'Attention, Self and *The Sovereignty of Good*' was published in Anne Rowe (ed.) *Iris Murdoch: Reassessment Palgrave Macmillan*, October 2006.

Attention, Self and The Sovereignty of Good

Christopher Mole

In 1958, Elizabeth Anscombe's 'Modern Moral Philosophy' forcefully launched the idea that character traits and states of mind are morally important in ways that could not be seen from the point of view of the contemporary moral psychology. '[M]oral philosophy', she wrote, 'should be laid aside [. . .] until we have an adequate philosophy of psychology, in which we are conspicuously lacking' (p. 1). This idea influenced many ethicists in Iris Murdoch's Oxford, and provides one of the motives for Murdoch's The Sovereignty of Good, and in particular, for Murdoch's attempt to establish that states of mind (construed broadly, so as to include character traits) are morally important for their own sake: that their moral importance is not exhausted by the importance they derive from the actions they might lead one to perform or the states of affairs that they might help to bring about.2

Murdoch's ideas about the moral importance of states of mind are closely related to her views about how one ought to think of oneself. Attention to oneself, she thinks, is a widespread source of moral failure: 'Goodness is connected with the attempt to see the unself';3 'In the moral life the enemy is the fat relentless ego' (p. 51). These two aspects of Murdoch's thought seem to be incompatible. It is natural to think that the states of mind whose moral importance Murdoch establishes are inner occurrences taking place on the private stage of consciousness. But this leads to the following problem. Trying to become good involves giving attention to things of moral importance, and so, if we understand the morally important states of mind to be private inner occurrences, it involves giving attention to private inner states. But this is a form of self-directed intellectual activity, and self-directed intellectual activity is the very thing that Murdoch wants to characterize as a source of moral failure.

There are at least two strategies for avoiding this problem. The first strategy avoids the problem by taking Murdoch's view of the self to entail something less than a complete prohibition on attention to the self. Perhaps self-directed attention comes in different forms, only some of which are prohibited, or perhaps the prohibition on self-directed attention only applies in certain circumstances. Samantha Vice employs this strategy in her contribution to this volume. The present paper pursues a different strategy. It avoids the problem by understanding the morally important states of mind as something other than inner occurrences taking place on the private stage of consciousness. Vice's view is that the strong prohibition on attention to the inner life is neither plausible, nor warranted by Murdoch's position. My view is that we can keep the strong prohibition on attention to the inner, but must reject the idea that the morally important states of mind and character are inner states, and think of them instead as being world involving.

The moral importance of the mind

The claim that the moral importance of states of mind is not exhausted by the importance of their effects can be interpreted in two ways: one strong and one weak. The strong position is that (independently of their effects) states of mind and character often have a crucial role in determining whether a person is doing well or badly, morally speaking. The weak position is that states of mind and character carry some weight, but vastly less than is carried by the moral importance of acts and states of affairs. Murdoch endorses the strong position. The parable of the mother-in-law, which is prominent in her discussion of this point, establishes only the weak position.

Murdoch asks us to imagine a mother-in-law who, by a process of reflection, comes to a positive view of the daughter-in-law whom previously she had regarded as vulgar and noisy (pp. 16-23). It is specified that the mother-in-law's new opinion is not accompanied by any change in outward behaviour (the daughter-in-law is dead or abroad). Murdoch thinks, and expects us to think, that there is something morally good about this change in the mother-in-law. Since, ex hypothesi, the only changes that take place are changes in states of mind, it must be these states of mind that make the moral difference, and so it must be that moral importance attaches to states of mind in a way that does not depend on their effects.

This argument for the moral importance of states of mind depends on the fact that the case of the mother-in-law is one in which action is out of the question. But, for that very reason, the example cannot establish

that states of mind are morally important when the possibility of action is in question. The parable of the mother-in-law shows, in the very special case of entirely mental conduct, that the question of whether the conduct is good or bad must be answered by reference to states of mind or character. It does not show that states of mind are morally important in general, only that moral importance attaches to inner states when nothing else is at stake.

There is another argument, however, that does support the strong position, according to which states of mind often have a crucial role in determining the morality of a person's conduct. Murdoch does not develop this argument in any detail, but we can reconstruct it in terms borrowed from the virtue-based approach to ethics that her work helped to revive. The virtue ethicist takes the proper starting point for ethical theory to be the fact that we should do what the virtuous agent would do were he in our circumstances. We act as we should only if we act as the virtuous agent would. The virtue-ethicist then claims that there is not usually any description of an act given in purely behavioural terms that allows us to settle the question of whether that act is one that the virtuous agent would do. The virtuous agent might lie if the circumstances called for it, but he would not lie callously. He might hurt others, but he would not hurt others brutally. It is all very well to say that "to copy a right action is to act rightly", says Murdoch, quoting Stuart Hampshire's 'Logic and Appreciation', 4 'but what is the form I am supposed to copy?' (p. 29). It cannot be the form of behaviour, considered independently of its motivation. In order to determine whether or not an action is one that the virtuous agent would perform, we need a description of the action that tells us more than is implied by a purely behavioural description. We need a description that tells us about the states of mind and character that the behaviour expresses. These states are important for determining the morality or immorality of a course of action, not because, as in the case of the mother-in-law, we can take away the possibility of action and retain a morally significant inner state - but because the descriptions of action under which we consider actions morally are already laden with commitments to the agent's being in a certain state of mind, or having certain character traits. Actions lose a crucial part of their moral character if we attempt to divorce them from these commitments.

Attention to the self

The argument we have just explored does, as the parable of the motherin-law cannot, provide Murdoch with a reason to believe that states of

mind have an important role to play in our moral thinking. It shows that it is not merely when nothing else is at stake that these states are morally important. But the argument also seems to show that Murdoch is committed to the very un-Murdochian view that the attempt to become good essentially involves attention to oneself.

The problem arises in the following way. In order to know whether we are acting as the virtuous agent would we need to know which aspects of our character we are exercising. It is not enough to know that we are hurting others, or that we are benefiting them. We need also to know whether we are being callous or manipulative. To know these things we must, it seems, pay close attention to ourselves. And that is precisely what Murdoch has told us we must not do.

The prima facie innocence of self-directed attention

What we have just seen is that Murdoch's best argument for the moral importance of states of mind entails a commitment to self-directed attention. This does not pose a problem by itself. The problem arises when this commitment is combined with Murdoch's view that selfdirected attention is a source of moral failure. No problem would arise if we were to reject that view of self-directed attention, and it is tempting simply to do so. Self-directed attention does not seem to be a moral failing. The forms of self-directed attention that we find ourselves committed to by the arguments of the previous section seem particularly innocuous when we consider their role in the moral reasoning at work in the following example: a man is wondering whether he should tell his wife about a minor indiscretion in his past. He recognizes that keeping the secret is a way of being untrustworthy and so he resolves to tell the truth. What moves him is the realization that he does not want to be the kind of person who would continue to lie. The distinctive feature of this form of moral reasoning is that the terms of evaluation it employs indict the agent rather than the act.

Self-indicting formulations often sound more natural than the alternatives from which all reference to the self has been removed, and they are not merely verbal variants on them. The belief that my wife ought to be told the truth rationally motivates me to tell her the truth if I have the desire that things be as they ought. To be motivated to act by the belief that 'I ought to tell her', I need only desire that I do what I ought.

The self-directed attention required for this form of moral reasoning does not seem to be objectionable. It may be clear that the agent who employs this sort of reasoning is not among the best of moral reasoners.

In our best moral thinking our reasons for acting are not provided by concerns about our own goodness. But this does not lead us to think that there is anything wrong with deliberately undertaking the task of acting well, or with being motivated by judgements about whether one is succeeding in that task: the self-directed attention that is required in the making of those judgements does not, on the face of it, rule them out of the attempt to become good, especially when, as in the example above, the judgements are negative ones.

Our argument for the moral importance of states of mind and character carries an apparent commitment to self-directed attention only to the degree that self-directed attention figures in the rather benign sort of reasoning sketched above. Murdoch herself seems to realize that self-directed attention can have a role in the attempt to become good. Her own telling of the parable of the mother-in-law, in fact, seems to involve self-directed attention:

The M of the example [the mother-in-law] is an intelligent and well-intentioned person, capable of self-criticism, capable of giving careful and just *attention* to an object which confronts her. M tells herself: "I am old-fashioned and conventional. I may be prejudiced and narrow-minded. I may be snobbish. I am certainly jealous. Let me look again"

(p. 17, emphasis Murdoch's).

The mother-in-law's praiseworthy change of opinion is clearly precipitated by self-directed attention of the sort that we met above. It is by attending to *herself* that M is in a position to know that she is old-fashioned, conventional and so on. In this passage Murdoch seems to recognize, even to endorse, the mother-in-law's self-directed attention as having a role in her moral progress. Why, then, does Murdoch also seem to think that self-directed attention is a source of moral failing? We have not yet seen any reason to think of self-directed attention in this way. It is neither selfish nor self-aggrandizing, and, moreover, it seems to figure in our everyday attempts to act well.

The rejection of self-directed attention

The self-directed attention involved in the sorts of reasoning discussed above is *prima facie* innocent. It even seems to figure in Murdoch's account of the mother-in-law's praiseworthy change of opinion. It is tempting, therefore, to say that self-directed attention is not always

a source of moral failure, and that it is morally permitted in the service of an attempt to become good. This temptation should be resisted. For Murdoch the struggle to be good is a struggle to keep attention *away* from the self: 'Goodness is connected with the attempt to see the unself' (p. 91).

It is important to be clear that Murdoch does not merely think that self-directed attention has no role in the fully good life. There would be no problem with thinking that the fully virtuous agent does not think of himself, while thinking that we should think of ourselves in order to become good. (Just as there is no contradiction in thinking that a good tennis player gives no attention to his ball-toss when serving, but that we need to attend to our ball-toss rather carefully in order to become good tennis players.) Murdoch's claim is that self-directed attention is absent from the pursuit of goodness, and not merely from the life in which goodness has been achieved. We see this most clearly in the concluding pages of 'The Sovereignty of Good Over Other Concepts' in the claim that 'although [the humble man] is not by definition the good man perhaps he is the kind of man most likely of all to become good (p. 101). The thoroughgoing rejection of self-directed attention is also prominent elsewhere: 'In the moral life the enemy is the fat relentless ego' (p. 51).

Murdoch's thoroughgoing opposition to self-directed attention is clear throughout *The Sovereignty of Good*, but *why* does she oppose it so? The following two passages provide some clues:

The difficulty is to keep attention fixed upon the real situation and to prevent it returning surreptitiously to the self with consolations of self-pity, resentment, fantasy and despair

(p. 89).

We are anxiety ridden animals. Our minds are continually active, fabricating an anxious, usually self-preoccupied, often falsifying veil which partially conceals the world. Our states of consciousness differ in quality, our fantasies and reveries are not trivial and unimportant, they are profoundly connected with our energies and our ability to choose and act. And if quality of consciousness matters then anything which alters consciousness in the direction of unselfishness, objectivity and realism is to be connected with virtue

(p. 82).

The prohibition on self-directed attention, then, is connected to the moral importance of accurately perceiving things as they really are. The accurate perception of things is central to Murdoch's conception of the good life, and of the movement towards it, both because it is good in itself, and because it enables us to make the right decisions about how to act: 'The love that brings the right answer [. . .] is an exercise of justice and realism, and really looking' (p. 89). We shall see, in the next section, that the connection between this imperative to accurate perception and the prohibition on self-directed attention is not an entirely straightforward one. But notice, for now, that attempts to perceive the self are prohibited, and that this is, in part, because they are attempts at a sort of perception that is particularly unlikely to be accurate, and that is likely to impair accuracy when it comes to perceiving things other than the self. The situation we are in is this: we have seen that a role for selfdirected attention in the attempt to become good seems to be entailed by Murdoch's emphasis on the moral importance of states of mind, and that the existence of such a role is incompatible with a strong prohibition on self-directed attention. We have also seen that, although it may be tempting to reject the strong prohibition, Murdoch does not reject it. The problem that we face is one that Murdoch is vividly aware of. It is a source of much of the moral drama that animates The Nice and the Good, published in 1968 (two years before the appearance of The Sovereignty of Good, but two years after the delivery of the lectures that the book of that title collects). The Nice and the Good gives an explicit statement of our problem as a 'great paradox of morality' experienced by the book's protagonist, John Ducane:

What Ducane was experiencing, in this form peculiar to him of imagining himself as a judge, was, though this was not entirely clear in his mind, one of the great paradoxes of morality, namely, that in order to become good it may be necessary to imagine oneself good, and yet such imagining may also be the very thing which renders improvement impossible, either because of surreptitious complacency or because of some deeper blasphemous infection which is set up when goodness is thought about in the wrong way. To become good it may be necessary to think about virtue, although unreflective simple people may achieve a thoughtless excellence. Ducane was in any case highly reflective and had from childhood quite explicitly set before himself the aim of becoming a good man.5

The problem that the authorial voice expresses here is the very problem that we have found in Murdoch's philosophical work. If being good is thought of as involving virtue, then a deliberate attempt to become good (like Ducane's 'explicitly set[ting] before himself the aim of becoming a good man') seems to involve self-directed attention, which 'renders improvement impossible'.

One way to deal with this problem is not to try to eradicate it, but simply to embrace the conclusion that it is not possible to become good by trying. Ducane is sometimes tempted by that view, but Murdoch has no sympathy with it and thinks that we must account for deliberate intellectual attempts to pursue goodness and account for the simple achievement of goodness. '[I]t must be possible to do justice to both Socrates and the virtuous peasant. In such "musts" as these lie the deepest springs and motives of philosophy' (p. 2).

A bad argument against self-directed attention

The escape from our problem is seen by looking more carefully at the way in which the emphasis on 'realism and really looking' provides the foundations for the prohibition on self-directed attention.

It is the imperative to 'realism and really looking' that leads Murdoch to prohibit self-directed attention, but it is not that Murdoch prohibits self-directed attention simply because the self is particularly difficult to really, realistically, look at. We should not interpret Murdoch in this way because if we were to do so we would have to credit her with the following patently invalid argument:

- 1. Accurate perception of the self is difficult.
- 2. We are morally required to perceive things accurately.

Conclusion: We are morally required not to attempt accurate perception of the self.

This argument is clearly not valid. If the first premise only cites the difficulty of accurate self-perception then all that follows is the entirely unremarkable conclusion that, when it comes to the perception of the self, the moral requirement of accurate perception is a difficult requirement to meet. That does nothing to justify the vilification of selfdirected attention. The struggle to become good is, after all, a difficult struggle.

There is a temptation to strengthen the argument by beefing up the first premise. Murdoch sometimes seems to use an argument that is a version of the one above, but one that is less obviously invalid because the first premise has been strengthened so as to say that accurate

self-perception is not merely difficult but impossible, and that the self prevents the accurate perception of other things:

That human beings are naturally selfish seems true on the evidence, whenever and wherever we look at them, in spite of a very small number of exceptions. About the quality of this selfishness modern psychology has had something to tell us. The psyche is a historically determined individual relentlessly looking after itself. [...] One of its main pastimes is daydreaming. It is reluctant to face unpleasant realities. Its consciousness is not normally a transparent glass through which it sees the world but a cloud of more or less fantastic reverie designed to protect the psyche from pain. It constantly seeks consolation, either through imagined inflation of self or through fictions of a theological nature

(p. 76).

The 'modern psychology' that Murdoch understands as revealing this picture of the self is Freudian psychology, and one who thinks that Freudian psychology is wrong-headed will be unmoved by Murdoch's argument here. But even if we were to grant this picture of the self as a source of fictions, we would still not have a premise capable of patching the hole in the argument sketched above. Even if we believe that the self is a source of fictions we still lack a reason for not trying to perceive it rightly. (A Freudian would say that we should try to perceive it rightly.) The premise that is needed to establish an absolute prohibition on selfdirected attention is not just that the self is a deceiver and that accurate self-perception is impossible. The needed premise is that the self is a deceiver, and that its deceptions are made worse by the attempt to perceive it correctly. This is a much harder premise to establish. It is not a premise that can be established on the basis of introspection, for to believe it on the basis of introspection is to undermine one's own basis for believing it. Nor is it a claim that is made obvious on the basis of our observations of others. When we observe others we do, perhaps, find them to be somewhat deceived about the way things are, but the matter of whether they would be any less deceived had they not tried to perceive themselves is not a matter that our normal encounters with people, however carefully conducted, enable us to decide. This claim is too strong to feature as an unsupported premise in the argument against self-directed attention.

There are no plausible strengthenings of the first premise that make the argument given above into a compelling one and so we should not understand Murdoch's prohibition on self-directed attention to be motivated by this argument, or by a version of it in which the first premise is strengthened. We get a more satisfactory interpretation of Murdoch's reasons for prohibiting self-directed attention by getting a better understanding of the strength of the second premise – the premise concerning the imperative to perceive the world correctly.

Realism and really looking

To understand the importance of really looking we must turn once again to the parable of the mother-in-law. The mother-in-law's morally praiseworthy change of opinion involved her seeing that her daughterin-law is 'not vulgar, but refreshingly simple, not undignified but spontaneous, not noisy but gay, not tiresomely juvenile but delightfully youthful' (p. 17). Such changes are said to result from 'realism and really looking' or, as Murdoch says elsewhere, from a 'just and loving gaze directed on an individual reality' (p. 33). Murdoch specifies that the case is one in which the mother-in-law is not deluding herself. The mother-in-law's vision is more accurate once she has revised her opinion upwards. But it must be a contingent fact that in this particular case the higher opinion is the more accurate one. Some daughters-in-law really are undignified, noisy and tiresomely juvenile, and the result of realistically looking at them would, presumably, be to see them as undignified, noisy and tiresomely juvenile.

Murdoch herself has given us the verdict from her realism and really looking: 'That human beings are naturally selfish seems true whenever and wherever we look at them' (p. 77). If 'really looking' is valuable for itself then it is valuable whether or not it shows us things in a good light. In being told to really look we are not being invited to optimistically misperceive the world; we are being asked to perceive it as it really is. 'Really looking' with 'a just and loving attention' is valuable whatever it is that one is doing justice to. It may be a daughter-in-law that one had underestimated, or the extent of human altruism that one had overestimated.

The emphasis on resisting fantasy shows us that it is respect for the real which is the achtung impelling Murdoch to act well: 'The authority of morals is the authority of truth, that is of reality' (p. 88). This is not an emotional reaction to the world as known, but the recognition of the world as something one must come to know (and thereby come to understand and respect):

The value concepts are here [in the case of imaginative art and the practice of a skilled craft patently tied onto the world, they are

stretched as it were between the truth seeking mind and the world. they are not moving about on their own as adjuncts to the personal will [. . .]. [W]e see it as natural to the particular kind of creatures we are that love should be inseparable from justice and clear vision from respect for the real

(p. 88f).

The world-involving nature of the morally important states of mind

The emphasis on the moral importance of character and of states of mind should be understood as an emphasis on the importance of modes of attention. The value that we pursue in developing craftsmanship, and in our engagement with the arts, is 'stretched between the truth seeking mind and the world' (p. 88) because it is value that inheres in attention. (There is a play on etymology here. 'Attention' comes from the Old French ad tendere: being stretched out.) Being loving and just, and possessing the other virtuous character traits, is not a matter of being in a particular sort of private, inner state. It depends on our mode of engagement with the world. The facets of our character take on the status of virtues only insofar as they involve particular faculties of attention that bring us into a virtuous relationship with the world.

Virtuous character traits involve particular propensities for varieties of valuable attention, while vices are tendencies towards inattention. Loving is (or at least, it essentially involves) an astute focus on the particularity of others. Kindness starts with an awareness of their needs. Pity pays heed to the origins of another's misfortune. Courage is the trait required for attention to produce action unimpeded. The angry man, the lazy man and the disrespectful man are all, in their various ways, negligent, careless, thoughtless, impulsive, tactless or rash, as is shown by the frequency with which the harm and offence they cause is inadvertent (tendencies Murdoch explored in her 1971 novel, An Accidental Man). The impatient man is the first to look away.

Possession of a character trait is not a kind of bias in the way that events are depicted in one's internal monologue. There is no distinctive profile to the inner phenomenology of acting from some particular character trait. What is distinctive about acting from a particular character trait is one's form of engagement with the world. It is the world that sets the standards. In trying to act well we must ask (as my argument in the first section demonstrated) questions about our character, but this does not commit us to problematically self-directed attention

because these are not questions that can be answered by directing attention onto oneself. To know whether one's character is virtuous is to know one's mode of attentive engagement with the world, and this cannot be known by looking inwards.

The question, 'Is this act loving?' is not a question about whether the behaviour is accompanied by a particular phenomenological twinge in the subject, but a question about (among other things) whether the act does its object any good and whether it is motivated by a proper recognition of what would do the object good. That is why 'Love needs to be expressed, it needs to do work'.6

Really looking does not get its value by revealing purposefulness and pre-existing value out there in the world: 'If there is any kind of sense or unity in human life, and the dream of this does not cease to haunt us, it [. . .] must be sought within a human experience which has nothing outside it' (p. 77). Nor does it involve an illusory projection of value from the self. Looking at the world is itself a bearer of value. Knowledge of the nature of one's character may be indispensable for the more or less reflective thinker's deliberate progressing towards becoming good, but this knowledge of character is not attained through the worthless unstretched-out attention involved in introspection. Even when introspection succeeds in being honest and astute, the features of ourselves that we learn about through introspection are features that are morally salient only on account of their relationships to things outside the self. Introspective meditations do not bring us into a proper relationship with the world, and they do not tell us whether we are in a proper relationship with the world. It is careful understanding of the world that reveals our failures of virtue as failures. If one takes our moral character to be partially constituted by the ways in which we attentively interact with the world, then one can hold that character traits are primary bearers of intrinsic value without thereby making one's own properties a focus of concern in one's pursuit of goodness. Insofar as the struggle to become good requires knowledge of one's own moral character it provides a further impetus for patient and careful attention to the world.

Notes

- 1. Anscombe, 'Modern Moral Philosophy' in *Philosophy*, XXXI (1958), 1–19.
- 2. Most of this work was completed while the author held the William Alexander Fleet fellowship at Princeton University. An enormous debt of gratitude is owed to Miss Julia Fleet, whose death, while the work was being prepared for publication, is an occasion of great sadness. Thanks for useful discussions are owing to David Sussman and Philip Pettit, to Jessica Boyd

- 3. All page references are to *The Sovereignty of Good*.
- Stuart Hampshire, 'Logic and Appreciation', in Language and Aesthetics (ed.) W. Elton (Oxford: Basil Blackwell, 1954).
- 5. The Nice and the Good (New York: Viking Press, 1968), p. 77.
- 6. The phrase is given to Willy Kost (p. 132). Murdoch probably endorses the view, although only with the additional observation that 'Love can't always do work' (p. 222).

Part III Revisiting The Saint and the Artist