

Moral Weakness, Self-Deception and Self-Knowledge

Denys Turner

In this paper I speculate about the idea of personal moral discovery. These speculations occur, however, only at the end, for my object is to present genuine, not rhetorical, questions. The substance of this paper concerns moral weakness and self-deception. These topics are relevant because the analyses of moral weakness and of self-deception throw light, in different ways, on what is meant by and also on what obstacles there are to the discovery for oneself of a moral point of view.

Briefly the light they throw is this: self-deception is one sort of obstacle to self-knowledge. Insofar, then, as discovering a personal morality is the process of discovering something about oneself—and I merely assert for the moment that it is in part this—self-deception is one kind of obstacle to personal moral discovery. It may follow from this that the contrary process of detecting one's own self-deceptions is one form of self-discovery which is relevant to moral discovery. If so, it will be important to be clear about how we do discover that we have been deceiving ourselves. I try to solve some of the difficulties there are about how this self-detection occurs in the second part of this paper.

The relevance of moral weakness to the topic of self-discovery is more complex. Again briefly, there are two main assertions which I make. The first is that any self-ascription of moral weakness is problematic. This is to say that the subjective evidence immediately available to the moral agent himself on any putative occasion of moral weakness is never *sufficient* evidence that he was, genuinely, morally weak on that occasion. We do not know, in other words, just from what is going on in our minds at the time, that our failure to act as we think we ought to be acting is moral weakness rather than something else.

I consider two ways in which we might be wrong in our self-ascriptions of moral weakness. The first is that we might be deceiving ourselves about the true nature of our moral beliefs and thus that our failure to act in accordance with them is, after all, moral weakness. I spell out what is meant by this in the first part of my paper. This discussion leads naturally into the second part which, as I have said, concerns how we would come to know that we had been in this respect deceiving ourselves.

The second assertion which I make about moral weakness takes the first assertion one step further. I argue that there are cases where the self-ascription of moral weakness is problematic even when we are sure that we are not deceiving ourselves about our moral beliefs. It can

happen that our failures to act as we think we ought to be acting are evidence not that our *actions* are wrong, being condemned, so to speak, by the authority of the moral beliefs, but just the reverse, namely that one's beliefs need revision on the evidence of one's failures. I discuss this paradox in the third part of my paper. The upshot of all this will be, I hope, to bring out some of the complexities of the process of self-, and therefore of moral, discovery. Deciding whether I have been morally weak is just a special case of discovering what my moral commitments are. And just as decisions about my moral weakness are partly decisions based on the evidence, partly decisions on how to view the evidence, so moral discovery generally is partly a matter of deciding *who I am* on the basis of a certain kind of evidence, and partly a matter of deciding *who to be*. Somewhat to disentangle the interaction of these two factors is the general aim of this paper.

I

First of all, then, there is my assertion that any self-ascription of moral weakness is problematic because, on the evidence immediately available to any moral agent, the question whether he is deceiving himself about his moral beliefs is undecidable. By the expression 'the immediately available evidence' I mean this: in every case of moral weakness there are, going on in the morally weak person's mind, certain thoughts and certain beliefs standing in certain relationships to one another, the structure of which relationships constitute the structure of the moral weakness. Let us say that the following formula correctly describes this structure of beliefs: *the morally weak man believes that he is doing something which he thinks that he ought not to be doing*. Since this structure of behaviour is a structure of *beliefs*, and since beliefs are intentional in character, I shall say that this formula describes the *intentional structure* of moral weakness. Furthermore, for the same reason that this structure is a structure of beliefs, this structure is available to a moral agent in the form of immediate evidence. For any man is, it seems, an authority on what, on a given occasion, he was *thinking*. I may be wrong, on an occasion, in believing that what I did was contrary to my moral convictions; but I cannot be wrong, surely, that I believed that I had acted so. Hence, a man's belief that he was doing something which he thought he ought not to be doing is an intentional structure of behaviour which has the character of evidence, immediately available to him, of moral weakness.

My assertion, therefore, that the self-ascription of moral weakness is always problematic is the assertion that just because my behaviour had the intentional structure defined by the formula it does not follow that I have actually been morally weak. Another way of putting this assertion is to say that moral weakness is not identical with its intentional structure. The contrary assertion is the assertion that self-ascriptions of moral weakness are incorrigible, the assertion, namely, not merely that only the moral agent in question can be sure that he has been morally weak, but also that the moral agent in question can never be wrong that he has been morally weak. This view has some plausibility, at least for

some cases, as the following example shows. I am the Count in *The Marriage of Figaro*. I believe that the woman in the bower is Susanna, with whom I plan to be unfaithful. As it turns out the woman is the Countess, my wife in disguise. In kissing her I think that I am being unfaithful to my wife. The fact that I kiss a woman I thought I ought not to kiss means that I have been morally weak. The fact which I later discover that the woman I kiss is a woman it is not wrong to kiss, does not permit me to withdraw the self-ascription of moral weakness. For I thought I was doing what I thought I ought not to be doing: my behaviour thus exhibits the relevant intentional structure. And, it might be thought, it is the intentional structure of the behaviour which gives it the character of moral weakness, not what extensional statements are true of it.

My answer to this example is that it shows only that a necessary condition for the self-ascription of moral weakness is the occurrence of this intentional structure. That it is not, however, a sufficient condition, is shown by the possibility that a person, the intentional structure of whose behaviour is described by the formula, might nonetheless be deceiving himself that he believes what he thinks he believes. Now if you allow this possibility at all the interest of it lies in the fact that there is no difference whatever, *in intentional structure*, between cases where the belief I have acted contrary to is a genuine moral belief and one which I deceive myself I believe. Thus, as I have put it, on the immediately available evidence, the question of my moral weakness is undecidable. This is because the self- or other-ascription of self-deception rests on the truth of certain extensional statements, on evidence, that is to say, other than what I have called the 'immediately available'. And so we must ask what that evidence is on the basis of which we decide between genuine and self-deceived belief or, generally how it is that we come to know that we have been deceiving ourselves.

II

It should not be thought that the ensuing remarks are an exhaustive account of self-deception. In particular, I have nothing to say either about *how* we deceive ourselves, still less about how it is *possible* for a man to deceive himself. I have something to say only about how we *undeceive* ourselves or uncover the self-deception of another. There is a problem about this process which is that the self-deceiver has a version of his behaviour—a story about what he has done and the reasons for which he acted—from which the undeceived observer dissents. The agent paints one picture of his behaviour, the observer another, each claiming that his picture is the truer and more successfully explanatory.

We do grant that it is in principle possible to detect a man's self-deception. Nonetheless it is difficult to say *how* this is done. For it is a peculiarity of some cases of self-deception that, for any one action which the observer claims an agent does for self-deceived reasons, the observer may not be able directly to falsify the agent's own account. And it *may* be the case even that in *every* instance of the behaviour, each instance

taken by itself, there is no way of directly falsifying the account which the agent gives of that behaviour.

This may seem odd. But in fact it is characteristic of the most complex forms of self-deception that what the agent claims as his motives are motives which he is aware of having and are, furthermore, motives which would explain, in normal circumstances, what he does. Thus, in this sort of case, the agent's own picture of any one occurrence of his suspect behaviour, may 'fit the facts' and explain it by normal, everyday standards of explanation. This should not be surprising. It accounts for the fact that the self-deceiver is sincerely convinced of his own version, which would itself be surprising if there were no facts which it fitted and if by no received standards it at all plausibly explained what he did.

Suppose, for example, that I believe that I smacked my child merely out of a father's natural desire to correct it before it acquires habits dangerous both to itself and to others. Now it may very well be true that I believe it to be right to smack my children in certain circumstances and that I believed these to be circumstances of that sort. It may even be true that these *were* circumstances of that sort. Given this we have data enough to explain, by ordinary deductivist canons, my smacking of the child, for the hypotheses are known to be true, the initial conditions hold, and the conclusion which follows is in accordance with the facts. I did act as a person with my beliefs would be expected to act in those circumstances.

Nevertheless, the observer claims (or else I afterwards come to see) that I was deceiving myself, that that was neither what I was doing nor were the motives I admitted to the actual motives which led me to do it. The observer claims, say, that 'what I was doing' (in smacking my child) was 'punishing my wife'; and that I was doing this out of jealousy for her preferring my child to me. The observer may admit that my version of my behaviour was true and explanatory. All the same, he denies that it gives the true explanation.

Now the observer's alternative explanation cannot be got from nowhere. It cannot, if he claims that his is *the* explanation, be just a fiction he invents which, like mine, covers the facts and leaves nothing unexplained. For its just *happening* to do this is insufficient to defeat my explanation, particularly since my account at least has the merit of mentioning motives I am correctly aware of having while his, the observer's, has the disadvantage of mentioning motivations of which *ex hypothesi*, I was unaware. From what source, therefore, is the observer to derive the evidence for his story: Obviously this source can only be that to which I appeal as evidence for mine, namely, my acknowledged beliefs and observable behaviour taken as a whole. It is this identity of sources for different and possibly inconsistent stories which appears to have an air of paradox about it.

The supposed paradox can be stated in quite general terms as deriving from a well-known property of deductive kinds of explanation. It is a well-known and logically indisputable truth that for any given event described by E there is an indefinite number of hypotheses from which, together with statements of initial conditions, E can be deduced.

Some of these hypotheses may be true, but they needn't be for E to be validly deducible from them. For, of course, a true conclusion may follow, in a valid deductive inference, from false premisses. Now if the only requirements for explanation were those for valid deductive inference, then any set of true statements or any set of false statements from which E can be validly deduced would equally 'explain' the event described by E.

This analogy suggests a distinction between two classes of self-deception. The first is of the sort already described, and is analogous to the explanation case where we have true premisses and a true conclusion and therefore *an* explanation, but not *the* explanation. The second kind of self-deception would be different from this and would be analogous to the explanation case in which, though we have premisses of a valid deductive inference yielding a true conclusion, the premisses, or at least one of them, is known to be false. This second kind of self-deception is therefore, of the sort where we have evidence that the beliefs and motives which the agent acknowledges as his are beliefs and motives which, though explaining what he does, he does not actually have. He deceives himself that he believes that corporal punishment is an effective or desirable device for correcting children; or else, though genuinely believing this, he deceives himself that this was an appropriate occasion.

Now in either kind of self-deception, the structure of a man's self-deceived story can sometimes be very complex and sophisticated, not to say devious. It can thus be very hard to penetrate the veil, for, as I have said, it may actually 'cover all the facts'. The self-deceiver's view of himself seems then to fit so very snugly that the observer is left only with the feeling that something is wrong with it, perhaps only with the feeling that the cover story is that bit *too* snug for belief. Otherwise the observer may have to admit that if the relationship between the behaviour and the covering explanation were conceived of as deductive, then it would be faultless, since the explanation does indeed provide premisses from which the behaviour can be validly deduced.

What, then, can be the justification for the observer's scepticism? In the second sort of case, that where the agent is deceiving himself about the beliefs he holds, the analogy with the deductive methods of natural science can be fairly easily extended. In the natural sciences we will often have readily available independent evidence that, though premisses in a valid deductive inference with a true conclusion, the explanatory statements are false. This evidence may either simply be that they are inconsistent with other *theoretical* statements which we know to be true, or at least better founded; or else it may be that though consistent with the truth of E they are inconsistent with other *observation* statements which we know to be true. A good part of the business of learning to do natural science consists in learning when and how to decide that a *formally* valid explanation has to be rejected in the light of stronger evidence.

Analogous to such scientific decision procedures are, for cases of self-deception of the second sort, our ordinary tests for the genuineness of beliefs. These are, however, very many and very varied in character,

ranging from very blatant inconsistency, which the agent appears not to advert to, among the beliefs themselves, to the apparently stronger sorts of test of action which prescriptivists, for example, stress so much. Inconsistency is less a test than a clue. We are entitled to suspect that a person might be deceiving himself in justifying a particular action by appeal to quite outrageous beliefs which, we think, he ought to be able to see do not square with what he ordinarily believes. Nonetheless, suspicion is not proof. A man is not necessarily deceiving himself just because he holds, unawares, inconsistent beliefs. On the other hand the test of action is not decisive either. From the fact that he does not always perform similar actions when the same reasons would hold as hold in this case—say, he only punishes his son, but never his daughter—does not *prove* self-deception. In itself, this fact is evidence as strong for moral weakness as for self-deception. To establish self-deception about moral beliefs rather than moral weakness in respect of them one needs to apply a whole battery of tests over a large amount of evidence derived from continuous stretches, not just episodes, of a man's moral career. And judgments based on such evidence are necessarily of a pretty uncertain nature.

However difficult it may be to detect such cases, self-deception of the first sort presents yet more difficulties. For in this sort of case all the evidence derived from the tests for genuineness of the beliefs is in their favour. On the other hand the claim is made, and we want to know how it could possibly be established, that he deceives himself that it is *these* beliefs, albeit genuinely held, which led him to act as he did, even though what he did was exactly what one would have expected him to do, given his avowed beliefs. So, for example, on one occasion, with Hamlet, who decides not to murder Claudius because, being at prayer, the victim would immediately go to heaven—no revenge at all. Well, it is true, let us say, that Hamlet believes that those who have confessed their sins go to heaven and that those who go to heaven are for ever happy and that he wants unhappiness for Claudius. Hence he has every reason for not murdering Claudius then. And yet it is traditional to regard Hamlet's reasoning as self-deceived, as being just that bit too clever to be true. What can be the evidence for this sort of assessment?

My earlier remark that the evidence *against* the self-deceiver's story must be derived from the same source as that which provides the self-deceiver with his evidence *for* his story must be taken to refer generically to his behaviour as the common source. The difference between the self-deceiver's and the undeceiver's stance does not, therefore, consist in their having available different *sources* of facts, but in the difference in story-line which they take the same facts to be telling. The difference does not concern what facts are relevant, but consists in a difference over *how the facts are to be viewed*. Briefly the difference can be reduced to the self-deceiver's insistence that each item in the list of evidence has to be treated individually and in isolation; the undeceiver, on the other hand, insists that the evidence has to be regarded cumulatively and systematically.

Thus, when the undeceiver has marshalled the evidence for his case, the case for the defence retorts by taking each item in the list separately and by showing how, in each case, there is always another way of viewing the matter. He deals with the prosecution's case as defence counsels deal with evidence which is circumstantial. It is only *taken as a whole* that circumstantial evidence is convincing. Item by item it can always be explained away. What, on the other hand, I have to do if I am to reveal to a person his self-deception is to show how, when taken as a whole, his behaviour tells a very different story from that which it tells when that behaviour is explained bit by bit. Thus what the self-deceiver is resisting is not the facts, but a way of viewing them: and he does this by denying that they fall under a single, systematic view at all.

It is, therefore, only when the self-deceiver is constrained by some decisive evidence from his own behaviour to *re-classify* the evidence of his actions, only when he can perceive the system in the apparently fragmentary, that he perceives his self-deception. So long as he does not perceive the system in the apparently random he can be maintained in his self-deceit. He does not have to deny *the facts* about what he did, or what happened, or how he reacted in each case. He may admit, on having it pointed out to him that he did smack then, and often does smack the child too hard and that he never smacks his daughter and so forth. But one case he explains, with perfect plausibility perhaps, as due to tiredness and consequent loss of temper; another case he puts down to simple lack of judgment about what would be fair punishment for the child. He explains his apparently unrelated failure to remember the child's birthday as 'sheer, unforgivable absent-mindedness' and the fact that he always manages to buy the child hopelessly frustrating toys as deriving from a desire to stimulate his imagination, and so forth. But the picture of overall hostility towards the child entirely evades him and thus a general account of the isolated instances. The mechanism of his self-deception is thus that of keeping the instances isolated. And when we say that he deceives himself we imply that he keeps the instances isolated *in order to avoid* having to acknowledge the overall picture.

Two points emerge from this rough sketch of self-deception. The first is that the reason why one may not be able to tell, in any given case, whether a man is deceiving himself, is that there are no standard descriptions of mental states or of behavioural symptoms which necessarily characterise every self-deceived action *as* self-deceived. This is, in fact, the same point as that which I made about moral weakness in the first part of this paper. It is that with self-deception as with moral weakness, the evidence is never direct and immediate, revealing, so to speak, by the expression on its face, how it is to be understood. In the examples which I have been discussing, there was nothing in the description of any one of the actions which the observer claimed were collectively and circumstantially evidence of self-deception, which, taken on its own, required one to say that it was self-deceived action rather than moral weakness. It was only when the observer had pieced together a different

sort of picture of the man's behaviour as a whole that the 'evidence' began to fall into place.

The second point which emerges is that the unmasking of one's own self-deception always has something of the character of a moral *revolution*. This is shown by the way in which as a result of this self-detection one is required, as with all revolutions, to re-write the past, one's history, in this case one's own personal history or autobiography. Perhaps it might help to make this clear if we compare this situation with that which Kuhn describes (with what degree of correctness I cannot say) as obtaining in periods of 'revolutionary' scientific change. For Kuhn, in a period of 'normal' science the work of the scientist is carried out within 'paradigms' or standards of acceptable problems and acceptable-in-principle solutions. Much is explained, knowledge is much extended within these standards of explanation, but inevitably much is left unexplained and at least some matters actually *resist* explanation. A period of revolutionary scientific change occurs only under two conditions, both important from the point of view of my comparison. First of all there accumulates a great deal of *anomalous* information, data or observational material which cannot be explained within the standardly accepted paradigms of problem-stating and problem-solving. And secondly a revolutionary change occurs only when a new paradigm is available in which both the established achievements of the previous era of scientific activity and the anomalous material inconsistent with that material can be drawn together within a single new picture of properly scientific activity.

A man detecting his own previous self-deception is in a parallel situation. His detecting it involves him in a task of re-writing his past as he had hitherto known it according to a new 'paradigm' or, as we might say, self-image. But this re-writing is not forced upon him simply by the observed discrepancies between his beliefs and his actions, as it were by his noticing that some of his actions would tend to refute the validity of his belief-claims. For, as with scientific theories, the weight of presumption is always conservative. It is not sufficient, therefore, to require this re-writing that I find that I do not always act as, on my professed beliefs, I could be expected to act. For a typical move of the self-deceiver is to acknowledge this discrepancy but deny that it actually requires revision of the belief-claims at all. Now one way in which a self-deceiver can avoid self-detection is, as I pointed out in the first part of this paper, by assigning the discrepancy to moral weakness. The devious purpose of this move of self-deception is secured by the fact that a man who describes his failure to act as he thought he ought to be acting as moral weakness can consistently maintain that he *really does* believe what he claims he believes. One form of self-deception, then, consists in the false admission of moral weakness, which, if an uncomfortable admission, may be less so than the admission of self-deception.

For moral weakness, if sitting uncomfortably with one's belief-claims, is at least not a complete scandal to them. One can say, therefore, that in some cases the *point* of admitting to moral weakness is precisely in

order to avoid having to alter one's view of one's beliefs and reasons for acting. In a situation requiring revolutionary revision the admission of moral weakness is a counter-move of reaction. What, on the other hand, is scandalous to one's belief-claims is the recognition of genuine anomaly, evidence of a relationship between belief-claims and actions which requires rejection of those claims, either, in the first sort of self-deception, as explaining the actions in question or, in the second sort of case, as being genuinely held beliefs at all.

It is the second of the two conditions for Kuhnian revolutions which has its parallel here: for to recognise that one's behaviour is in some respects anomalous vis-a-vis one's belief-claims is to recognise that some alternative account of one's behaviour is required, that some alternative view of oneself has to be pieced together. This recognition is the recognition that one is—and that all the time one *was*—a different sort of person from the sort of person one thought one was. It is in this non-arbitrary, factually grounded way that one is constrained to re-write at least in part, one's autobiography. In the new version the hitherto incidental and random evidence, marginalised under the description of being mere moral weakness, is drawn into the centre of the picture as evidence of a central motivational plot. 'All the time it was jealousy which led me to act as I did', we eventually admit. 'And I was unable to see this because I was assuming a picture of myself as a generous, open-hearted person, albeit with the average mortal's weaknesses'.

The upshot of this discussion of moral weakness and self-deception seems, therefore, to be this: the self-ascription of moral weakness is always problematic, since one has to allow for the possibility that such self-ascriptions are self-deceived. But the criteria we use for deciding for ourselves whether we have, genuinely, been morally weak or alternatively are deceiving ourselves about this, establish something about how we are to view ourselves as personal, autobiographical, beings. What we discover, in determining which is the true description, is something about *who we are* as evidenced by *what we have been*. And sometimes such self-discoveries can be characterised as revolutionary, since they require a radical revision of the way we are to view ourselves.

III

There is, however, another aspect to this scenario for moral weakness which I mentioned in the first part of this paper and I will conclude by making it explicit. This further aspect to moral weakness arises from the fact that though, normally, to say that one is morally weak is to commit oneself to the judgement that one's behaviour should be changed there are occasions when the judgment is reversed; where, in other words, one's moral weakness is evidence, not that the behaviour should be changed, but that one's moral convictions should be changed.

If this is not entirely clear examples are not hard to find which are. In *Jude the Obscure* Hardy describes the vacillations of Sue Bridehead in her relations with two men, Phillotson, the conscientious, morally unimaginative and physically repugnant school-teacher and Jude, the

explosively romantic, idealistic though unstable stone-mason. Sue, herself sexually inexperienced if not frigid, marries Phillotson but falls in love with Jude and is psychologically, though not physically unfaithful to her husband. At first, viewing her relationship with Jude in the light of conventional morality, she cannot but describe her love for Jude as a weakness. But this conventional view is slowly transformed into evidence for her not that she is morally weak, but on the contrary, that her marriage is unnatural, humanly destructive, as being a suffocation of what she discovers to be her best, most liberating desires. She thus discovers, via what she called her weakness, not false, but true desire—what, in one sense of that expression, she *really wants*.

Later, however, she returns, almost neurotically guilt-ridden, to her husband, as being what God and nature had always ordained. Nonetheless she is still, in her eyes, unfaithful with and deeply in love with Jude. But now that she once again accepts her marriage as right, she has to *redescribe* her infidelity, once again, as moral weakness.

This sort of case brings out, interestingly, two apparently quite different but in fact related points. The first has to do with the complexities of nuance there are concerning the way in which our moral judgments enter into our ascriptions of moral weakness. The second concerns the exactly parallel complexities in our ways of deciding what it is that we 'really want'. What is interesting about the connectedness of these two points is what they suggest (though by no means prove) about the connection between discovering what to praise and blame in one's own behaviour and the discovery of what it is that one 'really wants'—or, as we might put it, who one *really* is. I make no case for this connection, or for any significance which it may have, beyond saying that deciding whether one has been morally weak is partly based on the evidence of what one 'really wants' and is partly a way of determining what one 'really wants'.

I consider only some of the complexities concerning the idea of what a man 'really wants'. It is often said, with surprising indifference to the ambiguity of saying it, that moral weakness reveals more convincingly than any other text what a man is 'really like'—or which in the context amounts to the same thing, what he 'really wants'. This comment is usually meant in a derogatory sense, suggesting that, his professions to the contrary notwithstanding, a man's true colours are revealed in his failures. In this sense, what a man 'really wants' is shown by what he actually does; for in this sense, those wants are most 'real' which are the actual *causes* of a man's actions.

It is certainly true that what a man actually does can, sometimes, show what desires caused him to act, though the psychoanalytic evidence works against this as a general presumption. All the same, let us grant that there is a sense, if somewhat problematic, in which what a man 'really wants' is shown by what he does, and most revealingly in his failures. In this sense of 'real' a man's professions of moral belief can look somewhat less than 'real'. For, *ex hypothesi*, on occasions of moral weakness, those beliefs do not cause him to act.

Nonetheless this is only one way of viewing the contrast between 'less real' and 'more real' in respect of wants. In another sense of 'really want' a man who is morally weak can characterise those desires which caused him to act in a morally weak fashion as 'less real' than those which *he would rather had* caused his actions. For those desires he gave in to are, he may say, less really *his*, they express less precisely, or not at all, what it is that he is *committed to*. And so, though recognising that those were *his* desires and that they did lead him to act in a morally weak fashion, he *disowns* them as belonging to another 'self' than that which he 'really wants' to realise.

The two senses of 'really want' therefore trade upon quite different pairs of contrast. And in the ascription of moral weakness we are always called upon to *decide* which, of the conflicting wants which make up the structure of the situation, we are going to call 'real' and in which of the two senses we are going to do this. The possibilities here are aptly illustrated by the change in Sue Bridehead's description of her love for Jude. All along it is, if you like, one and the same desire which is being described. At first, thinking she would be faithful to her husband, but failing in this, she is aware of her love for Jude as a causally more real force in her life. As this love begins to reveal to her the unreality of her marriage to Phillotson, it becomes not just a more vigorous psychological force, but a flood of light shed upon her truest, most basic desires, a shift of moral axis. When once again, the moral axis shifts, it shifts also around the idea of what she most really wants. Her love for Jude is still the most potent force; but it is now more real only in the sense of what will cause her to act contrary to what she thinks best, contrary to her 'best' desires.

I am as aware as anyone who has read the novel that I have rather forced the story into a straight-jacket. Nonetheless, when fully nuanced, it does, it seems to me, illustrate an important aspect of the idea of moral discovery. All the occasions when a man is *prima facie* morally weak are in this way opportunities for self-discovery. Leaving aside the question of self-deception, it is for the reason that moral weakness both *raises* the question about what one really wants and only partially *decides* it, that self-ascription of moral weakness are always problematic. One could, without exaggeration, go so far as to say that whenever one fails to do something which one thinks one ought to be doing one is called upon to *decide* whether this is moral weakness. I may do this, if I am sure of my convictions, *on the evidence* of my failure. But this evidence of failure is not, *by itself*, conclusive. For what to one man is evidence of weakness is to another a discovery of a new, more compelling commitment. But in either case it is true to say that I did not do what I thought I ought to be doing.

In the end the importance of the fact of moral weakness in the moral life is that it provides such opportunities for deciding. It may even be argued that a good part of what is involved in moral discovery is bound up with such decisions. Incidentally it seems to me (as it also seemed to Hardy) that in this process of self- and moral discovery, the formulation

of moral principles has a relatively unimportant role. Indeed, in view of the way in which men can be hounded by their moral principles into spurious, though guilt-ridden, admissions of weakness, they are often a positive menace. In the meantime, in a culture dominated by the morality of moral principles, it is perhaps just as well that men often fail to live by them. For most of us it is only by this Kierkegardian route that we have a chance of learning what we might come to enjoy, as best answering to what we are, have been and might yet come to be.

Objections to Lonergan's Method

Fergus Kerr OP

Looking at Lonergan's Method is a collection of thirteen papers issuing from a conference held at Maynooth in the spring of 1973 at which scholars from differing traditions and disciplines gathered to assess the significance of the recent work of Bernard Lonergan.¹

These are not the papers written for the meeting; they represent the authors' reflections after it, in the light of the discussions that took place. Perhaps the momentum for the book was the convergence of fundamental doubts about the viability of Lonergan's method. At any rate, for all the respect and gratitude that some of the Catholic contributors voice for what Lonergan has done over the years to loosen the hold of a certain way of doing theology, it is very striking that all but one or two of these papers make what seem such irreparably damaging criticisms of his recent work that it becomes very difficult to regard it any longer as a promising trail in the reconstruction of Catholic theology

¹*Looking at Lonergan's Method*, edited by Patrick Corcoran SM. The Talbot Press, Dublin, 1975. 193 pp., £3.