MORAL INTUITION, MORAL THEORY, AND PRACTICAL ETHICS

<u>Moral Inquiry</u>

Suppose we wish to understand a particular moral problem - for example, abortion. How should we proceed? One approach, which I favor, is to reason on the basis of our existing substantive moral beliefs. We may, however, suspect that our moral beliefs about abortion, insofar as we have any prior to serious reflection, are unreliable. We may suspect, for example, that our beliefs about abortion reflect the influence of a religious education that we now repudiate, or that our feminist sympathies may make us insufficiently sensitive to the status of the fetus. Thus we may take as our starting point certain related moral beliefs about which we are more confident: for example, that killing a morally innocent adult human being whose continued life would be worth living is, except perhaps in the most extreme circumstances, seriously wrong, while painlessly killing a lower nonhuman animal (for example, a lizard) may often be permissible provided that the interests that are thereby served are reasonably serious. There is, of course, divergence of opinion even about these cases. Some people believe that killing an innocent adult can never be justified, while others believe that it can be justified whenever it is necessary to save the lives of a greater number of innocent people. And, while some believe that there is no objection whatever to killing a lizard independently of the effects this might have on human interests, others believe that killing a lizard is seriously objectionable just because of the effect on the lizard and requires a strong justification in order to be permissible. Nevertheless, everyone agrees that killing an innocent person is immeasurably more objectionable morally than killing a lizard, other things being equal.

We could therefore initiate our inquiry into abortion by exploring our confident sense that there is an enormous moral difference between killing people and killing lower animals - so that, for example, the killing of a lower animal might be justified by appeal

to considerations that would not constitute even the beginning of a justification for killing a person. With these beliefs as our starting point, we could work our way toward a better understanding of abortion. We could proceed by trying to understand why killing people is generally wrong and why it is generally so much more seriously wrong than killing lower animals. What are the relevant differences between normal adult human beings and lower animals? Are the properties of persons that make killing them generally worse all intrinsic properties? Or is part of the explanation of the greater wrongness of killing people that we bear certain relations to other people that do not exist between ourselves and animals? In addressing these questions, we may consult our intuitions about a range of particular cases and this may yield provocative results. We may notice, for example, that the extent to which an act of killing an animal seems wrong varies with the degree of harm the animal suffers in dying. Thus it seems more objectionable to kill a dog than to kill a lizard; and the explanation seems connected with the fact that the dog loses more by dying. But we may also notice that the extent to which it is wrong to kill a person does not seem to vary with the extent to which death is bad for the victim. Thus it seems no less wrong, other things being equal, to kill a dullard than to kill a genius, or to kill an elderly person with a reduced life expectancy than to kill a person in the prime of life. Further inquiry is necessary to understand the significance of these findings.

As our understanding of the morality of killing in general increases, we can begin to extract from our findings various implications for the morality of abortion. Suppose, for example, that we discover that there are certain properties that adult human beings generally possess that lower animals do not that seem to help explain the difference between killing people and killing animals. We can then consider whether these properties are possessed by human fetuses. If they are, then in that respect abortion is relevantly like killing an adult human being; if not, there is then reason to suppose that abortion should instead be assimilated to the paradigm of killing animals.

These remarks about abortion are intended only to provide a sketchy illustration of a certain approach to practical ethics, a certain general pattern of reasoning about moral problems. Its most conspicuous feature is that it treats certain substantive moral beliefs that we already have as reliable starting points for moral inquiry. It presupposes that at least some of our moral intuitions have a certain prima facie normative authority.

Intuitions

What are moral intuitions? As I will understand the term, a moral intuition is a spontaneous moral judgement, typically about a particular problem, a particular act, or a particular agent, though there are intuitive apprehensions of moral rules and principles as well. In saying that a moral intuition is a spontaneous judgement, I do not mean to suggest that it is evoked instantaneously, the way a sense perception is, when one is confronted with a particular case. If the case is complex, one may have to consider it at length in order to distinguish and assimilate its various relevant features - in much the same way that one might have to examine the many details of a highly complex work of art in order to determine whether one likes it or not. Rather, in saying that an intuitive judgement is spontaneous, I mean that it is not the result of inferential reasoning. It is not inferred from one's other beliefs but arises on its own. If I consider the act of torturing the cat, I judge immediately that, in the circumstances, this would be wrong. I do not need to consult my other beliefs in order to arrive at this judgement.

The belief that I cited as one of the possible starting points for an inquiry about abortion - namely, that killing adult human beings is generally wrong - may not or may not count as an intuition according to this understanding, though for most of us it has an intuitive basis. It counts as an intuition if one finds it immediately compelling but not if one accepts it as an inductive inference from one's intuitively finding that in this, that, and the other case, killing an adult human being is wrong.

In the history of moral philosophy, the idea that moral intuitions have normative authority has been associated, unsurprisingly, with a cluster of theories that have traveled

under the label 'Intuitionism.' Those doctrines are many and various and I do not propose to disentangle them. But two claims associated with certain historically prominent variants of Intuitionism have done much to discredit the appeal to intuitions. One of these is that intuitions are the deliverances of a special organ or faculty of moral perception, typically understood as something like an inner eye that provides occult access to a noumenal realm of objective values. The other is often regarded as a corollary of the first - namely, that intuitions are indubitable (that is, that their veridicality cannot be doubted) as well as infallible (that is, that they cannot in fact be mistaken). But a variety of considerations - such as the diversity of moral intuitions, the fact that people do often doubt and even repudiate certain of their intuitions, and the evident origin of some intuitions in social prejudice or self-interest - make it untenable to suppose that intuitions are direct and infallible perceptions of moral reality.

There are other features that are occasionally attributed to intuitions that are in fact inessential. It is sometimes said, for example, that intuitions are 'pretheoretical.' If all this means is that they are not derived inferentially from a moral theory, then it of course follows from the stipulation that they are not the products of any sort of inferential reasoning. If, however, it means that intuitions must be untutored or entirely unaffected by a person's exposure to moral theory, then the requirement is evidently too strong. Just as many people's moral intuitions have been shaped by their early exposure to religious indoctrination, so some people's intuitions are gradually moulded by their commitment to a particular moral theory. [Implications for the idea that core values are innate?]

Theory

Many philosophers reject the idea that moral intuitions have normative authority. Peter Singer, for example, suggests that we should assume 'that all the particular moral judgments we intuitively make are likely to derive from discarded religious systems, from warped views of sex and bodily functions, or from customs necessary for the survival of the group in social and economic circumstances that now lie in the distant

past.' On this assumption, he notes, 'it would be best to forget all about our particular moral judgments.' (Singer 1974, p. 516) It is, of course, possible to be rather less dismissive of intuitions and yet still regard them as lacking in normative authority. Some philosophers, for example, concede that intuitions may be tolerably reliable guides to action in most circumstances (since morality must ensure that people are equipped with dispositions to believe and act in certain ways in situations in which deliberation and reflection are not possible) but deny that they are a source of moral knowledge or have any proper role in reasoning about moral problems. They believe that practical reasoning about a moral problem must consist in determining what some favored moral theory implies about the problem. It is the theory that is the source of our moral knowledge concerning particular problems and cases. And the theory is itself validated by means other than its conformity with our intuitions.

According to this approach, if our concern is to understand the morality of abortion, our first task must be to discover the correct moral theory. Moral inquiry is initially and primarily theoretical; only at the end of the day is it possible to consider moral problems such as abortion, bringing the theory to bear and extracting from it the knowledge that we initially sought. This general approach therefore contrasts markedly with the first approach I sketched, according to which moral inquiry initially focuses intuitions about on problems and cases rather than on matters of abstract moral theory. According to the first approach, a moral theory about which we are entitled to be confident is something that we can hope to have only near the end of the process of inquiry into problems of substantive morality.

Let us refer to the two broadly defined patterns of moral inquiry that I have sketched as the Intuitive Approach and the Theoretical Approach. Both are richly represented in the history of moral philosophy. The Socrates of Plato's dialogues is an admirable exemplar of the Intuitive Approach, while Hobbes and Kant exemplify the Theoretical Approach. Each of the latter begins with a conception of the nature of

morality that he believes dictates a particular method for arriving at moral judgements about particular problems and cases. In recent years, most philosophers working on problems of practical ethics have largely followed the Intuitive Approach, but the Theoretical Approach also has many distinguished contemporary exponents, among them Richard Hare, Richard Brandt, and an assortment of theorists in one or another of the various traditions of contractarianism.

The Theoretical Approach is reformist in a rather radical way. People have always reasoned and argued about substantive issues in morality. According to adherents of the Theoretical Approach, however, people have been badly misguided insofar as their reasoning has diverged from the forms and patterns of moral reasoning prescribed by the correct moral theory. Richard Brandt, for example, suggests that 'is morally wrong' means 'would be prohibited by any moral code which all fully rational persons would tend to support, in preference to all others or none at all, for the society of the agent, if they expected to spend a lifetime in that society.' (Brandt 1979, p. 194) Assuming that this definition also states a test for determining whether an act is wrong, it seems clear that any convergence of the conclusions of most people's actual moral reasoning and the conclusions that might result from Brandt's proposed mode of reasoning would be fortuitous or coincidental. According to the Theoretical Approach, therefore, philosophical ethics is utterly different from, say, the philosophy of science. While the philosopher of science may criticize certain aspects of the practice of science, and may urge scientists to revise their understanding of the nature of their practices or the status of their conclusions, the philosopher does not presume to tell scientists that they have been utilizing the wrong method and would do better to adopt a different approach. The Intuitive Approach is in general more respectful of the modes of moral reasoning that people actually employ - though only because people in fact tend to reason about moral problems in the way it recommends. It too can be revisionist - for example, in its

condemnation of one very common mode of reasoning: namely, the deduction of moral conclusions from the supposedly infallible dicta found in one or another sacred text.

Theory Unchecked by Intuition

Could we really conduct our thinking about morality and moral problems in the way suggested by the Theoretical Approach, without building up from our moral intuitions or consulting those intuitions to test the plausibility of the implications of proposed moral theories? Even those who most vehemently deny that intuitions have any independent credibility nevertheless often build their arguments on the basis of appeals to common intuitions (for example, Rachels 1986, pp. 112-13 & 134-35, Singer 1993, p. 229). But, although this is suggestive of the difficulty of getting persuasive arguments off the ground without linking them to our preexisting moral beliefs, it is merely an ad hominem point and as such does little to support the appeal to intuition. An alternative point that may be urged against the Theoretical Approach is that our intuitions often compel belief in a way that, for most of us, no moral theory does. If an intuition that is highly compelling cannot be reconciled with what seems to be the best supported moral theory, can it be rational to abandon the belief in which we have greater confidence at the behest of the less compelling one?

It is important to be clear about the nature of this challenge. The claim is not simply that moral intuitions often strike us as more obvious or less open to doubt than it seems that any moral theory is. By itself, this would not be a strong consideration in favor of the intuition. The theories of modern physics tell us that many of our common sense beliefs about the nature of the physical world are mistaken. Many of these beliefs seem overwhelmingly obvious while the theory that disputes them may be so arcane as to be unintelligible to all but a few. Yet most of us recognize that at least certain scientific theories that overturn aspects of our common sense conception of the physical world are so well established by their powers of explanation and prediction and by the control they give us over the forces of nature that we readily acquiesce in their claims and concede

that our common sense views must be illusory. If a moral theory could command our allegiance by comparable means of persuasion, we might yield our intuitions to it without demur, even if it had none of the immediate obviousness in which our intuitions tend to come clothed. But the challenge to the Theoretical Approach is that no moral theory, at least at the present stage of the history of philosophical ethics, can have anything like the authority or degree of validation that the best supported scientific theories have. The lamentable truth is that we are at present deeply uncertain even about what types of consideration support or justify a moral theory. There are no agreed criteria for determining whether or to what extent a moral theory is justified. So when an intuition, which may be immediately compelling, comes into conflict with a moral theory, which can have nothing approaching the authority of a well grounded scientific theory, it is not surprising that we should often be profoundly reluctant to abandon the intuition at the bidding of the theory. We can, indeed, be reasonably confident <u>in advance</u> that none of the moral theories presently on offer is sufficiently credentialed to make it rationally required that we surrender our intuition.

It is instructive to consider how most of us respond when, on inquiring into a particular moral problem, we find that some moral theory has implications for the problem that clash with our intuitions. Our response is not to question how well grounded the theory is, on the assumption that we should be prepared to acquiesce if we find that the theory is well supported. If the theory generates its conclusion via a distinct argument, our tendency is to detach the argument from the parent theory and consider it on its own merits. According to R.M. Hare, for example, his universal prescriptivist theory of morality implies that we should reason about the morality of abortion by applying a variant of the Golden Rule: 'we should do to others as we are glad that they did do to us' (Hare, 1975, p. 208). When we discover that this principle implies (according to Hare) not only that abortion is wrong (if other things are equal) but also that remaining childless is wrong (again if other things are equal), we do not go back to

Hare's earlier books to check the arguments for universal prescriptivism. Instead we undertake an independent inquiry to try to determine whether and, if so, to what extent it matters to the morality of abortion that, when an abortion is not performed, there will typically later be a person who is glad to exist who would not have existed if the abortion had been performed. That is, if we are serious about understanding the morality of abortion, we will take seriously the considerations identified as relevant by the theory; and we may be grateful to the theory for helping us to see whatever relevance these considerations may in fact have; but we are generally not overawed by the fact that these considerations have been identified as relevant <u>by the theory</u>. Their provenance in the theory fails to impress.

One may even feel a certain puzzlement as to whether the norms and principles extracted from a moral theory with foundations wholly independent of our pretheoretical intuitions can claim to be constitutive of morality at all. Recall, for example, Brandt's claim that an act is morally wrong if it 'would be prohibited by any moral code which all fully rational persons would tend to support, in preference to all others or none at all, for the society of the agent, if they expected to spend a lifetime in that society.' (Note that Brandt cannot make antecedent knowledge of right and wrong a component of rationality, for that would make the account circular.) The sort of code that might be chosen by rational persons to govern their association (assuming that rational people could agree to accept the same code) might be very pleasant to live under, but is there any guarantee that it would coincide with what we would recognize as <u>morality</u>? Although I cannot demonstrate this, I suspect that any such code would omit a very great deal of what seems to be universal, or nearly so, in human morality.

For example, every society of which I am aware, past and present, has recognized a general moral difference between killing people and letting people die. But fully rational people, unencumbered by any prior moral ideas, who set out to devise a code to regulate their relations with one another might be unlikely to settle upon one that distinguished

between killing and letting die, at least in anything like the way that people's actual moral beliefs have done. For they would presumably find it more to their advantage to adopt a code that permitted the killing of innocent people when this would be necessary to save the lives of a sufficiently greater number of others. There is, of course, scope for debate about this. Perhaps rational people would prefer a code that would give protection of their autonomy priority over protection of their lives. But even if that were so, there are numerous other elements common to most moral actual moral codes that would be unlikely to occur to Brandt's rational moral architects: the notion of honor, various of the virtues, the idea that it is worse if a harm is an intended effect of one's action than if the same harm is a foreseen but unintended effect, the idea that one has duties of respect and beneficence to persons who are not members of the society governed by the code, and so on.

It is possible, of course, that some of these common elements of actual moral codes are in fact irrational - for example, a concern with honor or intention. Or, if they are not irrational, maybe there is some reason why it would be to the advantage of Brandt's moral architects to include them in their code. But these features of the codes that people actually live by are not obviously irrational and there is certainly no guarantee that they would have any appeal to rational choosers drained of all prior moral conviction. But if in fact these or other common features of our actual moral codes would not be included in the code that would be chosen, what reason would we have to recognize the chosen code as morality rather than as just the collection of rules that would make us all most comfortable? Is there really nothing more to morality than just a set of rules the function of which is to facilitate smooth interaction and cooperation to mutual advantage among the members of a particular society? The phenomenology of moral experience certainly suggests that this is an altogether shallow and reductive understanding of the nature of morality.

I have focused on Brandt's theory for the purpose of illustration; but similar points could be urged against all moral theories with foundations independent of our moral intuitions. To the extent that a theory leaves out major features of what most of us would recognize as morality, there is reason to wonder whether the theory is really a theory about morality at all.

Moral Epistemology

The remarks in the previous section are meant only to suggest certain reservations we might have about the Theoretical Approach; they are far from providing decisive reasons for rejecting that approach. Moreover, even if we had stronger grounds for skepticism about the Theoretical Approach than those I have offered, this would still be insufficient to throw us into the arms of the Intuitive Approach. For it is hardly a ground for confidence in our intuitions that there are reasons for doubting the approach to moral inquiry that denies them a role. Something more positive has to be said on behalf of our intuitions themselves. At a minimum, more needs to be said about the role they are supposed to have in the structure of justification in ethics. In what follows I will first offer a few general remarks about moral epistemology, after which I will briefly sketch an account of moral inquiry that explains the role that our intuitions should have in our moral thinking and also helps to elucidate their epistemological status. I will then conclude by noting that there is a conception of the nature of moral knowledge that has independent plausibility and, if correct, offers a deeper understanding of [...]. It should be noted that the issues I will be addressing are deep and difficult and I can do little more here than skim the surface.

Theories in epistemology may be theories either of truth or of justification. I will focus on the issue of justification and simply assume that there is a tight connection between justification and truth. Accounts of justification tend to be divided into two major approaches: coherentism and foundationalism. This is the case in the narrower

area of moral epistemology as well: accounts of moral justification tend to be either coherentist or foundationalist.

Coherentist accounts of justification hold that a moral belief is justified solely in terms of its relations, particularly its inferential relations, with other beliefs. It is justified to the extent that it coheres well with a set of beliefs that together form a coherent whole. By contrast, foundationalist accounts hold that some beliefs are self-justifying - at least in the sense that they are justified independently of their relations to other beliefs. According to foundationalist accounts, a moral belief is justified if and only if it is either self-justifying or bears an appropriate inferential relation to a belief that is self-justifying.

Of the two types of account, coherentism is generally thought to be more hospitable to the Intuitive Approach. The most commonly endorsed method of moral inquiry among contemporary moral philosophers is the method described by John Rawls under the label 'reflective equilibrium' (Rawls 1972, pp. 19-21 and 48-51). According to the method of reflective equilibrium, we begin with a set of moral intuitions about particular cases, filter out those that are the obvious products of distorting influences, and then seek to unify the remaining intuitions under a set of more general principles. We seek principles that both imply and explain our particular judgements. But the match between principles and intuitions will inevitably be very imperfect in the first instance. A candidate principle may imply a great many of our intuitions and yet have some implications that conflict with other intuitions. In that case we may modify or even abandon the principle; but, if the principle has considerable explanatory power with respect to a wide range of intuitions and cannot be modified without significant sacrifice of this power, we may instead decide to reject the recalcitrant intuitions. In this way we move back and forth between intuitions and principles, making reciprocal adjustments until our beliefs at the various levels of generality are all brought into a state of harmony or reflective equilibrium. This method is generally interpreted in coherentist terms, in that it is understood to make coherence with other beliefs the sole criterion of a belief's

credibility. Yet it obviously treats intuitions as potential sources of moral knowledge. Although intuitions arise noninferentially and thus, in coherentist terms, have no prima facie credibility on their own, those that survive the initial filtration and are compatible with the principles that emerge in the process of seeking reflective equilibrium turn out to be justified moral beliefs.

Foundationalist theories of moral justification tend to be favored by proponents of the Theoretical Approach. Typically, the foundational beliefs (that is, those that are not justified in terms of their relations to other beliefs) are held to be nonmoral; justified moral beliefs are all ultimately derivable via some process of reasoning that is based on the foundational nonmoral beliefs (for examples, see Timmons 1987). Those who attribute authority to our moral intuitions tend, understandably, to be more reluctant to embrace foundationalism. This is mainly because it seems implausible to regard our intuitions themselves as foundational. This seems to attribute to them too exalted a status. While our intuitions do seem to have a certain initial credibility, it seems exorbitant to suppose that they are self-evident or self-justifying. We recoil from the suggestion (advanced, as I noted earlier, by various traditional Intuitionists) that intuitions are the unshakable basis on which all moral knowledge rests.

There are, however, at least two ways of overcoming this ground of reluctance to combine foundationalism with the Intuitive Approach. The first is to recognize that a belief may be of the foundational sort and yet be defeasible. Suppose, for example, that sense perceptions are the foundations of empirical knowledge. Even if all empirical knowledge is derived immediately from sense perceptions or is ultimately traceable by chains of inference to sense perceptions, it does not follow that all sense perceptions are sources of empirical knowledge. Some may be distorted, illusory, or otherwise erroneous. And there is no reason why the same may not be true, mutatis mutandis, in the case of moral intuitions.

Second, a foundationalist account of moral knowledge may treat intuitions as reliable sources of moral knowledge without treating them as foundational or self-justifying. It is this possibility that I wish to explore in more detail.

A Sketch of a Foundationalist Conception of Moral Justification

It seems to me that the method of reflective equilibrium, or a process very much like it, is the best or most fruitful method of moral inquiry. Of the known methods of inquiry, it is the one most likely to lead to justified moral beliefs. It does not, however, have to be interpreted within the coherentist framework. It is compatible with a foundationalist conception of moral justification.

Let me outline in more detail how the method works. Again, we begin with intuitions about particular problems, particular cases. If our initial interest is in a problem about which we have no intuitions, or about which our intuitions are weak or conflicting, we should, as I have suggested with reference to the problem of abortion, find closely related cases about which we have confident intuitions and work from these. The question immediately arises, however, why we should carry the inquiry any further. Why cannot we rest content with our intuitions, allowing ourselves to be guided by them on a case-by-case basis? Part of the answer, of course, is that there are many moral problems about which we have no intuitions, or about which the intuitions we have are weak, conflicting, or obviously suspect or dubious. We need a method for determining what we should believe and what we should do in cases such as these.

There are many reasons why an intuition may appear tainted or suspect, even to the person who has it. It may, for example, seem to be a remnant of a religious or metaphysical system accepted during childhood but now rejected; it may be a product of social prejudice or self-interest (e.g., the belief among slave owners in the acceptability of slavery); it may reflect the operations of an aberrant state of mind; and so on. If one's intuition about a particular moral problem is doubtful for any such reason, one will need a method of addressing the problem other than merely consulting one's intuition.

Even if one's intuition presents itself compellingly, one knows that moral beliefs are both fallible and important and that it is thus imperative to subject them to scrutiny. One should seek to determine whether the intuition can be justified or defended. How should one proceed? When one's moral intuition is challenged by another person, it is natural to respond by appealing to claims of a higher level of generality that imply or explain the intuition. The assumption here is that the credibility of the intuition is enhanced if the intuition can be subsumed under a plausible moral principle. So, for example, the intuitive judgement that it would be wrong to torture the cat for fun might be defended by appealing to the principle that it is wrong to cause suffering without good reason. Private moral reflection may follow the same dialectical pattern as moral disagreement between persons. We should challenge our own intuitions in much the way that an opponent might challenge them; but we may also respond in much the same way, by trying to bring them within the scope of a plausible principle.

But why suppose that the credibility of one's intuition is enhanced when it is shown that the intuition follows from a moral principle (together with the facts of the case)? One suggestion is that the principle may elucidate the intuition by identifying the features of the case that are morally salient. If, for example, one feels intuitively that it is wrong to kill animals for sport, one's objection is sharpened or focused if it is seen to follow from the more general view that it is wrong, other things being equal, to deprive any individual of a good that that individual would otherwise have. The principle brings out more clearly exactly what one finds intuitively objectionable.

It is also important to note that no one supposes that just any principle will do. For the principle to support the intuition, it must have independent credibility. (The principle itself may have intuitive appeal. Presumably it does if one is inclined to accord it some initial credence - that is, if it counts as a belief.) According to coherentism, of course, the mere fact that the principle implies the intuition provides some minimal epistemic support for each; for mutual coherence among beliefs is the criterion of justifiability.

But, even according to coherentism, the principle will provide no more than token support for the intuition unless it is itself well integrated within a larger network of mutually coherent beliefs. Hence the method of reflective equilibrium demands that the principle itself be tested for consistency and coherence with other moral beliefs. Its implications about particular cases should not conflict with one's intuitive judgements about those cases and, to the greatest extent possible, its implications should not contradict the implications of other principles one accepts. It is, of course, too much to require that moral principles not have conflicting implications: conflict is the price of pluralism. But conflicts should, in principle, be resoluble, in that one recognizes the necessity of one value's yielding to another, though not without some irreducible loss.

So the defense of one's initial intuition by subsuming it under a more general principle is only the beginning of moral inquiry. The principle must itself be assessed by testing its implications for consistency with one's other beliefs. One need not accept coherentism in order to appreciate the importance of this test. There are practical reasons why inconsistent moral beliefs are problematic: they may, for example, provoke indecision and, ultimately, paralysis of the will. More importantly, the achievement of greater coherence among one's beliefs diminishes the likelihood of error by helping one to identify and weed out moral beliefs spawned by self-interest, faulty reasoning, failure of imagination, illusory metaphysical beliefs, impaired faculties, and other sources of distorted or mistaken belief.

But there is a deeper basis for trying to subsume an intuition under a principle that is itself supported by its power to unify and explain a range of other intuitions. This is that the process of achieving increasing coherence among principles and intuitions facilitates the discovery of one's deeper values and also brings one's surface beliefs about particular cases into alignment with those deeper values in a way that reveals and illuminates the connections between them. When one seeks to formulate a moral principle that implies and illuminates one's intuition about a particular problem or case,

one is in fact groping or probing for one's own deeper values. The expectation that the principle will illuminate and explain the intuition assumes that the intuition is in fact an expression or manifestation, in a particular context, of a moral belief that is deeper, more basic, and more general than the intuition itself. One's efforts to formulate the principle and to revise and refine it in a way that brings more and more of one's intuitions within its scope are attempts to capture or articulate some core moral belief in its full generality, to get its form exactly right, omitting nothing, however subtle.

This process, as I have described it, is indistinguishable from that endorsed by the coherentist practitioners of reflective equilibrium. One seeks support for an intuition by appealing to a principle, then seeks to support the principle by demonstrating its compatibility with other intuitions, and so on. Why not understand the method as most people do, in coherentist terms?

There are various general objections to coherentist accounts of justification in ethics (e.g., Gaus 1996, ch. 6). I will not rehearse them here. (There are also, of course, general objections to foundationalism; I will not discuss those either.) I will simply note two problems that I find particularly disturbing. One is that, according to coherentism, no belief is immune to rejection, no matter how compelling it may be. If its elimination from the network of beliefs would enhance overall coherence within the network, the belief must go. Indeed, it seems possible, though not likely, that a coherentist approach to the pursuit of reflective equilibrium could lead ultimately to the rejection of every belief with which one started. Both these suppositions, however, are utterly alien to moral life and moral reflection. There are some moral beliefs that we simply cannot give up just for the sake of greater coherence. Sometimes we must hold tenaciously to certain convictions even at the cost of diminished coherence or systematicity. According to the coherentist, one's moral beliefs are like pieces in a game: one shuffles them around, sacrifices some, and acquires others, all for the sake of achieving certain relations among them. No piece has any significance in itself; it has significance only in relation to the

other pieces and in particular in the contribution that it makes to the whole of which it is a part. If moral reflection were really a game like this, in which our moral beliefs had no claim to our loyalty and thus were readily expendable in the service of coherence, coherence would be fairly easily and painlessly achievable. It is because some of our moral beliefs compel our allegiance independently of their inferential relations to other beliefs that coherence always seems a distant, perhaps impossible goal.

A closely related worry about coherentism is that it assigns the same epistemic status to our intuitions about particular cases as it does to the deeper principles of which the intuitions are expressions. They stand in relations of reciprocal support: the principles imply the intuitions and we can therefore infer our way to the principles on the basis of the intuitions. But in fact the relations of reciprocal support seem asymmetrical: the principles seem to be epistemically more basic, more secure. They articulate our core values which unify, explain, and justify our intuitive judgements. Our intuitions do not so much justify the principles as merely provide evidence of their existence and guidance as to their nature. In short, the principles are foundational with respect to the intuitions. Insofar as our intuitions are reliable sources of moral knowledge, they are so because they are expressions of, and point back to, a range of deeper, more general values that lie at the core of our nature as moral beings.

As I noted earlier, foundationalism is distinguished by the view that some beliefs are justified independently of their relations to other beliefs. I will refer to these beliefs as 'foundational.' Among those who accept a foundationalist moral epistemology, there is a rough division between those who take certain nonmoral beliefs to be foundational and those who identify certain moral beliefs as foundational. Among the latter, there is a further division between those who take intuitions to be foundational (e.g., Ross 1930 and Gaus 1996) and those who take some general principle or principles to be foundational (e.g., Sidgwick 1907). Thus, according to Geoffrey Sayre-McCord, 'many have treated the privileged [i.e., foundational] moral beliefs as roughly on a par with perceptual

judgments and suggested that the justification of our various moral principles parallels the kind of justification our scientific principles receive from perception. Others have thought that our privileged moral beliefs concern, instead, the most general and abstract principles of morality, and that these in turn serve to justify (or not) our other beliefs deductively' (Sayre-McCord 1996, p. 150).

The view that I have suggested is of this second sort. But it is distinguished from many views of this sort in that it does not regard the foundational principles as self-evident or accessible through the exercise of intuition. Many philosophers (e.g., Sidgwick 1907 and Unger 1996) have regarded our intuitions about principles as more reliable than our intuitions about particular cases. But, insofar as a moral principle is substantive in character rather than merely formal (for example, 'treat like cases alike'), I believe that it is a mistake to have confidence in our intuitive apprehension of the principle. To be justified in accepting a moral principle, we must first understand what it commits us to in particular cases. As William James noted in a letter written long before he became a practicing philosopher, 'no one sees farther into a generalization than his own knowledge of the details extends' (Barzun 1983, p. 14). So, while I regard the principles rather than our intuitions as foundational, I deny that moral inquiry proceeds by deducing conclusions about particular cases from self-evident moral principles. Rather, it seems to me that <u>the order of discovery is the reverse of the order of justification</u>. Although the deeper principles are explanatorily prior, we have to work our way to them via our intuitions in much the way that scientists work towards general principles via our perceptual data. The process of discovering and formulating the more general principles is evidently difficult and intellectually demanding, rather in the way that discovering the syntactically structures that govern our use of language is. As this familiar analogy suggests, as we grope our way towards the principles, we are discovering what we antecedently believe, albeit below the level of conscious awareness. The principles that we hope to uncover express deep dispositions of thought and feeling

that operate below the level of consciousness to regulate our intuitive responses to particular cases.

If this is right, it explains why we experience the process of moral inquiry as a process of discovery rather than an exercise of choice or will. It also explains why the foundationalist approach I am describing should coincide with coherentism in holding that we may expect to arrive at a moral theory, if at all, only near the end of reflection about particular problems and cases rather than coming to the problems with a theory ready to be 'applied'. It explains why it is more than a little suspicious when philosophers emerge from graduate school already believing themselves to be in possession of the correct moral theory.

Challenges

This brief sketch of an account of moral justification raises far more questions than can be answered, or even addressed, in the remainder of this short essay. The central question is, of course, what reason there is to suppose that the principles that might emerge through the pursuit of reflective equilibrium are reliable. One response, which is hardly satisfying but may nevertheless be true, is that they are simply ultimate: there is nothing more basic, nothing by reference to which they can be justified. When we have worked our way back to the principles or values that underlie and best explain the countless judgements we make about particular problems and cases, there is simply nowhere deeper to go.

This would perhaps be easier to accept if the principles at which we might arrive would be luminously self-evident; but there is no guarantee that they would be. Indeed, if experience is any guide, there is even reason to suspect that they may not be. Consider, for example, our intuitions about the moral differences between doing harm and allowing harm to occur, or, more specifically, between killing and letting die. Various philosophers have sought to trace these intuitions, which emerge very strongly in particular cases, to their source in certain deeper principles. In no instance of which I am

aware do the principles cited have anything approaching the immediate plausibility of the intuitive judgements that they were supposed to explain. The principles are too abstract, too schematic, to carry immediate conviction. But perhaps the problem is that the principles that have been identified are not really ultimate; perhaps there is a deeper layer of explanation that has yet to be uncovered.

There is reason to conjecture that at least some of our foundational or core values are biologically based. One piece of evidence for this is the surprising uniformity of our intuitions about particular cases. We have been impressed for so long by the claims of anthropologists, postmodern relativists, undergraduates, and others about the diversity of moral opinion that we are inclined to overlook how much agreement there in fact is. Interestingly, what one finds is that moral disagreements tend to widen and intensify the more we abstract from particular cases and focus instead on matters of principle or theory. When the partisans of different schools of moral thought turn their attention to particular cases, there is far more intuitive agreement than their higher-level disputes would lead one to suspect. Although this is merely anecdotal, those who teach courses in moral philosophy will, I think, recognize in the following observation a familiar phenomenon. When I teach a course on normative ethics, my students tend to be ethnically quite heterogeneous. Many were raised in remote parts of the world and speak with a discernible accent from their native language. Yet, in spite of having been raised in diverse cultures, they tend to have the same intuitive responses to the cases we discuss (which is, of course, the expectation of those writers in ethics who appeal to intuitions about cases; otherwise, if a case fails to elicit the same intuitive response from most readers, it provides no basis for moral argument).

Here is a pair of examples that teachers often employ.

[1] Trolley: A runaway trolley is careering down the mainline track where it will soon hit and kill five people trapped on the track. A bystander can flip a switch that would divert the trolley onto a branchline track before it reaches

the five; but it would then hit and kill a single person who is trapped on the branchline track.

[2] Transplant: Five people in the hospital need organ transplants to survive. It would be possible to save them by killing a healthy patient and dividing his organs among them.

In both of these cases there is a choice between letting five innocent people die and killing one innocent person. Virtually all of my students agree that it would be permissible to kill the one to save the five in the Trolley case; but virtually all also agree that it would be impermissible to do this in the Transplant case.   there are two considerations that make this consensus rather remarkable. First, while the students discriminate intuitively between the two cases, the basis of their own discrimination typically eludes them; they are unable to identify the factors that distinguish the two cases morally. (As those familiar with the relevant literature are aware, it is possible to introduce variations into these cases that challenge the ability of even the most subtle philosophers to identify the bases of their own intuitive discriminations. This is not just a game one plays with undergraduates.) Second, the cases are beyond the range of the students' experience. Choices of these sorts are, indeed, very rare and it is unlikely that the students have ever previously thought about them at all. Yet their intuitive responses are quite strong. This, together with the fact that the students tend to come from widely divergent cultural backgrounds, suggests that the intuitions are not learned responses.

An alternative explanation appeals to the analogy, noted earlier, between the deep values that we unconsciously consult in making intuitive moral judgements and the deep syntactical structures that govern our use of language. Colin McGinn attributes this conception of moral knowledge to Noam Chomsky, who originally formulated the parallel account of our knowledge of language. 'According to Chomsky,' McGinn writes, 'it is plausible to see our ethical faculty as analogous to our language faculty: we acquire ethical knowledge with very little explicit instruction, without great intellectual

labour, and the end-result is remarkably uniform given the variety of ethical input we receive. The environment serves merely to trigger and specialise an innate schematism. Thus the ethical systems of different cultures or epochs are plausibly seen as analogous to the different languages people speak - an underlying universal structure gets differentiated into specific cultural products' (McGinn 1993, p. 30).

McGinn goes on to comment that 'perhaps the innate system of commonsense psychology, installed to negotiate our social relations, contains the resources for generating the basic principles of ethics' (McGinn 1993, p. 30). Certainly the sociobiological literature has familiarized us with accounts drawn from evolutionary biology that purport to explain the presence of those deep values that give rise to our intuitions about such matters as parental responsibilities, marital fidelity, and so on.

But this kind of account leaves ample room for skepticism. Even if we are biologically programmed to think and feel in a certain ways, we can still ask whether it makes sense to think and feel in those ways. There is, for example, no obvious reason why natural selection would have left us with much concern for the lives and well-being of nonhuman animals: evolutionary biology would thus predict that our attitudes towards animals and our practices that involve them would be 'speciesist' - that is, would give little or no weight to the interests of animals simply because animals are unrelated to us in certain ways. And indeed our beliefs and practices are speciesist, for the most part. But only for the most part. Many reflective people who have tried to work through their moral beliefs about animals in the way recommended by the method of reflective equilibrium have concluded that our traditional beliefs and practices are unconscionable. A plausible account of moral epistemology must be able to accommodate and make sense of our ability to make radical revisions of moral beliefs that are apparently quite deeply situated in our biological nature.

Another worry about the innatist conception of our core values is that it may suggest a disturbingly subjectivist conception of value. If our moral judgements are

merely the products of biologically programmed dispositions of thought and feeling, then it may seem that values are entirely within us rather than in the world - that, as various subjectivists have put it, values are projections of our minds onto a world that is itself entirely value-neutral. On this view, values are like tastes, essentially arbitrary and ultimately uncriticizable. This is, of course, a common enough position, but many of us are strongly inclined to resist it.

The innatist conception does not, however, lead inevitably to a subjectivist account of value. It is not out of the question that our minds have evolved in such a way that our moral sensibilities are capable of apprehending what is actually there. It is true that our sensibilities contribute something to our apprehension of value (positive or negative). Some philosophers have suggested that our apprehension of value is analogous to our perception of secondary qualities (colors, tastes, etc.). They have argued that, just as color is a real property of surfaces that is nevertheless analyzable in part in terms of its disposition to produce certain sensations in individuals with appropriate perceptual faculties in certain conditions, so too value may be a genuine feature of acts, events, and so on that is nevertheless analyzable in part in terms of its tendency to evoke certain characteristic responses in those endowed with an appropriately sensitive moral sensibility. On this view, we do not create value any more than we create color. Although we may be innately disposed to respond with approval or disapproval to acts and events of various kinds, it does not follow that value is created only by the act of valuing. (The analogy between values and secondary qualities is of course highly controversial and the relevant literature is extensive. See, for example, McDowell 1985 & 1997, Wiggins 1991, Wright 1988, Goldman 1988, and Johnston 1989.)

A further concern about the foundationalist account I have sketched is that, if our intuitive judgements are traceable to deeper principles that are themselves articulations of dispositions of thought and feeling that may be innate, one would expect there to be greater harmony among our intuitions than one in fact finds. Our intuitions, indeed, seem

bewilderingly chaotic and we find, as soon as we begin to try to unify them under principles of greater generality, that inconsistencies abound.

There are several explanations for this. One is that our moral intuitions undoubtedly stem from numerous diverse sources: while some derive from biologically programmed dispositions that are largely uniform across the species, others are the products of cultural determinants, economic or social conditions, vagaries of individual character and circumstance, and so on. Given the heterogeneity of these sources, it is hardly surprising that there are conflicts. The process of achieving greater consistency and coherence is, as I noted earlier, in part an effort to weed out those intuitions that derive from tainted sources. But even among those intuitions that seem to derive from sources that we may regard as reliable, there is no a priori guarantee of harmony. There are some areas of morality - for example, the area concerned with causing new people to exist - in which there seem to be deep incompatibilities that do not seem to be the result of error. In these areas there may be no way of achieving coherence that does not involve a rejection of something that is true. I suspect that the richness of moral life and experience will continue indefinitely to resist our efforts to reduce it to a unified and comprehensible system.

BIBLIOGRAPHY

Audi, Robert: 'Intuitionism, pluralism, and the foundations of ethics', <u>Moral Knowledge?</u>, ed. Walter Sinnott-Armstrong and Mark Timmons (New York: Oxford University Press, 1996), pp. 101-36.

Barzun, Jacques: <u>A Stroll With William James</u> (New York: Harper & Row, 1983).

Brandt, Richard B.: <u>A Theory of the Right and the Good</u> (Oxford: Clarendon Press, 1979).

Brink, David: <u>Moral Realism and the Foundations of Ethics</u> (Cambridge: Cambridge University Press).

Brody, Baruch: 'Intuitions and objective moral knowledge', <u>The Monist</u> 62, (1979), 446-56.

Daniels, Norman: <u>Justice and Justification: Reflective Equilibrium in Theory and Practice</u> (Cambridge: Cambridge University Press, 1996).

DePaul, Michael R.: <u>Balance and Refinement: Beyond Coherence Methods of Moral Inquiry</u> (London: Routledge, 1993).

Gaus, Gerald: <u>Justificatory Liberalism: An Essay on Epistemology and Political Theory</u> (New York: Oxford University Press, 1996).

Goldman, Alan H.: <u>Moral Knowledge</u> (London: Routledge, 1988).

Griffin, James: <u>Value Judgement: Improving Our Ethical Beliefs</u> (Oxford: Clarendon Press, 1996).

Hare, R.M.: 'Abortion and the golden rule', <u>Philosophy and Public Affairs</u>, 4 (1975), 201-22.

Johnston, Mark: 'Dispositional theories of value', <u>Proceedings of the Aristotelian Society</u>, supplementary volume, 63 (1989), 139-74.

Lewis, David: 'Dispositional theories of value', <u>Proceedings of the Aristotelian Society</u>, supplementary volume, 63 (1989), 113-37.

McDowell, John: 'Values and secondary qualities', <u>Morality and Objectivity; A Tribute to J.L. Mackie</u>, ed. Ted Honderich (London: Routledge, 1985), 110-29.

_____: 'Projection and truth in ethics', <u>Moral Discourse and Practice</u>, ed. Stephen Darwall et al. (New York: Oxford University Press, 1997), 215-25.

McGinn, Colin: 'In and out of the mind', <u>London Review of Books</u>, (2 December 1993), 30-31.

Rachels, James: <u>The End of Life: Euthanasia and Morality</u> (Oxford: Oxford University Press, 1986).

Rawls, John: <u>A Theory of Justice</u> (Oxford: Clarendon Press, 1972).

Raz, Joseph: 'The relevance of coherence', <u>Ethics in the Public Domain</u> (Oxford: Clarendon Press, 1994), 277-325.

Ross, W.D.: <u>The Right and the Good</u> (Oxford: Clarendon Press, 1930).

Sayre-McCord, Geoffrey: 'Coherentist epistemology and moral theory', <u>Moral Knowledge?</u>, ed. Walter Sinnott-Armstrong and Mark Timmons (New York: Oxford University Press, 1996), pp. 137-89.

Sidgwick, Henry: <u>The Methods of Ethics</u>, seventh edition (London: Macmillan, 1907).

Singer, Peter: 'Sidgwick and reflective equilibrium', <u>The Monist</u>, 58 (1974), 490-517.

_____: <u>Practical Ethics</u>, second edition (Cambridge: Cambridge University Press, 1993)

Timmons, Mark: 'Foundationalism and the structure of ethical justification', <u>Ethics</u>, 97 (1987), 595-609.

Unger, Peter: <u>Living High and Letting Die: Our Illusion of Innocence</u> (New York: Oxford University Press, 1996).

Wiggins, David: <u>Needs, Values, Truth</u>, second edition (Oxford: Blackwell, 1991).

Wright, Crispin: 'Moral values, projection, and secondary qualities', <u>Proceedings of the Aristotelian Society</u>, supplementary volume, 62 (1988), 1-26.