

## 6

# Moore on the Right, the Good, and Uncertainty

*Michael Smith*

In *Principia Ethica* G. E. Moore famously argues that there is an analytic connection between facts about which actions are right and wrong and facts about the goodness and badness—that is, the value—of actions' outcomes.

All moral laws, I wish to shew, are merely statements that certain kinds of actions will have good effects. The very opposite of this view has been generally prevalent in Ethics. 'The right' and 'the useful' have been supposed to be at least *capable* of conflicting with one another, and, at all events, to be essentially distinct. It has been characteristic of a certain school of moralists, as of moral common sense, to declare that the end will never justify the means. What I wish first to point out is that 'right' does and can mean nothing but 'cause of a good result,' and is thus identical with 'useful'; whence it follows that the end always will justify the means, and that no action which is not justified by its results can be right. That there may be a true proposition, meant to be conveyed by the assertion 'The end will not justify the means,' I fully admit: but that, in another sense, and a sense far more fundamental for ethical theory, it is utterly false, must first be shewn. (Moore 1903: 146–7)

Though Moore here claims that it is analytic that right acts maximize value, he later came to amend this under the influence of Russell (Russell 1910, Moore 1942). He subsequently suggested that it is a priori, but not analytic, that right acts maximize value. But though even this weaker claim can be and has been challenged (Rawls 1971, Scanlon 1998), I will take the stronger claim that it is analytic for granted in what follows. For what interests me is not the status of the claim that right acts maximize value, but rather what those who accept it, whether as analytic or a priori, should say when they are reminded that we cannot be certain about the consequences of our actions. This is the question that Moore goes on to address in the relevant section of *Principia Ethica*. Before proceeding, however, let me anticipate a couple of misgivings.

First, as we will see, Moore moves seamlessly between talk of defining rightness in terms of the maximization of value and talk of defining duty in such terms. But, it might be objected, to say that an act is right is to say that it is permissible, and hence that it is not the case that one has a duty not to do it. It is not to say that it is one's duty to do it. I will, however, overlook this difference in what follows. If it is one's duty to maximize value then acts that are permissible, but not one's duty, are presumably those that produce at least as much value as some other act that maximizes value. Right acts maximize value just the same as duties. As between the options that one faces, one's duties are simply those acts that are uniquely right.

Second, far from being analytic, or even a priori, some might think that Moore's consequentialist account of right action isn't so much as true. As such, they might say that they have no interest in what one who accepts it should say in the light of uncertainty. But this would be a serious mistake. For the kind of consequentialism that is true, even if what Moore says is analytic, is a kind of consequentialism that deontologists can and should happily accept (Dreier 1993, Louise 2004). While Moore's own view was that value is a simple property, and hence would presumably have favoured a formalization of his definition of right action along the following lines:

(x)(x's  $\phi$ -ing at time  $t$  is right iff x's  $\phi$ -ing at  $t$  maximizes value)

we can allow that values might be relativized to persons and times in a way congenial to deontology (Smith 2003). A formalization of the following kind might therefore be more accurate:

(x)(x's  $\phi$ -ing at time  $t$  is right iff x's  $\phi$ -ing at  $t$  maximizes value <sub>$x,t$</sub> )

My keeping my promise now might maximize value<sub>me,now</sub> and so be the right thing for me to do, even though my now breaking my promise would maximize value<sub>me,later</sub> and perhaps even maximize value<sub>me,over time</sub>. Moore's puzzle about the impact of uncertainty on the definition of right action is thus a puzzle for everyone.

The paper is in four main sections. In the first I spell out Moore's view of the way in which uncertainty affects the proposed definition of rightness in terms of the maximization of value. In the second section I compare Moore's view with an alternative put forward more recently by Frank Jackson (1991). In the third and fourth sections I offer my own account and say why it should be preferred to both Moore's and Jackson's views. To anticipate, it turns out that Moore and Jackson are both right about something and wrong about something. The correct view combines elements from both.

## 1. Why Moore thinks that uncertainty is a problem

Why is uncertainty an issue, given the proposed definition of 'rightness' in terms of 'maximization of value'? Moore explains the problem this way.

In order to shew that any action is a duty, it is necessary to know both what are the other conditions, which will, conjointly with it, determine its effects; to know exactly what will be the effects of these conditions; and to know all the events which will be in any way affected by our action throughout an infinite future. We must have all this causal knowledge, and further we must know accurately the degree of value both of the action itself and of all these effects; and must be able to determine how, in conjunction with the other things in the Universe, they will affect its value as an organic whole. And not only this: we must also possess all this knowledge with regard to the effects of every possible alternative; and must then be able to see by comparison that the total value due to the existence of the action in question will be greater than that which would be produced by any of these alternatives. But it is obvious that our causal knowledge alone is far too incomplete for us ever to assure ourselves of this result. Accordingly it follows that we never have any reason to suppose that an action is our duty: we can never be sure that any action will produce the greatest value possible. (Moore 1903: 149)

The problem is thus supposed to be that, since we can never be certain of all of the effects of the actions we perform, still less all of the effects of the actions which we don't perform but which were options for us, it follows that 'we never have any reason to suppose that an action is our duty.' But if this is the problem, then it seems to be rather overblown. If we cannot be certain of the effects of our actions then, let's agree, there can be no *conclusive* reason to suppose that any particular action is our duty. But that doesn't entail that there are no reasons at all to suppose that any particular action is our duty—or rather, that would follow only if reasons for belief had to be conclusive reasons, and there seems to be no good reason to suppose that this is so.

As subsequently becomes clear, however, the claim that there is no reason at all to suppose that any particular action is our duty is not crucial to the problem Moore wishes to raise. He goes on as follows.

Ethics, therefore, is quite unable to give us a list of duties: but there still remains a humbler task which may be possible for Practical Ethics. Although we cannot hope to discover which, in a given situation, is the best of all possible alternative actions, there may be some possibility of shewing which among the alternatives, *likely to occur to any one*, will produce the greatest sum of good. This second task is certainly all that Ethics can ever have accomplished: and it is certainly all that it has ever collected materials for proving; since no one has ever attempted to exhaust the possible alternative actions in any particular case. Ethical philosophers have in fact confined their attention to a very

limited class of actions, which have been selected because they are those which most commonly occur to mankind as possible alternatives. With regard to these they may possibly have shewn that one alternative is better, i.e. produces a greater total of value, than others. But it seems desirable to insist, that though they have represented this result as a determination of *duties*, it can never really have been so. For the term duty is certainly so used that, if we are subsequently persuaded that any possible action would have produced more good than the one we adopted, we admit that we failed to do our duty. It will, however, be a useful task if Ethics can determine which among alternatives *likely to occur* will produce the greatest total value. For, though this alternative cannot be proved to be the best possible, yet it may be better than any course of action which we should otherwise adopt. (Moore 1903: 149–50)

The real problem that concerns Moore is that there is no plausible task for Practical Ethics in the absence of certainty about which actions are right. There is no way of getting from the abstract knowledge of what it is about an act that is right that makes it right to a decision about what to do in concrete situations in which we have only limited knowledge of consequences.

There are two parts to Moore's response to this problem. In the first he details a 'humbler' task for Practical Ethics. Even though we can never be certain which acts maximize value, there is, he tells us, still something 'useful' we can establish, namely, 'which among the alternatives, *likely to occur to any one*, will produce the greatest sum of good.' In other words, in the absence of certainty about which actions maximize value, we can still say which acts are *likely* to maximize value, given our limited knowledge. And then in the second part he argues that, since what we would conclude if it were to emerge that an act that we classified as 'right' on the basis of our limited knowledge didn't in fact maximize value is that our classification of the act as 'right' was in error, it follows that acceptance of this humbler task for Practical Ethics gives us no reason to question our original definition. When we call an act 'right' what we mean to be saying of it, even when we engage in the humbler task, is that that act maximizes value, not merely that (say) to the best of our knowledge at the time of speaking it maximizes value.

Moore appears to be on strong ground in giving this two-part response. For a similar set of worries arises in the case of less controversial definitions, and the solution, in such cases, is the very solution Moore proposes. Consider the definition of 'bachelor' as 'unmarried male.' The worry in this case is that, since we have only limited knowledge of the things that people have done in the past, we can never be certain whether any particular person is an unmarried male. There is, however, still something 'useful' we can do in our search for bachelors in the absence of such certainty, for we can classify people on the basis of the evidence available to us. In compiling a list of bachelors we thus decide whom to

put on the list on the basis of the *likelihood* that they are unmarried males. But if it subsequently comes to light that someone we classified as a bachelor was married at the time that we made our list, then we conclude that it is our list that is in error. To say of someone that he is a 'bachelor,' even when we engage in the humbler task of classifying people on the basis of the evidence available to us, is thus to say of him that he is an unmarried male, not merely that (say) to the best of our knowledge at the time of speaking he is an unmarried male.

For all the appearance that Moore is on strong ground, however, both parts of his response have been challenged. In the following section I will compare Moore's definition of right action, and his resultant account of the 'useful' task to be undertaken in Practical Ethics, with a competing definition of right action and account of the task of Practical Ethics proposed by Frank Jackson. Though Jackson does not explicitly mention Moore, the arguments he gives look like they could well have been formulated with the aim of refuting both parts of the Moorean view.

## 2. Jackson versus Moore

As we have seen, Moore's view has two parts. The first is the claim that right acts are those that maximize value. The second is his account of the 'humbler' task of Practical Ethics, given that we cannot be certain which acts maximize value. We decide what to do by figuring out which acts are *likely* to maximize value. Let's begin by focusing on the second component.

Frank Jackson provides what seems to me to be a quite decisive refutation of Moore's account of the task of Practical Ethics. The refutation takes the form of a counter-example. He asks us to consider Jill, a physician, who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint.

Jill has only two drugs, drug X and drug Y, at her disposal which have any chance of effecting a cure. Drug X has a 90% chance of curing the patient but also has a 10% chance of killing him; drug Y has a 50% chance of curing the patient but has no bad side effects. Jill's choice is between prescribing X or prescribing Y. It is clear that she should prescribe Y, and yet that course of action is not the course of action most likely to have the best results. (Jackson 1991: 467)

What Jackson's counter-example suggests is that, at least when we do Practical Ethics, we have no special interest in which actions are most likely to maximize value. Though Jill's prescribing drug X has a 90 per cent chance of maximizing value—that is, of bringing about a complete cure—and prescribing drug Y

only has a 50 per cent chance of having this result, our intuitive response is that Jill should prescribe drug Y, not drug X. Jill should prescribe drug Y, notwithstanding the fact that it *is not* most likely to maximize value.

Jackson makes explicit his alternative conception of the task of Practical Ethics in his discussion of another variation on the example of Jill, the physician, and her patient, John. This time

Jill... has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way she can tell which of the two is the perfect cure and which the killer drug. What should Jill do?

The possible outcomes we need to consider are: a complete cure for John, a partial cure, and death. It is clear how to rank them: a complete cure is best, followed by a partial cure, and worst is John's death... But how do we move from that ranking to a resolution concerning what Jill ought to do? The obvious answer is to take a leaf out of decision theory's book and take the results of multiplying the value of each possible outcome given that the action is performed, summing these for each action, and then designating the action with the greatest sum as what ought to be done. In our example there will be three sums to consider, namely:

$$\begin{aligned} &Pr(\text{partial cure/drug A taken}) \times V(\text{partial cure}) + Pr(\text{no change/drug A taken}) \times V(\text{no change}) \\ &Pr(\text{complete cure/drug B taken}) \times V(\text{complete cure}) + Pr(\text{death/drug B taken}) \times V(\text{death}) \end{aligned}$$

and

$$Pr(\text{complete cure/drug C taken}) \times V(\text{complete cure}) + Pr(\text{death/drug C taken}) \times V(\text{death})$$

Obviously, in the situation as described, the first will take the highest value, and so we get the answer that Jill should prescribe drug A. (Jackson 1991: 462–3)

What this example shows is that there is a clear alternative to Moore's suggestion that when we engage in Practical Ethics, we should try to figure out which action is most likely to maximize value. For, much as decision theory tells us that the right action to choose is the one that maximizes *expected utility*, so we might suppose that when we engage in Practical Ethics the right action to choose is the one that maximizes *expected value*. This allows us to make good sense of our reaction to the two drugs example. For though prescribing drug X is more likely than prescribing drug Y to maximize value—90 per cent versus 50 per cent—the 10 per cent chance that prescribing drug X will have a very

very bad outcome—it will kill John—as opposed to the 50 per cent chance that prescribing drug Y will merely leave him with his minor but not trivial skin complaint, means that the expected value of prescribing drug X is much lower than the expected value of prescribing drug Y. This well explains why we think that Jill should prescribe drug Y rather than drug X.

So far I have portrayed Jackson as offering an alternative task for Practical Ethics. But in fact Jackson thinks that—with a qualification to be mentioned presently—the two- and three-drugs examples show not just that the second component of Moore's view is mistaken, but that the first component, his definition of right action, is mistaken as well. Here is the relevant passage.

The other possible account of how to recover what a person ought to do from consequentialism's value function that we need to consider holds that a person's beliefs... do not come into the picture. What is crucial is simply which action in fact has, or would have, the best consequences. Many consequentialists write as if this was their view...

There are two problems with this proposal. First, it gives the intuitively wrong answer in the [three] drugs case. In the [three] drugs case, either it is prescribing drug B or it is prescribing drug C which is the course of action which would in fact have the best consequences—and Jill knows this, although she does not know which of the two it is—but neither prescribing drug B nor prescribing drug C is the right course of action for Jill. As we observed earlier, it is prescribing drug A which is the intuitively correct course of action for Jill despite the fact that she *knows* that it will *not* have the best consequences. We would be horrified if she prescribed drug B, and horrified if she prescribed drug C.

The second problem arises from the fact that we are dealing with an ethical theory when we deal with consequentialism, a theory about action, about what to do. In consequence we have to see consequentialism as containing as a constitutive part prescriptions for action. Now, the fact that an action has in fact the best consequences may be a matter which is obscure to the agent. In the drugs example, Jill has some idea but not enough of an idea about which course of action would have the best results... We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics, and having the best consequences is a story from the outside. It is fine for a theory in physics to tell us about its central notions in a way which leaves it obscure how to move from those notions to action, for that passage can be left to something which is not physics; but the passage to action is the very business of ethics. (Jackson 1991: 465–7)

The two related problems with Moore's proposed definition of right action are first, that it gives the wrong answer in the three-drugs example, and second, that since it defines a concept that is not and cannot be action-guiding, it cannot be a definition of our ordinary concept of right action.

However, as subsequently becomes clear, the second objection is much more important than the first—here we come to the qualification I mentioned earlier.

... I need to note an annoying complication. I have been arguing for an interpretation of consequentialism which makes what an agent ought to do the act which has the greatest expected moral utility, and so is a function of the consequentialist value function and the agent's probability function at the time. But an agent's probability function at the time of action may differ from her function at other times, and from the probability function of other persons at the same or other times. What happens if we substitute one of these other functions in place of the agent's probability function at the time of action? The answer is that we get an annoying profusion of 'oughts'...

I think that we have no alternative but to recognize a whole range of oughts—what she ought to do by the lights of her beliefs at the time of action, what she ought to do by the lights of what she later establishes..., what she ought to do by the lights of one or another onlooker who has different information on the subject, and, what is more, what she ought to do by God's lights, that is, by the lights of one who *knows* what will and would happen for each and every course of action. ... I hereby stipulate that what I mean from here-on by 'ought,' and what I meant, and hope and expect you implicitly took me to mean when we were discussing the examples, was the ought most immediately relevant to action, the ought which I urged to be the primary business of ethical theory to deliver. When we act we must perforce use what is available to us at the time, not what may be available to us in the future or what is available to someone else, and least of all what is available to a God-like being who knows everything about what would, will and did happen. (Jackson 1991: 471–2)

Jackson's concession here is both subtle and important.

Though our concept of right action is the concept of an action that maximizes expected value, Jackson suggests that this concept is, as such, incomplete. In order to classify actions we need to know whose expectations are at issue. There are therefore as many concepts of right action as there are creatures with expectations: the agent at the time of acting, an observer, the agent in retrospect, and even, at the limit, God with his perfect knowledge of everything. The latter is, in effect, Moore's view. But though it follows that Moore is right when he says that 'the term duty is... so used that, if we are subsequently persuaded that any possible action would have produced more good than the one we adopted, we admit that we failed to do our duty'—there is indeed a concept of right action relative to which this is so—we must not conclude, on this basis, that our primary or most central concept of right action is the concept of an action that (say) maximizes value.

In support of Jackson's 'profusion of "oughts"', it is perhaps worth emphasizing that, in the sorts of situations Moore imagines in which we evaluate

actions with hindsight, expectations still matter. An action performed by an agent at  $t_1$ , but evaluated subsequently at  $t_2$ , is still appropriately judged to be right just in case that action maximizes *expected* value, it is just that the expectations in question are those the evaluator has subsequently at  $t_2$ . Imagine the original three-drugs example evaluated in retrospect when (say) it is believed that the original assignment of probabilities was mistaken. Instead of drug A being very likely to relieve John's skin condition but not completely cure it, and one of drugs B and C being almost certain to effect a complete cure, while the other will kill the patient, the new view is that it is drug B that is very likely to relieve the skin condition, but not completely cure it, and one of drugs A and C which is almost certain to completely cure the condition, while the other will kill the patient, with there being no way to tell which will do what. In this case the subsequent evaluator should surely suppose that the right thing for Jill to have done was to give John drug B, not drug A. In other words, by the subsequent evaluator's lights, it is still virtually certain that the right thing for Jill to do is not what will maximize value. It is just that it is her subsequent expectations, not Jill's original expectations, that are crucial for evaluating rightness.

But notwithstanding the 'profusion of "oughts"', Jackson argues that one 'ought' stands out from the others as the 'ought' that is 'most immediately relevant to action.' What I take it he means by this is that only one of the 'oughts' grounds genuine *criticism* of the agent. Only one links up, in the right kind of way, with a story about what the agent can appropriately be held responsible for doing, and this should therefore be identified as our primary concept of right action. Imagine again the revised three-drugs example. The subsequent evaluator might well think that Jill ought to have given John drug B, not drug A, as this is what maximizes expected value where the expectations in question are the subsequent evaluator's own. But if the information about the drugs wasn't available to Jill, then there is no sense in which she failed to live up to her responsibilities as an agent in giving him drug A. Indeed, if Jill had given John drug B, then there is surely a sense in which the subsequent evaluator would still be totally horrified. Jill would have done what the evaluator thinks is the right thing to do, but her right conduct could at best have been a complete fluke, relative to her own reasons for acting. But if this is right—if Jill can only be held responsible for doing the best she can, given the information available to her—then, Jackson argues, the 'ought' that it is 'the primary business of ethical theory to deliver' is the 'ought' defined in terms of maximization of expected value where the expectations are those of the agent at the time of acting.

Let me sum up. Contrary to Moore, Jackson argues that the task of Practical Ethics is to establish which acts maximize expected value, not which acts are most likely to maximize value. Moreover and much more importantly, again

contrary to Moore, Jackson argues that though there are alternative concepts of right action—there are as many concepts of right action as there are expectations relative to which we could assess the maximization of expected value—right action in the sense of action which maximizes expected value, where the expectations are the agent's own at the time of acting, is the primary ethical concept, the one that hooks up in the right kind of way with what we can hold an agent responsible for doing. Moore is thus doubly wrong when he says 'right' does and can mean nothing but "cause of a good result." 'Right' can mean 'maximizes expected value' and this meaning is primary.

### 3. The definition of 'right'

Moore and Jackson give us competing definitions of right. But who is right and who is wrong?

According to Jackson, the primary meaning of 'right' is the meaning that is 'most immediately relevant to action.' As I have said, I take it that his idea is that right acts, in this primary sense, are those we can legitimately expect agents to do and criticize them for failing to do. As Jackson sees things it follows that the primary definition of right action must be given in terms of the maximization of expected value, not, as Moore thinks, in terms of the maximization of value *tout court*, and it must be given in terms of the maximization of expected value where the expectations are the agent's at the time of action, not those of the agent at some other time or someone else. Unfortunately, however, there is a gap between Jackson's premise and conclusion. We can no more legitimately expect agents to maximize expected value than we can expect them to maximize value.

A striking feature of Jackson's definition of 'right' is an asymmetry in the treatment of evaluative and non-evaluative facts as determinants of right action. Suppose I observe an agent behaving and wonder whether she is failing to live up to her responsibilities as an agent, so leaving herself open to criticism. What exactly does this involve? Jackson rightly points out that, on the non-evaluative facts side of things, all we can reasonably expect is that an agent does the very best she can, given the information available to her: the full exercise of such rational capacities as she has. Putting to one side cases of impaired rationality, this means that she must form her beliefs about the means to her desired ends in a responsible manner, given the evidence available to her, and she must subsequently act on such beliefs in an appropriate manner as well. When she believes that there are alternative ways of realizing her desired ends, she must prefer the more certain option, and when she is equally certain that

acting in different ways will realize different desired ends, she must prefer the option that realizes the desired end that she desires more. But since no mere exercise of such rational capacities as an agent has will ensure that the non-evaluative facts manifest themselves to her, it follows that the mere fact that she lacks knowledge does not, as such, render her liable to criticism.

Yet while this well explains why Jackson insists that right action is action that maximizes *expected* value, it doesn't explain why he thinks that right action maximizes expected *value*. Indeed, anyone impressed by Jackson's argument on the non-evaluative facts side of things should surely suppose that an equally impressive argument could be made for the conclusion that right action consists not in the maximization of expected *value*, but rather in the maximization of expected *value-as-the-agent-sees-things*. For no mere exercise of such capacities as an agent has looks like it will ensure that what is really valuable will manifest itself to her either.

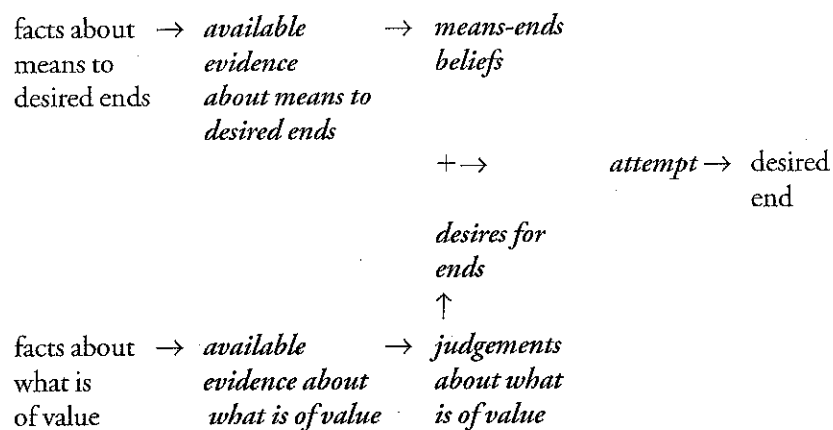
There are, after all, cultural circumstances in which it would be wildly optimistic to suppose that agents could, merely through the exercise of their own rational capacities, come to judge to be valuable what's really valuable. Cultures that are dominated by oppressive religions, and the like, would seem to make it very difficult, perhaps even impossible, for ordinary people even to contemplate alternatives to what is portrayed within their culture to be valuable. More generally, to the extent that we think that our own values are an improvement on those of our parents, the wrong thing for us to think is that we have been more responsible than our parents in the formation of our evaluative beliefs. Indeed, this would be a remarkably arrogant thing to think. The right thing to think is rather that the evidence available to us, as opposed to our parents, is different, and, as a result, we are simply better placed to form evaluative beliefs.

If this is right, however, then it seems that the most that we could ever expect of a normal agent—that is to say, again putting to one side cases of impaired rationality—is that they form their evaluative commitments in a way that is sensitive to such evidence as is available to them and that they form their desires in a way that is sensitive to their evaluative commitments. Probabilities or confidence levels come into the picture twice over. They come in once because we have different levels of confidence that various means are means to our desired ends. And then they come in again because we have different levels of confidence about what we judge to be intrinsically valuable. When an agent is equally confident about the intrinsic value of two things, but she judges one more valuable than the other, then she must desire more strongly the one that she judges more valuable, and when she judges two things to be equally intrinsically valuable, but is more confident of the value of the one than the other, then



she must desire more that about which she is more confident. The upshot is that whether we would criticize an agent for failing to maximize expected *value* turns very much on that agent's epistemic circumstances: whether or not they had available to them evidence of what is really valuable.

The situation can be diagrammed as follows (Smith 2004).



What we hold agents responsible for is, at most, the elements in bold italics, for these are the elements whose relations are governed by the agent's exercise of such rational capacities as she has. Has the agent attended to all of the available evidence as regards both what is of intrinsic value and means to her desired ends in the formation of her evaluative judgements and means-end beliefs? Do the agent's desires for ends reflect her judgements about what is intrinsically valuable? Is what the agent attempts to do an appropriate reflection of his desires for ends and means-end beliefs? But if this is right then, contrary to Jackson, we simply do not hold agents responsible for failing to maximize expected *value*. At most we hold them responsible for failing to maximize expected *value-as-they-see-things*. Maximizing expected value-as-they-see-things is the right thing for them to do, in the sense that is 'most immediately relevant to action.' This is what we can legitimately expect them to do and criticize them for failing to do.

Once this becomes clear it seems to me that we see much better the attractions of Moore's definition of right action. Unlike Jackson, Moore adopts a completely symmetrical approach to both evaluative and non-evaluative facts as determinants of right action. Even though we can only hold agents responsible for failing to maximize expected value-as-they-see-things, that just goes to show that our concept of a right action isn't tied so closely as Jackson thinks to what we can legitimately hold agents responsible for doing. If people live in an oppressive culture in which they are epistemically cut off from what's really

valuable then the very best that they can do is to maximize expected value-as-they-see-things. But since, in so doing, they fail to maximize expected value, even Jackson must conclude that they fail to do the right thing, albeit through no fault of their own. The concept of a right action, in the sense relevant to ethics, must therefore be the concept of a certain sort of ideal, one that abstracts away from the possibility of this kind of error for which we cannot hold an agent responsible.

But once we have distinguished the concept of right action from the concept of what we can legitimately hold agents responsible for doing in this way, there would seem to be no stable stopping point short of the Moorean definition. People who maximize *expected* value, but fail to maximize value *tout court*, also fail to act in accordance with an ideal. Moreover the explanation is much the same as the explanation in the case where people maximize expected value-as-they-see-things, but fail to maximize expected value. In each case they are epistemically cut from a relevant domain of facts, in this case, facts about what the means to their desired ends really are. It is therefore this latter concept, the concept of an action that maximizes value *tout court*, a concept that incorporates two ideals—one on the evaluative side of things and the other on the non-evaluative side of things—that properly captures the concept of right action. We should therefore accept Moore's definition of right action, not Jackson's.

#### 4. The task of Practical Ethics

So far, so good; but what about the task of Practical Ethics? How do we get from abstract knowledge of what it is about an action that makes it right—that it maximizes value—to a decision about what to do in a concrete situation? Isn't Jackson right that the Moorean view will either force us to conceive of the task of Practical Ethics in the completely wrong way, ensuring that we get the wrong answer to which action we should choose—this is what Moore himself does—or else make that transition altogether opaque? The answer is that it doesn't, and that it is instructive to see why not.

The fundamental problem with Moore's conception of the task of Practical Ethics is not his conception of right action, but rather his conception of moral motivation. Suppose that right actions are those that maximize value. Does this imply that, as a right-minded agent, I will be moved, at bottom, by an intrinsic desire to maximize value? Though Moore is not explicit on the issue, this does seem to be his view, for having this intrinsic desire is precisely what's required for agents to take an interest in which action is most likely to maximize value when they decide what to do. Imagine again the two-drugs example.

If what Jill cares about, fundamentally, is that she maximizes value, then of course she will prescribe drug X rather than drug Y. For prescribing drug X has a much better chance of achieving what really matters to her. In so far as we have the reaction that we do to the two-drugs example, we implicitly assume that this is not what we fundamentally care about. But in that case what do we fundamentally care about? And how does what we fundamentally care about lead us to make a more sensible decision, a decision that squares with our account of right action?

Consider again the relations charted in the diagram above. If I believe that experiencing pleasure is intrinsically valuable then, insofar as my desires are formed correctly in response to my evaluative beliefs, I will have an intrinsic desire that I experience pleasure. And if I believe that being autonomous is intrinsically valuable, then, insofar as my desires are formed correctly in response to my evaluative beliefs, I will have an intrinsic desire that I be autonomous. Being right-minded thus requires not that I have an intrinsic desire to maximize value but rather, since it requires that my desires are appropriately sensitive to my evaluative judgements, that I have intrinsic desires for the things I judge to be valuable themselves. This suggests a completely straightforward explanation of why, if I am right-minded, I will decide to perform the action that maximizes expected value-as-I-see-things.

Suppose, to begin, that I believe that experiencing pleasure is more valuable than being autonomous. If we abstract away from the levels of confidence associated with each of these beliefs—let's assume I am equally confident—then, insofar as my desires are formed correctly in response to my evaluative beliefs, my intrinsic desire that I experience pleasure will be stronger than my intrinsic desire that I be autonomous. And if we suppose that I am more confident of the value of pleasure than the value of being autonomous, but we abstract away from the degree of value I assign to each—let's assume that I assign them equal value—then, insofar as my desires are formed correctly in response to my evaluative beliefs, my intrinsic desire that I experience pleasure will again be stronger than my intrinsic desire that I be autonomous.

Putting these two conclusions together, suppose that I am very confident that being autonomous is valuable, but not very valuable, and that I am fairly confident—not as confident, but still quite confident—that experiencing pleasure is valuable too, but more valuable than being autonomous. Insofar as my desires are formed correctly in response to my evaluative beliefs the relative strengths of my intrinsic desires to experience pleasure and be autonomous will then depend entirely on the levels of confidence and associated degrees of value. The strength of my intrinsic desires will track the product of the levels of confidence and associated degrees of value: the greater that product, the greater

the strength of my intrinsic desire. My intrinsic desire to be autonomous might therefore even be stronger than my intrinsic desire that I experience pleasure if the difference in the relative levels of confidence is greater than the difference in the associated degrees of value.

Finally, let's suppose that we plug these intrinsic desires into a standard expected utility calculation, a calculation that takes account of the different levels of confidence I have about the various means to my intrinsically desired ends. It then emerges that, in so far as I am right-minded, I will decide and choose to act not just so as to maximize expected utility, but so as to maximize expected value-as-I-see-things.

This, it seems to me, is the right thing to say about the task of Practical Ethics. We must decide to do what maximizes expected value-as-we-see-things, and, when we offer advice to others, we must in effect tell them what we would decide to do if we found ourselves in their circumstances. The ambiguity of what it means to find ourselves in another's circumstances—whether we imagine ourselves with our own values and the other person's expectations, or both our own values and our own expectations—supports and explains a profusion of 'oughts' of the kind noticed by Jackson.

Importantly, however, note that even if I manage always to act so as to maximize expected value-as-I-see-things, and so act responsibly and avoid all criticism, and even if others manage always to do exactly what I would do if I found myself in their circumstances, we can still quite happily admit that such actions may yet fail to meet an ideal. For it is a simple fact of life that we act in an environment in which we are epistemically cut off from all sorts of facts about both what is really valuable and what the means to ends really are. We can therefore quite happily admit that the right action for anyone to perform in their circumstances, the one against which their own perfectly responsible behaviour in those circumstances may well and quite appropriately be judged to be a failure, is the act that maximizes value.

## Conclusion

At one point Jackson says, in defence of his own definition of right action:

We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics, and having the best consequences is a story from the outside. It is fine for a theory in physics to tell us about its central notions in a way which leaves it obscure how to move from those notions to action, for that passage can be left to something which is not physics; but the passage to action is the very business of ethics. (Jackson 1991: 467)



One way of putting the argument of this paper is that Jackson gets things exactly the wrong way around. We already have a story from the inside of an agent, a story that is not ethics. This is the story of rational decision-making as outlined in the diagram above, the story that is formalized—or anyway partially formalized—in decision theory. The passage to action is thus not the business of ethics itself; or, at any rate, it is not the business of that part of ethics itself whose concern is to provide a definition of right action. The business of that part of ethics is to come up with a definition of right action which dovetails in the right kind of way with the theory of rational decision-making. My suggestion is that that definition will be an idealization of the story of rational decision-making and that this is what Moore's definition of right action provides.

## References

- DREIER, JAMES (1993). 'Structures of Normative Theories', *The Monist*, 76: 22–40.
- JACKSON, FRANK (1991). 'Decision Theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics*, 101: 461–82.
- LOUISE, JENNIE (2004). 'Relativity of Value and the Consequentialist Umbrella', *Philosophical Quarterly*, 54: 518–36.
- MOORE, G. E. (1903). *Principia Ethica* (Cambridge: Cambridge University Press).
- (1942). 'A Reply to my Critics', in P. A. Schilpp (ed.), *The Philosophy of G. E. Moore* (Evanston, Ill.: Northwestern University): 543–667.
- RAWLS, JOHN A. (1971). *Theory of Justice* (Cambridge, Mass.: Harvard University Press).
- RUSSELL, B. (1910). 'The Elements in Ethics', repr. in *Philosophical Essays* (London: Routledge, 1944): 13–59.
- SCANLON, T. M. (1998). *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press).
- SMITH, MICHAEL (2003). 'Neutral and Relative Value after Moore', *Ethics*, 113: 576–98.
- (2004). 'The Structure of Orthonomy', in J. Hyman and H. Steward (eds.), *Action and Agency*, Royal Institute of Philosophy Supplement, 55 (Cambridge: Cambridge University Press): 165–93.

## 7

# Scanlon versus Moore on Goodness

Philip Stratton-Lake and Brad Hooker

G. E. Moore's *Principia Ethica* started twentieth-century moral philosophy with the open question argument. At the end of the century, T. M. Scanlon, in his *What We Owe to Each Other*, says that he is led by Moore's open question argument, not to Moore's own view of the relation of goodness to reasons, but to the buck-passing account of goodness. The buck-passing account of goodness is the view that goodness is not a property that itself provides practical reasons (i.e. reasons to desire, to admire, to pursue, etc.) but rather is the purely formal (higher-order) property of having some other properties that provide reasons. According to the buck-passing account, the power to provide practical reasons is passed from goodness itself to the properties on which goodness is based. In this paper we will explore the question of whether there is a good argument from the open question argument to the buck-passing account of goodness. We will argue that the case for the buck-passing account of goodness is stronger than some of Scanlon's critics make out. We will not, however, offer a full defence of the buck-passing account of goodness. Such a defence would have to respond to the various objections to the buck-passing view itself, and that is beyond the scope of this paper.

## 1. The open question argument

Moore presents the open question argument as an argument for the view that naturalists commit a fallacy. Unfortunately, Moore was rather vague about what this fallacy is supposed to be (see Frankena 1939). Moore most often portrays the naturalistic fallacy as the mistake of thinking that 'good' is definable. But what is most important to Moore (1993: 19) is to insist not that 'good' cannot be defined, but that it cannot be defined in *naturalistic* terms, and we