Learning About Right and Wrong: Ethics and Language

OSWALD HANFLING

The inculcation of moral values is often thought to be a matter of choice. We hear people speak (sometimes with satisfaction, as if it were to their credit) about having taught their children 'the difference between right and wrong'. But what does this mean? What would the teaching consist in? And to what extent would it be a matter of choice? 'Even the wisest and best of our countrymen', observed Socrates to Protagoras, 'are unable to hand on to others the virtue that they possess'.

> Pericles, for instance, the father of these two boys, gave them the very best education in everything that depends on teaching, but in his own special kind of wisdom he neither trains them himself nor hands them over to any other instructor: they simply browse around on their own like sacred cattle, on the chance of picking up virtue automatically (*Protagoras* 319e–320a).

But is virtue really something that cannot be taught? Of course it cannot be taught in the ways in which, say, history or science are taught, but there are ways of teaching a child to be brave or generous, for example; and this is something that a parent may choose or not choose to do. On the other hand, there is also such a thing as 'picking up virtue automatically'. One way of doing so is by emulation. A child growing up in the shadow of a virtuous parent might be expected to acquire virtue in this way. But there is another way of 'picking up virtue', which is not dependent on the virtue of a parent. What children necessarily pick up, if they are normal, is their native language, and this includes moral language. The acquisition of language happens in two main ways: that of being spoken to by people competent in the language, and that of being present when such people speak to one another. Given these conditions, it follows as a matter of human nature that the child will pick up the use of language, and this includes the language of right and wrong, vice and virtue.

It might be objected that Socrates's question was not about the *language* of vice and virtue, but about virtue itself. These questions are not, however, as distinct as might be thought. Knowing the dif-

doi:10.1017/S0031819103000020 ©2003 The Royal Institute of Philosophy

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

ference between right and wrong, between virtue and vice, is nothing other than knowing what acts are *describable* as 'right', 'wrong', 'virtuous' etc.: it is a matter of being able to apply these words correctly.

It is sometimes thought, however, that these words are devoid of descriptive meaning, so that any action whatever might be described as right or wrong, good or bad. But this free-for-all view is true neither to how these words are learned nor to how they are used thereafter. We learn such words by hearing them applied to typical examples, as when we are told that what we did on one occasion was wrong, while what we did on another was a good deed. These associations are enlarged and refined as we grow up. At this stage we learn to think and reason about right and wrong, and our mastery of moral language extends beyond the simple cases with which it began. We may also find, or decide, that some of the 'typical examples' to which we first heard moral terms applied are in fact of doubtful validity. This does not mean, however, that there is no longer any connection with typical examples, for such examples are still definitive of the terms in question. The much-maligned 'paradigm case argument' is relevant here (as it is, indeed, elsewhere)[1] 'This is what we call "a good deed"', we might say to someone who does not understand the meaning of this phrase; and again, using the rhetoric of the paradigm case argument: 'If this isn't a good deed, what is?' The truth of moral statements, as applied to familiar paradigms, is a matter of meaning: someone who questioned the truth of such a statement would be held not to understand the meanings of the relevant words.

There is a similar connection between the meanings of aesthetic predicates and suitable paradigms.[2] It is sometimes thought that aesthetic statements are 'purely subjective', so that anything whatever might be described as, say, ugly or beautiful without any transgression of logic. But this is not so. Of course aesthetic judgements vary widely, but there are limits. Someone who used the words 'ugly' and 'beautiful' in excessively surprising ways would be regarded as lacking a proper understanding of their meanings. Again, if an anthropologist told us that people in a remote society describe the most surprising things as beautiful or ugly, he might be advised to reconsider his translation of their language.

It must be admitted, however, that the connection with paradigm cases is looser in the case of aesthetics than in that of morals, and

[1] See O. Hanfling, *Philosophy and Ordinary Language* (Routledge 2000), ch. 5.

[2] I once saw a cartoon in which a mother, pointing to some greeting cards in a shop, instructs her child: 'There, now you know what bad taste is.'

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

that there is more room for personal disagreements here. Such disagreements are tolerable, because aesthetic descriptions are not as central to our lives as moral ones, and do not impinge on behaviour to the same extent. For the same reason we are not particularly disturbed if a person lacks aesthetic sensitivity, whereas such a lack in matters of morality would be regarded as a serious fault.

What I said about moral descriptions is incompatible with the 'open question argument' that is commonly but erroneously attributed to G. E. Moore. Moore was concerned to deny that the meaning of 'good' could be *identified* with that of terms such as 'pleasure'. If it could, he argued, then the question 'Is pleasure good?' could amount to no more than 'Is pleasure pleasant?'; and this is obviously not so. It does not follow, however, that the question 'Is X good?' is always an open one. Consider the question whether, in the absence of countervailing reasons,[3] it is good to help people in need. Someone who treated this as an open question would not be regarded as understanding the meanings of the words concerned. But this is not, of course, to say that the meaning of 'good' is *identical* with that of 'helping people in need', or with any other expression—which is what Moore was concerned to deny.

Wittgenstein's point about 'agreement in judgements' (PI I/242) is applicable here no less than in the case of other descriptive words. People may disagree about whether an object is to be described as 'red', 'flat', 'a game', 'a number' etc.; but there must be a bedrock of cases that are beyond dispute—where speakers of the language must agree in their descriptions. If this were not so—if every application of a given word were open to question—then the word could not function as an instrument of communication. People have disagreed about whether, for example, the earth is flat; but agreement and disagreement about this is possible only against the background of uses of 'flat' that are *not* open to disagreement. Thus if someone described, say, a tennis ball as 'flat', or denied that my table top is flat, he must be either ignorant of the meaning of this word or introducing a new meaning for it.

[3] I use this expression in preference to the more elegant but less accurate 'other things being equal'. The latter leaves it unclear what is to be included among 'other things', and what is supposed to be equal to what. I take it that the relevant 'other things' must be reasons of a moral kind. The phrase I use may be understood in the context of a situation in which an accused person is trying to justify his transgression of some moral principle by bringing forward reasons, of a moral kind, against following the principle in that particular case. If this cannot be done (or if the reasons are inadequate), then he is to blame for his transgression.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

The same is true of moral words. There is of course much disagreement about the applications of 'right' and 'wrong' and other moral words, but, again, there are limits. Someone who denied that, in the absence of special circumstances, it is good to be honest, or who described 'clasping one's hands three times in an hour' as 'a good action'[4] would be betraying a failure to understand what these words mean, and not merely expressing an unusual opinion.

So far, the situation as I have described it is no different from that of mastering any other descriptive language. In the case of moral language, however, there is also an emotive element. It is largely by way of *admonition* that basic moral concepts are imparted to children, and these admonitions are reinforced, more or less spontaneously, by suitable sanctions: smiles and affection on the positive side; frowns or physical interference on the other. These sanctions are akin to imperatives rather than to the descriptive 'right' and 'wrong'. 'Don't do that!', we say to the child: this is the original and immediate verbal counterpart of frowns and physical prevention. The introduction of 'right' and 'wrong' is a further development: these words indicate objectivity, going beyond parental approval and disapproval. In due course the child itself comes to have suitable feelings towards the relevant acts. The concept of *shame* is important here. When, with a suitable facial expression, we put it to the child: 'Aren't you ashamed of yourself?', the child is apt to acquire a tendency to have suitable feelings—of shame, remorse and guilt—about the wrongs it has done, and this tendency is likely to persist in later life.

Moral reasons are sometimes felt to have a peculiar authority—to be 'overriding' in some way. This cannot mean, of course, that they override other reasons in practice. But how does that feeling come about? The feeling comes into play when we are tempted to do what we ought not to do. The experience of temptation, of 'want' versus 'ought', enters our consciousness at an early age. This happens because parents need to curb the unsocial behaviour of children—to get them to see that one ought not to behave like that. (A similar process takes place with the 'ought' of prudence, assuming that the parents are interested in their children's health and safety.) The approved patterns of behaviour are instilled in us by a kind of conditioning, with good behaviour attracting approval and affection, and bad behaviour the reverse; and these associations remain to influence us in later life, when 'the voice of conscience', reinforced by feelings of shame and guilt, takes the place of parental approval

[4] Philippa Foot, *Virtues and Vices* (Blackwell 1978), 118–9.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

and disapproval. In this way that the demands of morality can appear as if they came from a superior authority.

Such feelings, as I said, are not confined to the sphere of morality; the temptation of 'want' versus 'ought' also occurs where one's own welfare is concerned. Thus a person who succumbs to a temptation to smoke or drink when he knows it will be bad for him may be assailed by feelings of guilt and remorse, even if no harm to others is likely. The question 'Aren't you ashamed of yourself?' is also appropriate here. In this case, as in that of moral temptation, we have been trained, by the emotional pressures I have described, not always to give way to our immediate desires—that this can be something to be ashamed of.

Since the work of Wittgenstein it has been widely recognized that the learning of language cannot be wholly accounted for by the fact that children are exposed to speech in suitable circumstances. What has further to be assumed is that the child is able to 'catch on' to what is happening, linguistically, to him and around him. That this does happen, in the case of normal children, is a fact of human nature; but it is only if this fact is in place that the child will come to describe and classify things in the correct manner. In the case of moral language, however, a further point must be made. A normal child will not only learn to describe and classify things correctly; he will also come to respond with appropriate feelings and attitudes towards the relevant actions and people.[5]

The connection of moral description with appropriate feelings and attitudes is, again, comparable to what happens in the case of 'beauty' and its opposite. How is the concept of beauty acquired? We are probably introduced to it by ostension. 'What a pretty flower', says the parent to the child in the push-chair; 'isn't it beautiful?' Gradually the child comes to repeat these words when looking at pretty flowers and other beautiful objects. He 'catches on' to the appropriate use of these words when applied to a variety of objects. (He also discovers, in due course, that judgements of beauty are often disputed.) But what needs to be added to this account is that a certain *attitude* to beauty is normal. This attitude is likely to be expressed when the adult says 'What a pretty flower' etc., for he

[5] 'If we consider what in detail a person who has learned the difference between right and wrong has learned, we do not normally draw a line between some things, namely, what he has learned to say and do, and other things, namely, what he has learned to feel, and relegate the latter to the class of mere after-effects of his learning to say and do the proper things.' (Gilbert Ryle, 'On Forgetting the Difference between Right and Wrong', *Collected Papers,* 387).

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

will typically say this in a tone of voice that is expressive of pleasure, and perhaps with a smile; and we expect the same attitude to be acquired and expressed by the child. Beauty is something we love, and to describe a thing as beautiful is to value it accordingly.

I have described the ways in which children are inducted into the moral language-game—including both its vocabulary and the appropriate feelings and attitudes. To what extent do parents have a choice in this process? Of course they can and do exercise some choice in the values that they encourage and discourage in their children, and in their ways of doing this. But there are limits. A child, as Bradley remarked, 'is born not into a desert, but into a living world'.[6] Now the living world into which a child is born is, in the first place, that of a family. But in this world interaction with others, both harmful and beneficial, is inevitable; and the inculcation of certain moral values follows from this by a kind of necessity. Children are liable to be a nuisance. This is a fact of human life, and so are their parents' reactions to such behaviour. Parents cannot put up with behaviour that is troublesome to themselves and other members of the family. Hence they cannot help expressing, or being inclined to express, pleasure and affection when the child is good and displeasure when it is troublesome. There would be a kind of contradiction in telling a child with an expression of pleasure that he is troublesome, or telling him with a grim expression that he has been a good boy. (The contradiction would be comparable to that of announcing with a smile that one has a toothache.)

The original imparting of moral values is, I have argued, a matter of natural reactions, involving both parent and child, rather than of teaching. It might be thought that the latter comes into its own when moral misbehaviour is distinguished from misbehaviour in general. This seems to be the view of Elizabeth Wolgast. She gives examples of a child being scolded for playing with electricity or for not dressing properly in cold weather: these would be ways—but not moral ways—of being naughty.[7] In such cases, as in a case of stealing, 'the first consideration may be to stop the conduct and prevent its repetition'; but, Wolgast continues, 'if this were all that is involved…, teaching children not to steal would be no different from training them not to play with electricity'.

It is a mistake, however, to suppose that we *teach* children not to steal. What we teach them—or rather, what they pick up in the course of learning their native language—is the meaning of the word 'steal'. And having mastered this, they already know that

[6] F. H. Bradley, *Ethical Studies* (OUP 1962), 171.

[7] *The Grammar of Justice* (Cornell 1987), 170.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

stealing is wrong—that one ought not to do it. In a later passage Wolgast speaks of 'a responsible parent' giving 'a moral lecture on why the action was wrong' (171–2). It is true that such lectures are sometimes given or attempted by parents, but the child's moral development is not dependent on this, and neither is its grasp, in due course, of the difference between the naughtiness of playing with electricity and the naughtiness of stealing. This comes about when the child is able to distinguish between actions that are liable to harm itself and those that are liable to harm others—between such reasons as 'You may get hurt' and 'You may hurt so and so', and between 'It will be bad for you' and 'It will upset your sister'.

The error of supposing that the wrongness of stealing is *taught* to us also affects the idea of moral instruction by a religious authority. The Ten Commandments are sometimes held up as a paradigm of such instruction. But what could we learn from such instructions as 'Thou shalt not steal' or 'Thou shalt not bear false witness'? Knowing the meaning of 'steal', we already know that one ought not to steal. In religious apologetics such instructions (or 'teachings') are sometimes cited as if they were reasons for embracing a particular religion. But this presupposes that we already recognize the values in question, in which case we do not need the instruction.

So far I have dwelt on the negative aspects of morality—the ways in which the wrongness of certain kinds of behaviour is inculcated. A similar account can be given of positive behaviour, such as kindness and compassion. In this case the parents' interest is not in avoiding trouble within the family, but in promoting a pleasant atmosphere, in which smiles rather than frowns, love rather than animosity, predominate. Such preferences, which are natural among human beings, find natural expression in the parents' encouragement of the relevant behaviour in their children. And thus it comes about that instances of behaviour that is kind, generous, etc. are definitive of *'good* behaviour'. 'That's a good boy', we say, with a smile, to the child who has behaved well; and this is one way in which it learns the meaning of 'good behaviour' and 'being good'. Here again, there is a natural reciprocity on the part of children. It is natural for children to feel affection and sympathy for others, and it is natural for these feelings to be capable of development by parental encouragement.

It is also natural for them to be extended, in due course, to people outside one's family and immediate surroundings. (The extent of these developments will, of course, vary greatly with different individuals and the circumstances in which they are brought up.) Here again there is a role for 'natural reactions' on the part of the

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

parents. I said earlier that parents are bound to want to curb the behaviour of children when it interferes with peace inside the family. But the point can be extended to contexts outside the family, for it is natural for parents to want to be at peace with their neighbours; and this desire is unlikely to be satisfied if one's children misbehave towards the neighbours or the neighbours' children. Hence it is not surprising that when we admonish our children to behave in certain ways, this is not confined to behaviour within the family.

There is a corresponding process on the positive side. It is natural for parents to be pleased when their children are praised by others; such praise is flattering to the parents as well as to the child. And one way of attracting it is that of moral behaviour. In this respect, again, the inculcation of moral behaviour is a 'natural reaction' on the part of parents.

Given that the family situation is as I have described, one moral value, in particular, must be of central importance in the child's entry into the moral world, and thereafter. It is that of *consideration for others*. This value is also of central importance to many of the familiar virtues and vices. This is obviously so in the case of altruistic virtues, such as kindness and benevolence; it also applies to respect for property (as in not taking away what belongs to another child), and fairness (as when disputes occur among children). *Dishonesty*, again, is likely to create trouble for other people, while *patience* may take the form of forbearance towards them. In all of these cases the admonition to the child may, and often does, take the form of pointing out that someone in the household would or would not like it if the child behaved in such and such a way. (The question 'How would you like it?', which is commonly used in parental admonition, is a way of directing the child to pay attention to the wishes of others: we want the child to 'put itself in the place' of the other person.)

Children and animals

Let us return to the original inculcation of desirable behaviour, which I described as 'a kind of conditioning'. Negative behaviour is likely to attract negative stimuli from the parents and positive behaviour the reverse; and human children are such that these associations are likely to be 'internalized' and exert their influence even when parents are not present. Now the same points can be made about the inculcation of desirable behaviour in animals. A dog who misbehaves can be conditioned (trained) to stop doing so by regular

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

negative sanctions. Here, as in the case of the child who is a nuisance, the owner (in place of the parent) is bound to react with disapproval, and this is likely to be communicated to the dog and affect it in suitable ways.

We may also compare the situation in the human family with that *within* an animal family, say of chimpanzees. In such families too the young may be observed, at times, to be a nuisance to their parents, provoking rough treatment from the latter; and this treatment may be presumed to curb the troublesome behaviour of the young. Where, we might ask, is the difference? Why should we not regard animals as moral agents, subject to blame and other moral evaluations? This notion is not even as far-fetched as it may seem. There is a history of animals being put on trial for various misdemeanours and being sentenced accordingly.[8]

The crucial difference lies in the fact that human beings can *know what they are doing*. What does this mean, and how does this knowledge manifest itself? It manifests itself when we ask such questions as 'What are you doing?' and 'Why are you doing that?', and are given suitable replies. If a person is unable to give such replies, then we may conclude: 'He doesn't know what he's doing'. But this condition cannot be applied to non-linguistic creatures—or pre-linguistic ones, as in the case of human infants. Observing a dog's behaviour, we can tell that it knows it is going to go for a walk and we can also say that this is the *reason* for its behaviour. But though the words 'know' and 'reason' are appropriate here, they do not have the same implications as in the case of human adults. A reason, in the case of human adults, is essentially what a person would *give*, in answer to the question 'Why?'; and this condition cannot be applied in the absence of human language.

The same contrast can be made in terms of *responsibility*. This quality, too, can properly be ascribed to dogs and other animals. A mole, for example, is responsible for the heap of soil on my lawn. Responsibility can also be ascribed to inanimate things, as when we say that the weather is responsible for the bad harvest. What is different about human responsibility is the ability to be responsive to questions, such as 'What are you doing?' and 'Why?'.

These points about responsibility, and knowing what one is doing, apply also, of course, to moral responsibility. An animal or pre-linguistic child, for instance, cannot be said to know what they are doing in the sense I have explained. And hence our attitudes to their misdeeds cannot be same as in the case of beings who are capa-

[8] E. P. Evans, *The Criminal Prosecution of Animals* (Faber 1987).

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

ble of such knowledge. This is not to say that we cannot feel *angry* towards a dog that has misbehaved—or for that matter, a piece of machinery that has let us down; or that we cannot *blame* them for doing so. But the moral attitude of *resentment* is not appropriate here.[9] This attitude is appropriate when a miscreant can understand the difference between what he is doing and what he ought to do or has been told to do. 'I said you shouldn't do X, and you did it.' 'Why are you doing Y, when I told you not to?' It is when a child is able to understand such questions that he can be said to 'know what he is doing'—in this case, that what he is doing, or has done, is contrary to what he has been told to do. At this stage too the child acquires the ability to ward off blame by means of *excuses*—including 'I couldn't help it', which expresses an understanding of the difference between having, and not having, *a choice*.

A related development takes place when the child learns to respond to a parent's request with 'I don't want to'. There is more to this than a mere failure to obey. A dog may fail to obey the command 'Sit!', but in this case we need only say that the conditioning has not worked. But 'conditioning' would not cover the case of a child who has mastered the language in the ways I have described. When the child responds with 'I don't want to', this is not a failure of conditioning; it is a case of deciding to act contrary to what one has been told to do.

The relation of a dog's behaviour to the command 'Sit!' is more akin to the kind of regularity that occurs in nature than to the obedience or disobedience of a child who has mastered the use of language. There is a regular, though not invariable, connection between the appearance of heavy clouds and the falling of rain; and there is the same kind of connection between the command 'Sit!' and the dog's sitting. It is true, however, that in the case of the dog the regularity has been produced by a process of conditioning—a process in which the desired regularity is itself *created* by a regular combining of behaviour with sanctions. And this process is similar to that which takes place, as I have argued, in the case of a parent and child. It is essential to this kind of training that there are regular correlations both in the training and in the desired result. But with further development in the understanding of language a new kind of correlation appears. It now becomes possible for the child to obey or disobey a request that it has not heard before. There is a correlation between 'Don't go near the water' and the child's not

[9] Cf. P. F. Strawson, 'Freedom and Resentment' in *Freedom and Resentment and other Essays* (Methuen 1974).

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

going near the water, but it is not that of conditioning: it has not been brought about by regular combining of that behaviour with sanctions. And if the child *refused* to obey, by word or deed, its behaviour would also be correlated (though negatively) with that instruction.

It is important to the new, fully linguistic kind of correlation that it has both a positive and a negative side. Thus the child may express its understanding of that instruction by conforming, but also by refusing to conform. He may do the latter by replying 'I don't want to' or by arguing: 'Why should I?' etc. And with this development the child has joined the community of those who can be held morally responsible—who can be asked to justify what they have done or propose to do and whose reasons may be subject to further questioning, etc.

This language-game is, so to speak, grafted on to a more primitive pattern of natural reactions, by a process that is comparable to the way in which an infant's primitive cry of pain has added to it what Wittgenstein called 'a new pain behaviour' (PI 244). The child has hurt itself and it cries. Its parents speak to it as parents naturally would, and in due course the child learns from this that there is another way of expressing what it feels—that of language. And with this development it enters the realm of moral choice, for being able to use language entails that one is able to speak both truly and falsely—to lie, for example, about being in pain, when this is to one's advantage. 'Why?', asks Wittgenstein, 'cannot a dog pretend to be in pain? Is he too honest?' (PI 250). The words 'honest' and 'dishonest' are not applicable to dogs and other non-linguistic creatures because they cannot be called to account for what they have said or done.

In this way animals and pre-linguistic infants may be described as 'innocent'. According to Mill, 'a large proportion of all animals pass their existence in devouring and tormenting other animals. They have been lavishly fitted out with the instruments necessary for that purpose; their strongest instincts impel them to it…'[10] Such behaviour, cruel though it is, is ascribed to 'instinct' because the perpetrators cannot be called to account for it. They do not know what they are doing.

Is 'Stealing is wrong' analytic?

I have argued that learning 'the difference between right and wrong' is a matter of learning how these and other moral terms are

[10] J. S. Mill, 'Nature', in *Three Essays on Religion* (London 1875). He saw here 'evidence to blacken the character of the Creator' (ibid.).

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

used. In learning the meaning of 'steal', for example, we learn that this behaviour is wrong, as well as acquiring a negative attitude towards it. Hence it cannot be an 'open question' whether stealing is wrong, whether it is good to be honest, etc. There *are* open questions in ethics, about which opinions may differ, but these are not among them.

Should we then describe such statements as 'Stealing is wrong' as analytic? This would not be correct because, as with other moral principles, we must allow for cases in which there are countervailing moral reasons. In some cases—stealing the weapon of a murderer, for example—it would not be wrong to steal; while in others opinions might differ. Whether stealing would be wrong in these cases would be an open question. Let us say, then, that stealing is wrong unless there are countervailing moral reasons, and that this is an analytic truth.

But is this really so? What if someone denied that stealing is wrong? Such statements have actually been made by professional thieves in interviews by researchers. 'There's nothing wrong with stealing', such a person might say (perhaps adding '…if you can get away with it'). We could hardly object that such people are ignorant of the meanings of 'wrong' or 'steal'. Must we then give up the claim that 'Stealing is wrong, etc.' is analytic? Should we say that on this question the thief's statement represents a genuine opinion on a question which is, after all, a matter of opinion? This would have a paradoxical consequence. We would have to conclude that the thief who had that opinion would be innocent of malicious intent: in stealing, he would be doing something regarded as wrong by the majority, but regarded as innocent by him. But this would be a curious defence.

But should we take the thief's denial at face value? Perhaps his point is really about countervailing moral reasons; and this might emerge if we questioned him further. Perhaps he thinks that *his* stealing is justified because his needs are greater (or sufficiently greater) than those of his victims; or that the latter, being insured perhaps, would suffer no harm. Or perhaps he would invoke some abstract principle about the injustice of rich and poor, etc. Such replies might of course be challenged and might simply be wrong. But however that may be, the thief's preparedness to give such replies would show that what he means is merely that stealing is not wrong *if* there are justifying circumstances, and not that it is so in the absence of these.

Another point that the thief might be making is that he *feels no compunction* about stealing. This is certainly intelligible, but it

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

would not entail that he sees nothing wrong in stealing. It is a regrettable but familiar fact that people sometimes have no compunction about doing wrong, but this does not mean that they don't know they are doing wrong, or that, because of such ignorance, they are not to blame.

That the thief could not literally mean, or stand by, his statement that stealing is not wrong, would appear if we reminded him of the position of the *victims* of stealing. Such persons would be justified in *blaming* the culprit for the harm he had done them, calling for measures to be taken against him, etc. Now the thief, if he understands what is meant by 'blame' and 'justified', would have to admit that this is so. But then he could not also claim that he had done nothing wrong. If he had done nothing wrong, then the blame and the resentment felt towards him would not be justified.

Those who claim that there is nothing wrong with stealing and other kinds of dishonesty sometimes add '…unless you are found out' or '…if you can get away with it'. But this is a confusion. The wrongness of stealing does not depend on whether the culprit is found out and punished, but on whether, if he were found out, such punishment would be justified. *This* is what the thief would have to deny if he were to stand by his claim. But a thief who denied this would not be credited with understanding the meaning of 'justified'.

Here again the principle of consideration for others is of central importance. Given the statement 'There is nothing wrong with stealing', we may ask ourselves: *What* is wrong with stealing? What does 'wrong' mean in this case? And the answer, in straightforward cases of stealing, is that it is a way of harming others. That is why the victim of a theft has a justified grievance against the perpetrator. And this must be clear to anyone, including the thief himself, who understands the words 'justified grievance' or some equivalent expression. The very concept of *a reason* is at stake here. Someone who did not recognize that a person who has been harmed by another has a reason to feel resentment against the latter, could not be credited with understanding what a reason is.

It might be objected that the conversation with the thief, as I have imagined it, is an idealised one. Who knows what an actual thief might say in such an interview? Perhaps he would not know what to say; perhaps, having become bored or suspicious, he would break off the conversation prematurely. My argument, however, does not depend on risky conjectures about what would actually be said in such a situation. If a thief, or some other person, denied that stealing is wrong, but was unwilling or unable to say what this denial

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

amounts to, then we would not be obliged to take what he had said as a counter-example to what has been claimed. In such cases, however, it is appropriate to consider what the thief's statement *might* amount to, or what a more willing and articulate person *would* say in response to appropriate questions. Only in this way can we establish whether the simple denial was really a threat to the claim that 'stealing' entails 'wrong'.

The concept of a reason; morality and self-interest

I have discussed the ways in which values are 'picked up' in the course of learning one's native language, which includes its moral vocabulary. But a similar account can be given about the very concept of a reason. What is a reason? How do we learn to recognize certain kinds of statements as answers to the question 'Why?'—'Why did you do that?', 'Why should I do this?' and 'Why do you feel resentful against so-and-so?'.

A child cries in the night. 'Why are you crying?'—Silence.—'There must be some reason; tell us what's the matter.' In this and a hundred other situations the child is badgered to produce a reason, and in due course it catches on to this language-game, as it catches on to others. It learns to produce reasons for what it says and does, and in this way it becomes a member of the community of rational beings. But what counts as a reason? What is needed is an answer that is suitable for the behaviour in question. From the child that has cried, for example, we demand an answer that is appropriate for crying. Thus, in learning the concept of a reason, the child must learn what counts as an *appropriate* reason.

Another context in which the concept of a reason—of 'why?' and 'because' and 'therefore'—may be introduced is that of admonishing a child how to behave. 'You shouldn't eat so many sweets.'—'Why not?'—'Because it's bad for your teeth.' Or again: 'Why can't I have that toy?'—'Because you've got one already'; or 'Because it doesn't belong to you'. It is by means of such paradigms that we learn what counts as a reason.

Altruistic reasons are also among these paradigms. Suppose a child asks: 'Why should I stop making that noise?' A suitable answer would be 'Because Daddy doesn't like it', or even 'Because I asked you to'. It might be thought that the second answer would not be of moral significance, but this would be a mistake. The second answer, no less than the first, would introduce consideration for the wishes of others as a reason for acting (or abstaining), and this is certainly

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

a moral reason. (It is this kind of reason that underlies the 'preference' version of utilitarianism.)

The point that altruistic reasons are among the paradigms by which we acquire the concept of a reason is especially important if we turn to the question of conflict between morality and self-interest. It is often assumed that self-interested behaviour is logically prior to altruism, the first being the real paradigm of rationality, while the second has to be justified, somehow, on that basis. But this is not so: altruistic reasons are no less basic than those of self-interest.

Consider an everyday example. Someone stops me in the street to ask for directions. Here I have a perfect reason for giving directions. This is so, not because I enjoy giving directions, or because I would feel guilty if I didn't give them: the fact that someone needs directions is, just by itself, a perfect reason for giving them to him. Here again, using the rhetoric of the Paradigm Case Argument, we might ask: If this isn't a good reason, what is?

It might be thought that this example of giving directions is so trivial as not to deserve the name of 'altruism'. But that would be a mistake. A trivial case of altruism is still one of altruism. We should not think of moral behaviour as special or remarkable (though of course it can be so), but as something that pervades ordinary life. This is so in the original learning situation, in which we first acquire the concept of a reason, as well as in later life.

There is an important difference between learning about right and wrong, as discussed earlier, and learning the language-game of reasons. The former, as I said, is largely the result of admonition, accompanied by suitable sanctions for good and bad behaviour. But in learning how to give and recognize reasons we learn to make explicit, to ourselves and others, that it is reasonable to do what is right even in the absence of sanctions. Such behaviour is among the paradigms by which we acquire the concept of reasonable action.

The reality of moral reasons has been questioned by some philosophers. David Brink introduces us to an 'amoralist', who asks: 'Why should I care about moral demands?'. He 'recognises the existence of moral considerations and remains unmoved' by them.[11] Now it is true that people are sometimes unmoved by moral considerations (though it may be doubted whether anyone could be wholly unmoved by them). But what is the significance of this? Does it mean that to establish the reality of moral reasons, we must (as some 'realists' and 'anti-realists' have supposed) find some

[11] *Moral Realism and the Foundation of Ethics* (Cambridge 1989), 46.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

external standard or entity, beyond the facts of human nature? And if this cannot be done, must we conclude that morality might be, as Kant put it, a mere 'phantom of the brain'? No: the fact that someone remains unmoved by a reason does not impugn the reality of that reason. Suppose that someone—the amoralist, for example—has promised to do X. Then the question 'Why should I care about doing X?' has a straightforward answer: 'Because you promised'. This is a perfect answer, 'than which no better can be conceived', to the question why he should care. This answer does not, of course, entail that he *will* care, but that does not affect its validity. That people sometimes do not do what they have reason to do makes no difference to the fact that they have, and know they have, those reasons.

It is sometimes thought that moral reasons are more exposed than reasons of self-interest to the question 'Why should I care?'. But this is not so. Suppose my larder is empty and you remind me that I ought to get some food before the shops close. It is certainly conceivable that (to paraphrase Brink) I 'remain unmoved by considerations of self-interest' and challenge you to tell me 'why I should care'. The possibility is not even far-fetched. People in a depressed state of mind do express themselves in such ways, and behave accordingly. (Tolstoy, in his 'Confession', is a famous example: 'What is that to me?') But this does not alter the fact that they *have* reasons to act in the appropriate ways.

In his discussion of 'the ultimate ends of human action', Hume pointed out that the quest for reasons cannot go on indefinitely.

> Ask a man *why he uses exercise*; he will answer, *because he desires to keep his health*. If you then enquire, *why he desires health*, he will readily reply, *because sickness is painful*. If you push your enquiries farther, and desire a reason *why he hates pain*, it is impossible he can ever give any…
> Perhaps to your second question, *why he desires health*, he may also reply, *that it is necessary for the exercise of his calling*. If you ask, *why he is anxious on that head*, he will answer *because he desires to get money*. If you demand *Why? It is the instrument of pleasure*, says he. And beyond this it is an absurdity to ask for a reason (*Enquiry* 293).

Here we are given several examples of a reasonable answer to the question 'Why?', but it is, as Hume points out, absurd to push the question beyond its accepted limits. What could it mean to ask someone *why* he hates pain or desires pleasure? To ask such questions would itself be unreasonable.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press

It would be a mistake, however, to assume that the 'ends of human action always come down to these two. Human beings desire a variety of things for which no extrinsic reasons can or need to be given, and altruistic reasons have as good a claim to being 'ultimate' as have the pursuit of pleasure and the avoidance of pain.[12]

The Open University

[12] I am grateful to David Cockburn, Roger Crisp, Peter Hacker, Bob Sharpe and David Walton for comments on previous drafts.

https://doi.org/10.1017/S0031819103000020 Published online by Cambridge University Press