believe that an action ought to be done have some incentive to perform this action. But this does not mean that the property of obligatoriness *as such* has any inherent magnetic force attached to it. It is because most people *care* about doing their duty that they are somewhat inclined to do a certain action, once they acknowledge that it is obligatory.

This means that Mackie was not completely wrong when he claimed that normative properties are in a manner of speaking motivating. However, this motivating force is not always present, and it is certainly not inherent in the moral properties as such. This property should rather be explained along the following lines: It is a remarkable fact that some normative judgements, such as most typically the expression of obligations, have, not only truth-conditions, but satisfaction conditions as well. The proposition that I ought to do X can be satisfied or not satisfied. It is satisfied, if I do X, and it is not satisfied, if I don't. And we take an interest in these satisfaction-conditions. These satisfaction-conditions, and our interest in them, explain why we feel that moral properties are, in a sense, 'prescriptive'.

In fact, the notion that norms have satisfaction values takes care of *one* aspect of Mackie's idea of prescriptivity. He explains what he means by prescriptivity, with a reference to Plato's idea of the Forms, and stresses *two* rather different aspects.

The Form of the Good is such that knowledge of it provides the knower with both a *direction* and an *overriding motive*; something's being good both *tells* the person who knows this to pursue it and *makes him* pursue it.⁸

The fact that norms have satisfaction values explains why they give us direction. This means that the proposition that I ought to do F can be, not only true or false, but satisfied or not satisfied. In this sense moral properties *are* prescriptive. In this a normative claim is different from other truth-claims, such as the claim that the sky is blue. If there are moral properties, then this means that there are properties that are different from ordinary natural ones, of course. But there is nothing *queer* about that.

I have also taken moral realism to imply that answers to the question what I ought to do *simpliciter* has a special *authority*. Is that queer? I think not, we are perfectly familiar with this phenomenon. This follows from the fact that moral properties are prescriptive and *real*.

However, if these properties also provide us with an overriding motive, then this is different. If this is the case, then it is queer. It is very queer indeed to believe that the fact that an action ought to be done should make me prone to perform it. This would be queer if for no other reason that I may be ignorant of its normative status. But even the idea that my *belief* that an action is obligatory should, as such motivate me, would be queer, I submit. But why believe in this kind of internalism?

Let us return to the satisfaction values and truth-values of a norm and consider them more in detail. Consider a simple example: the circumcision of a young girl. Let us assume that one person claims that the circumcision of this girl ought to take place. This claim is true, on my objective understanding of it, if, and only if, the girl in question ought to be circumcised. This is a way of explaining the truth-condition of this normative judgement. But the same judgement has a satisfaction-condition as well. The judgement that the young girl ought to be circumcised is satisfied if and only if the young girl *is* circumcised.

Most people have an *interest* in having their normative judgements (about obligations) satisfied. We have been socialised to have this interest. This means that many moral discussions have a practical dimension. When we disagree about the truth conditions of normative judgements we also have a practical conflict about the satisfaction conditions of the very same judgements.

If you are interested in having a certain girl circumcised (you believe that she ought to be circumcised and you want to have this norm satisfied) while I believe that she ought not to be circumcised (and want this norm to be satisfied), then we face a practical conflict. This captures the grain of truth in Mackie's thought that normative properties are motivating.

However, not all normative judgements have necessary satisfaction-conditions in this straightforward manner. In order to see which judgements do, and which judgements don't have both sufficient and necessary satisfaction-conditions, we may account for the normative categories in the following manner. An action is right if it is not wrong. An action is wrong if it is not right. An action ought to be performed if and only if it is wrong not to perform it. As we saw above, the first two claims about the relations between actions being right and wrong are not *logical* truths, but for a moral realist, it is natural to assume that they are *true at our world and at worlds that are similar to ours*. Obviously, judgements about obligation (you ought to do so and so) and about wrongness (it is wrong to do so and so) have satisfaction-conditions (the latter is satisfied if, and only if, you do not perform the action in question), but judgements about rightness (it is right to do so and so) do not have any necessary satisfaction conditions. They only have sufficient ones, and they can be viewed as *permissions*.

⁸ Ibid. p. 40.

Suppose you ought to F, then your obligation is satisfied if you F otherwise it is not satisfied. This prepares the room for the following kind of practical conflict. If you believe that you ought to F and want to do what you believe you ought to do, then you want your obligation to be satisfied. Suppose you come across a person who believes that you ought not to F. This person too wants to see to it that what she believes to be obligations be satisfied. She wants you not to F. So she has a practical conflict with you. There is no way to satisfy both your conflicting (and inconsistent) norms.

Here is another example. One person believes that it is right to F. Another person believes that it is wrong to F. According to the latter person, it is obligatory not to F. If you F, then you are in conflict with this person. He does not get his favoured obligation satisfied. However, if this person, who believes that it is wrong to F, does not F, you (believing that it is right to F, need not have any problem with this. And if, for some reason, you don't F, the person who believes that it is wrong to F, is perfectly happy with this. So not all intellectual moral conflicts translate easily into practical ones. However, many do, if not directly, so at least indirectly, and I believe that this is this phenomenon that Mackie has conflated with the queer property of an inherently motivating force of moral judgements.

Furthermore, not all people want to have their favoured norms (the norms they believe are true) satisfied. We all know a kind of person (we know him or her from our philosophy seminars) who is exclusively interested in the truth-conditions of various different norms. He or she doesn't give a damn about the satisfaction of these norms. We may look askance at such a person and yet, for all that, this kind of person does exist. But the mere possibility that there are such persons — and we may safely infer that they are possible, once we have met with them — shows that there is no inherently motivating force in moral properties, in Mackie's sense — i.e. no 'intrinsic to-be-pursuedness' — built into them. They give us *direction*, all right, and this is what it means for them to be 'prescriptive', i.e. they have satisfaction-conditions — but they do not *move* us to action.

But even if our moral (normative) judgements are objective, having both truth-conditions and (some of them) both necessary and sufficient satisfaction-conditions, and even if we believe that *most* people, *most* of the time, are interested in having their normative judgements satisfied, why should we believe that, by issuing such judgements, we hit upon anything existing objectively, to be captured by our judgements? Why not believe, as Mackie did, that all moral judgements are false?

Well, if there is nothing strange with the prescriptivity of moral (normative) properties (we are perfectly familiar with the phenomenon and we can explain it with reference to satisfaction-conditions), why think that they do not exist? Why assume that our moral judgements do not correspond to anything in reality?

In his argument from queerness Mackie adds another respect in which he finds that objective moral properties are queer, the fact that they are 'consequential' or 'supervenient':

Another way of bringing out this queerness is to ask, about anything that is supposed to have some objective moral quality, how this is linked with its natural features ... The wrongness must somehow be 'consequential' or 'supervenient'; it is wrong because it is a piece of deliberate cruelty. But just what in the world is signified by this 'because'? ... It is not even sufficient to postulate a faculty which 'sees' the wrongness: something must be postulated which can see at once the natural features that constitute the cruelty, the wrongness, and the mysterious consequential link between the two.⁹

This is most plausibly taken as two different objections, one to do with the word 'because', and one directly concerned with supervenience. I will start with the query about the word 'because', which seems to be Mackie's main worry. Is it difficult to understand what is referred to by this word? I think not. The answer to Mackie's query what in the world is signified by this 'because', is given with reference to the fact that in the world there are true moral principles, capable of explaining (together with certain facts) particular obligations. And these principles are not different with respect to the necessity implicit in them, than laws of nature. They state right-and wrong-making conditions. Our notion of objective moral reasons hence helps us to answer Mackie's argument.

The answer to his more general question about supervenience must be given with reference to some metaphysical principle. I once toyed with the idea that we should give up supervenience altogether. Two worlds could be similar in all non-moral respects, and yet be different in moral respects. I happened to discuss this idea, many years ago, with John Mackie. He was not convinced. As a matter of fact, he was angry with me! He could not understand how on earth I could hold such a bizarre view.¹⁰ And I realise that if moral properties do not in any way supervene upon non-moral properties, it is hard to see how we could have any knowledge about them. If we jettison supervenience, we end up as moral sceptics.

Now I tend to believe that the most natural response from the moral realist to Mackie's argument would be, not a denial of supervenience, but a denial of the claim that supervenience as such is any strange phenomenon.

⁹ Ibid. 41.

¹⁰ It must have been around 1980, when he visited our philosophy department in Stockholm.

We know it from many other parts of our human understanding. The mental seems to be supervenient upon our neurophysiology, causal relations seem to presuppose supervenience, and so forth. The details of this thesis, however, are far from clear.

There is a strong version of supervenience, according to which moral principles are necessarily true. So what makes something right or wrong in the actual world makes it right or wrong in all possible worlds. This claim may appear to be queer. Can the claim that moral properties are both *sui generis* and supervenient upon natural properties in this sense be upheld?

Frank Jackson has put forward an argument with reference to Simon Blackburn to the effect that this is not possible. I here answer Jackson's concern. I am at loss with respect to Blackburn's own version of the argument, since I cannot understand why he finds it difficult to explain why there cannot be a 'mixed' world, where sometimes a naturalistic characteristic is right-making, and at another time not right-making. That this is impossible follows from an ordinary understanding of what a moral principle (connecting the two properties) is. Unless it holds without exceptions in a world, and with counter-factual implications, it is no principle (the same is true of putative laws of nature).¹¹

The view that moral properties are both *sui generis* and supervenient upon natural properties can be upheld, if we adopt a fairly weak view of supervenience, according to which a certain natural property brings with it a corresponding moral property only in all possible worlds that are sufficiently like ours. This seems to me a fairly plausible metaphysical view. I do not want to rule out the possibility that there are possible worlds where different laws of nature and morality obtain. Even if utilitarianism is actually true, or close to the truth, there may be worlds where Kantianism is true. There may even be worlds where moral nihilism is true, and there may be worlds where moral particularism is true (worlds where particular actions are right or wrong but where there are no true moral principles). If I am right about all this, then this means that the moral realist is under no pressure to explain how it is possible that certain moral properties and certain natural properties co-exist in all possible worlds and are yet, for all that, distinct. We have no reason to believe in the claim of co-existence in all possible worlds in the first place. We may even concur in the assessment that this claim is *queer*.

On a weak notion of supervenience, we can accept the claim that if a world is exactly like ours in all natural respects, it must also be like ours in moral respects. And this weak notion of supervenience is consistent with the claim that, in each world, if there is a moral difference between two cases, there must also be a natural (non-moral) difference.

Why, then, should we not believe in strong supervenience? Why should we not believe that true moral principles are necessarily true (true in all possible worlds)?

The rationale behind the belief that moral principles have (merely) the same kind of necessity built into them that natural laws have, helps us not to have moral realism collapsing into naturalism. And the reason to believe in it is the observation that moral principles play the same role in moral explanations that natural laws play in scientific explanations.

Furthermore, as will be noted below, we seem to arrive at our belief in moral principles in a manner that is similar to how we arrive at our belief in laws of nature, i.e. through a hypothetical deductive method of investigation.

Hence, if laws of nature can differ between possible worlds so could moral laws (principles). This view of the kind of necessity exhibited by moral principles allows us to acknowledge that, even if moral realism is true, it could have been false.

Finally, there are views in normative ethics inconsistent with strong supervenience. I think for example of actualism in population ethics, the view that only the fate of actual people matter to the moral rightness and wrongness of actions. On this view, an action performed in the actual world may be wrong while, at the same time, there is a possible world, from the point of view of which the very same action is right. This view may sound strange, I believe it is false, but it does not strike me as incoherent.

This standard and rather weak notion of supervenience, combined with the idea of (merely) a weak necessity built into moral principles, implies that there is a possible world where it is all right to torture an innocent child for no reason at all. I suppose that this world is very different from ours, but my imagination cannot help me to a concrete understanding of what it would be like. Moreover, when this is said, it is crucial that we note that this talk of possible worlds is a mere metaphor. No existential claim with respect to merely possible worlds is here being made. If the metaphor is taken literally, if we claim that all possible worlds exist, then we are in deep moral trouble. I have discussed this elsewhere,¹² so I will not elaborate on this point in the present context.

So much for the standard notion of supervenience. To go for strong supervenience, according to which moral and natural properties are coextensive in all possible worlds, threatens, of course, to have one's view collapsing into the naturalistic claim of the Cornell variety that moral properties are identical to some natural properties. I will say more in the next chapter about why we should avoid moral naturalism in all its varieties.

¹¹ See Blackburn, 'Supervenience Revisited'.

¹² In my 'The Moral Import of Modal Realism'.

To sum up, moral (normative) properties are not queer. They are not queer in the sense that they can directly make us want or desire certain things (internalism). The answer to internalism is externalism: they don't have this kind of motivational force built into them. They are not queer because they have prescriptivity built into them either. They do have prescriptivity built into them, but we know this phenomenon quite well. It has to do with the fact that normative judgements have both satisfaction-values and truth-values. We are familiar with, and can specify, criteria of both truth and satisfaction for any particular norm. The notion of supervenience associated with moral properties, may well involve us in deep metaphysical problems, but we are familiar with the very notion from many other areas of science. Finally, and most importantly, we know what it means to say that an action is right 'because' so and so. By saying this, we imply that there are true moral principles helping us to explain its rightness.

Of course, the particularists may be right when they insist that there are no moral principles. We still know what it means to say that it is 'because' of the existence of a true moral principle that something is right or wrong. However, if the particularists are right, all such claims are false. Moreover, if, as I think, the particularists are wrong, and there are indeed true moral principles, then there are such claims that are not only meaningful, but also true. On the other hand, when the particularists speak of actions being right or wrong 'because' so and so, Mackie's argument from queerness does seem to stick. They (the particularists) cannot explain to us how they use the term 'because'. But that is their problem, not mine.

So why believe that moral properties are projections of ours? It doesn't seem as though Mackie's argument from queerness has forced us to accept that conclusion. But Mackie's has also put forward another argument, to the same nihilist effect, his argument from relativity.

3. THE ARGUMENT FROM RELATIVITY

The argument from relativity is supposed to show that our positive moral judgements, pretending to be objective, are all false. This is how he states it:

The argument from relativity has as its premiss the well-known variation in moral codes from one society to another and from one period to another, and also the differences in moral beliefs between different groups and classes within a complex community ... radical differences between first order moral judgements make it difficult to treat those judgements as apprehensions of objective truths.¹³

This argument is problematic in many ways, however. Mackie himself is well aware of one of them. Even if this disagreement exists, this is not different from what we see in sciences such as history, or biology, or cosmology, he concedes. We do not, because of this disagreement, conclude that there is no objective truth in history, biology, or cosmology. But there is a difference, Mackie claims:

... such scientific disagreement results from speculative inferences or explanatory hypotheses based on inadequate evidence, and it is hardly plausible to interpret moral disagreement in the same way. Disagreement about moral codes seems to reflect people's adherence to and participation in different ways of life.¹⁴

Could this be correct? I do not deny that many moral judgements are, in this way, conventional. However, this is consistent as well with the possibility that some people have been able to transcend their actual conventional moral universe and pose the more fundamental questions about right and wrong action. This is the possibility I want to defend.

Here it is not far-fetched to assume that, while there is much disagreement about speculative moral principles, intended to 'explain' particular moral judgements, just as there are disagreements in history, or biology, or cosmology, there is also much agreement. There are indeed moral facts, and, if we succeed in taking up an impartial stance, we are sometimes capable of recognising them.

All decent and impartial people, who consider the question, would agree that it is wrong to torture an innocent child, at least unless there are any pressing reasons of any sort to do so. The latter clause is elliptical and though of as taking care of objections such as: 'But by torturing this child you make the world a better place in the long run', 'But by torturing this child you do what you have undertaken to do', and so forth. If we come across a person, who doesn't share our judgement about this case, we would say that he doesn't understand what it means for an action to be wrong — he has not been able to take up a critical stance with respect to the norms that have been inculcated in him, or he suffers from some serious personality disorder.

¹³ Ibid. p. 36.

¹⁴ Ibid.

Perhaps there are only a few examples of this kind that can be produced, and perhaps they are a poor foundation for speculative inferences or explanatory hypotheses. In that case, we may here see the explanation of why people who have thought about them, and tried to find general explanations of them, have moved in different directions. And yet, for all that, these cases exist. And we may add that the kind of transcendental and distanced approach to morality that I am here speaking of is perhaps a rather new phenomenon. As Derek Parfit has insisted, 'both human history and the history of ethics may be just beginning'¹⁵

It should be noted that when I claim that it is *obvious* that it is wrong to torture an innocent child 'for no reason', this additional 'no-reason' clause is crucial to my argument. For according to some ethical theories, held by clever people who have thought hard about these things, it is indeed all right to torture an innocent child – if the consequences are good enough to compensate for the suffering felt by the innocent child. I have myself argued that it might be right to torture an innocent child, if this means that a great many people, who are all of them already very happy, will undergo a sub-noticeable increase in their respective hedonic situation.¹⁶ This is a controversial view, of course. All decent and clever people who have thought hard about it do not agree. Derek Parfit, for one, has nicknamed my idea 'the Ultra Repugnant Conclusion'. The fact that he and many others disagree means that I have to concede that I *may* be wrong about this conclusion. However, this is not enough for me to give it up. After all, this conclusion follows from, and is explained by, a theory (hedonistic utilitarianism) for which I think there exists adequate evidence. Yet, for all that, this conclusion is far from self-evident, in the way it is self-evident that it is wrong to torture an innocent child for *no reason whatever*. I find it hard to see that there is *any* way that I could come to give up this belief – unless I would become morally perverted or go insane.

But is it not because we have been socialised to believe that this is wrong, that we find it wrong to torture an innocent child for no reason? This is the main point in Mackie's argument from relativity, even if it has little to do with relativity as such, and everything to do with the *origin* of our moral beliefs.

4. HARMAN'S EMPIRICIST ARGUMENT

It is a remarkable fact that the argument I have here attributed to Mackie, to do with how we should *explain* our moral beliefs, was also put forward by Gilbert Harman, in the opening chapter of his book, *The Nature of Morality*. Mackie's and Harman's books were published the same year, in 1977. I now turn to this argument. It should first be noted, however that, this kind of argument was alluded to already by R.B. Brandt when, in a footnote to his *Ethical Theory*, he compares the role of intuition or feeling in morality to observation in science and makes the following qualification:

There is one difference. Physical theory, taken with a description of the experimental setup, may logically imply 'The ammeter will point to 30' and we can observe whether or not this is the case. Whereas, although in ethics we may reject 'There is no obligation to do X' by appeal to the fact that we feel a strong obligation to do X ... we cannot say that ethical principles entail anything about how we shall feel — at least not in any direct way.¹⁷

According to Harman's more elaborated development of the argument, it relies on an explicitly stated empiricist principle to the effect that we ought not to countenance something as real unless we need to have recourse to it in the best explanation of some observation. And, according to Harman, when we explain why people make moral observations we never need to have recourse to moral facts. It is sufficient (we get a better explanation) if we refer rather to psychological facts to do with the agent, such as his moral sensibility, and so forth. This is the somewhat more elaborate way in which Harman makes his point:

... observation plays a role in science that it does not seem to play in ethics. The difference is that 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 ... 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. In the scientific case, theory is tested against the world.¹⁸

It might seem that the assumption upon which Harman's argument relies is problematic. It seems *itself* to be normative. Harman assumes that we should not acknowledge something as real, unless we have to refer to it in our best explanation of some judgement we make (such as the report of an observation). But as will be shown in Chapter 8, we may think of rules such as this one as stating mere Humean, instrumental norms. They are to the

¹⁵ *Reasons and Persons*, p. 453.

¹⁶ See my *Hedonistic Utilitarianism*, p. 77.

¹⁷ R.B. Brandt, *Ethical Theory*, p. 249n.

¹⁸ *The Nature of Morality*, p. 6.

effect that, if you want to further an epistemic goal you hold, if you want to have a realistic picture of the world, then you ought to abide by the rule. This is how I will understand it. Under this interpretation I accept it. The moral realist should bite the bullet and answer Harman's (and Mackie's) challenge, then. So it is pertinent to find out if Mackie is correct when he claims that the moral codes we adhere to reflect our various different ways of life. And we need to know whether Harman is right when he claims that our moral observations depend on our moral psychology, not on moral facts.

These assumptions seems to me *extremely* far-fetched, when applied to the example we *now* discuss. It is more plausible to assume that we take it to be wrong to torture innocent children for no reason whatever because it *is* wrong. And it is very plausible to assume that we have been socialised into believing that this is wrong because it *is* wrong.

Can we not find a better, evolutionary explanation, why we hold this belief, than concluding that we have been able to grasp its truth?

In general, if we can find an evolutionary explanation of most our moral beliefs, this means that we should become agnostic with respect to their truth. This means less of a trouble than if Mackie is right in his insistence that normative notions are queer. If he is right, then all our moral beliefs are *false*. It is also less serious, than if expressivists of Stevenson's and Gibbard's variety would succeed in showing that all our moral expressions are used, not to express propositions with truth-values, but our emotions or something of the kind. If they were right, then our moral beliefs would lack cognitive content altogether. And yet, an evolutionary explanation would be a 'debunking' one indeed. It would show that we should be less confident in our moral beliefs. We would realise that we are not justified in holding them (even if they *make sense* and even if they *may* be true). Moreover, if we could find such an evolutionary explanation it should not only make us agnostic with respect to the truth of our deeply cherished opinions. We should also, as is stressed by Richard Joyce, become agnostic with respect to the moral notions (concepts) we employ.¹⁹

But all his means no problem automatically for the moral realist. Quite to the contrary, such a suspicion about conventional morality is really the point of departure for the realist. Now, the suspicion about the moral notions will naturally lead to a close examination of them. And even if this examination somewhat unexpectedly leads to the conclusion that the moral beliefs we hold are indeed true, since we are to understand them as elliptic, making an implicit reference to the system of norms upheld in our society, this is only the beginning of further queries. For now it may transpire that the moral realist is not at all satisfied with conventional morality, nor with notions of right and wrong action used in her society. Then a kind of transcendental move is natural to expect. The words 'right' and 'wrong' take on a new, absolute sense, and critical moral views are put forward with the aid of these notions. And the debunking explanation of the content of common sense morality is now a reason to bracket these beliefs, in our pursuit for the truth. Here our cognitive and emotional capacities, provided for by evolution, are put to use.

How, then, are we to think of these new critical moral views, partly in conflict with conventional morality? Well, to the extent we conceive of them as the result of progress in our moral thinking, it is impossible to just set them to one side. We treat them as true in an absolute sense. We take them to track moral facts.

But can we not once again try to explain, and hence debunk, the content of these beliefs with reference to our biological evolutionary history? I think not; at least with respect to some of them this would seem far-fetched to assume. This is true already of the idea that one should not torture an innocent child for no reason. It is even more obvious when it comes to the recent insight that the same is true, not only of children, but of all sentient creatures, human or non-human.

We learn from sociobiology that we can be expected to behave altruistically towards those who share our genes. We also learn that we can be expected to have a capacity for cooperation, conditional altruism, and mutual accommodation of our interests. If such behaviour is called moral behaviour there exists an evolutionary explanation why it exists. However, it is hard to see why evolution should have put altogether impartial moral ideals into our heads. It is more reasonable to assume, I submit, that evolution has provided us with a capacity for understanding that allows us to grasp some basic moral truths.

In the final analysis, we may also turn the plausibility of moral claims themselves against various nihilist or agnostic arguments in a manner adopted by A.C. Ewing, Ronald Dworkin and Thomas Nagel.²⁰ But this move *presupposes* the truth of an objectivist analysis at least of *our* use of moral language, of course. If our moral opinions are mere emotions, the nihilist claim that there are no moral facts existing independently of our conceptualisation and thought does not contradict any one of *our* moral opinions.

It might be tempting to argue that *general* moral facts such as the one that it is wrong to torture an innocent child for no reason cannot be grasped, since they lack location in space and time. To be in 'contact' with such a fact must be utterly different from, say, observing that there is a cat in front of you. And the *content* of the act of

¹⁹ Richard Joyce, *Evolution of Morality*, p. 181.

²⁰ This can be seen as the main thrust of Nagel's book *The Last Word*. A similar move is made in Ronald Dworkin, 'Objectivity and Truth: You'd Better Believe It'. And long before them, this kind of argument was used by A.C. Ewing in *The Definition of Good*, pp. 32-33.

apprehension, to wit, *that it is wrong to torture an innocent child*, can thus not form part of any explanation of why you come to hold the belief.

But this argument is not convincing. It is certainly true that the fact that it is wrong to torture an innocent child lacks location in space and time. But *your contemplating this putative fact* does take place at a certain place and time. And part of the explanation why you end up believing this to be a fact might be that it is a fact (together with the fact that you are an ordinary receptive and clever person).

Compare when you come to believe that $5+7=12$! Harman claims that this cannot be explained with reference to the fact that $5+7=12$:

Observation does not seem to play the role in mathematics it plays in physics. We do not and cannot perceive numbers, for example, since we cannot be in causal contact with them. We do not even understand what it would be like to be in causal contact with the number 12, say. Relations among numbers cannot have any more of an effect on our perceptual apparatus than moral facts can.²¹

This does not make Harman a nihilist with respect to mathematics. But, on his account, mathematics is saved indirectly, so to speak. We need maths when we state our theories of physics. But a more straightforward argument is available to the mathematical realist (Platonist). At the right time and place in your early mathematical education, this truth, that $5+7=12$, may well dawn upon you. And the truth of it forms part of a causal explanation of the origin of your belief. Karl Popper has noted this remarkable fact:

... the human mind can see a physical body in the literal sense of 'see' in which the eyes participate in the process. It can also 'see' or 'grasp' an arithmetical or a geometrical object; a number, or a geometrical figure. But although in this sense 'see' and 'grasp' is used in a metaphorical way, it nevertheless denotes a real relationship between the mind and its intelligible object, the arithmetical or geometrical object; and the relationship is closely analogous to 'seeing' in the literal sense.²²

This is how a Platonist about numbers *should* argue. This is how a consistent moral realist of a non-naturalist slant *must* argue, with reference to moral realities, as well.

Finally, assuming that there are objective moral (normative facts), making our moral (normative) judgements true or false, how can we ever pretend to have *knowledge* about these facts? Is there any reason to believe that one opinion rather than another about these facts is any better founded? Can we have *justified* moral (normative) beliefs without using any faculty of moral perception or intuition 'utterly different from our ordinary ways of knowing everything else', to use Mackie's words?

5. MORAL KNOWLEDGE

True and justified moral belief (knowledge) is possible. Let me here just indicate how I think we obtain *justified* moral (normative) beliefs. We *can* reach justified moral belief. When we do, and when our justified beliefs are true, we do possess moral knowledge. We go about in roughly the following way, using our ordinary emotional and intellectual capacities, produced by the evolutionary process.

We grow up in a society where a certain moral code is taken for granted. At first we concur in taking it for granted. Then some of us come to a point where, from a detached perspective, we come to question what the code urges us to do. We become suspicious of common sense morality, not entirely, certainly, but partly. We find that there are many considered moral beliefs that we are hard put to give up, however, such as the one about not torturing innocent children for no reason. But there are also other beliefs, taught to us, that we find it hard to believe. If we are of a philosophical bent, we try to articulate speculative hypotheses explaining the truth of the judgements that strike us as plausible and explaining the falsity of those we find suspect. When articulating such speculative moral principles, we go for those that are simple, general, fruitful, and so forth. We select the one or the ones that we claim to give the best explanations of (the content of) our more particular judgements.

We also put our particular judgements, our moral intuitions, to scrutiny. We search knowledge about how we may have obtained them. There is now much evidence from psychology and even from neuroscience, including the results of scans of our brains, to turn to. We submit our intuitions to cognitive psychotherapy and retain only those who survive our knowledge of how they may have come to exist, as still subjectively justified. We should rely mainly on those that have been generated by our cognitive capacities, I submit. We should rely on them, even when they run counter to our moral 'gut feelings'; in particular we should do so if we can understand why it is rational to have these gut feelings, even if, now and then, they lead us astray.

²¹ *Ibid.*, pp. 9-10

²² *Objective Knowledge*, p. 155.

In the famous trolley cases, we should rely on our rational insight to the effect that it is better to save lives than to sacrifice lives, even when it means that we have to kill in order to do so. The gut feeling against active and physical killing is rational, however; it is rational in the sense that, in *most* cases, it is morally correct to rely on it. It is rational in the same sense that it is rational to fear all snakes, irrespective of whether they are poisonous or not. By fearing all we err on the right side and we stay alive. Hence we ought to cherish this kind of primitive attitude in our *lives*. And yet, in rare cases it is wrong to abide by it, so we should not count on it in an *epistemic* context. Our knowledge of its origin and function here has, and should have, a debunking function. It is not to be relied on in our pursuit for the moral truth. This is similar to our emotion with respect to snakes. If we want a correct classification of them, into those that are poisonous and those that are innocuous, we should not rely on our gut feelings (our primitive emotions).

In this manner we take (the content of) our considered moral judgements or intuitions as evidence for the more general principles. It is crucial here, of course, that we have not simply consciously *inferred* our more particular beliefs from our general ones, supposed to explain them.²³

The parallel to a scientific investigation is obvious. The process sometimes goes the other way round. We find that we have to let go of some more particular and considered moral judgements, in spite of the fact that they have survived cognitive psycho-therapy, since they conflict with principles that we have found have a strong support (they are supported by other particular considered judgements, as well as by methodological considerations; in addition to this, they may just feel intuitively right to us). When the process ends we have reached what John Rawls has famously called a 'reflective equilibrium.' And there seems to be genuine progress in this process, allowing us to think that we do come closer to the truth.

Deontological ethics takes our gut feelings seriously. Utilitarianism takes them into account too, but gives them a different, secondary place, in a two level approach to ethics. In my opinion, this means that utilitarianism makes a theoretical advance over deontology. It explains the content of our considered intuitions better than deontology, it contradicts deontology, and it explains the *relative* success of deontology.

In moral philosophy we are used to focus on disagreement, but there exists a growing common moral understanding, exemplified with the growing understanding that non-human sentient beings have a moral standing. I would even go so far as to claim that there is something close to a general recognition, that there must be some truth in utilitarianism, if utilitarianism is seen as more of a research program than a definite doctrine, and if it is acknowledged that there *may* be circumstances where utilitarianism may give the wrong answer (such as situations where its demands are overly demanding, where they mean that there is conflict with deontological constraints and individual rights, and so forth). In those cases where there are no such problematic aspects few would have any problems with adhering to the utilitarian solution of the moral problem. Few are prepared to argue that it is right to make the world a less happy place, unless there are any pressing reasons to do so. Those who are prepared to do so are simply wrong.

What are we to say when we find that people, who have thought equally hard about these matters, as we have done, have come to different conclusions? Does this mean that we must give up our endeavour?

Now, if they think differently about the torture example given above we are unperturbed by this finding. We set their opinions to one side; their moral opinions are just bizarre, and we try to find an explanation to do with some disorder or distortion within the dissenting person.

What then if the dissenting party agrees with us about a hard core of 'simple' and obvious moral judgements, but disagrees completely in the 'organisation' of morality? This person believes in moral principles that we reject. And we find no obvious flaws in these people's emotional or intellectual characteristics. They are at least as intelligent and sympathetic as we are. Does this mean that we give up our favoured moral principles, arguing that it would be far-fetched to assume that we, rather than those competent and decent people with whom we are in disagreement, should have reached moral knowledge?²⁴

Obviously, we don't, and we shouldn't. To the extent that we come across such disagreement, where we cannot explain in any simple manner why it occurs, we become more humble. We realise that our own way of accounting for what we believe is a uniquely true morality may have gone wrong in many places. This is something we may come to find out in the future. However, each piece of our moral understanding has the argumentative support it has, so it would be irrational for us, given our epistemic goal (to be explained in Chapter 8), to wit, to believe true propositions and disbelieve false propositions, to give it up merely because other people have reached other results.

²³ I elaborate on this idea in *Understanding Ethics*, Chapter 9. For literature about the results from scientific studies of the origin of moral intuitions, see for example Greene et al, 'An fMRI Investigation of Emotional Engagement in Moral Judgment', and Greene and Haidt, 'How (and Where) Does Moral Judgment Work?'. The literature on this subject is rapidly growing.

²⁴ For a discussion of this possible reaction, see George Sher, 'But I Could Be Wrong'.

6. CONCLUSION

I do not pretend to give any knock-down argument in this chapter for the thesis that objective moral facts or reasons exist, independently of our thoughts and actions. My main positive argument in defence of moral (normative) realism is the observation that it seems as though our making of some moral judgements, once we have been able to transcend conventional morality, is best explained with reference to moral facts. This is a rather weak support. But it is the kind of support one can give for a thesis of this kind. And I do not think any stronger support is really needed. I tend myself to agree here with Thomas Nagel when he wants to show that pain and pleasure provides us with agent-neutral reasons, and starts his argument by saying:

In arguing for this claim, I am somewhat handicapped by the fact that I find it self-evident.²⁵

As was also noted above with reference to the same author, we may also turn the *content* of our moral beliefs — in form of particular judgements as well as in form of theories, and in form of the support we have gathered for them — against attempts to show a priori that they are false (moral nihilism) or that we must be agnostic with respect to them (because of a putative debunking explanation of why we hold them). Incidentally, I am prepared to go even further and claim that, unless moral facts did exist, independently of our way of conducting our lives and passing moral judgements, life would lose ‘meaning,’ in an important sense.²⁶

This final claim is of course no argument in defence of the *truth* of moral realism, or even an argument why it is rational to *believe* that it is true; or if it is after all a reason for a person who like me *wants moral realism to be true in order to find meaning with life*, to believe that moral realism is true, this reason has nothing to do with *epistemic* justification, of which more will be said in Chapter 8.

However, the fact that moral realism can give meaning to our lives does indeed indicate why the discussion of it is of the utmost importance.

²⁵ *The View from Nowhere*, pp. 159-60.

²⁶ I defend this claim in the concluding chapter of my book *Moral Realism*.

