

Agency, Character and The *Real* Failure of Consequentialism

Kevin C. Klement
University of Iowa

1. Introduction

Consequentialism continues to make ethicists nervous. There is widespread agreement even among its critics that consequentialism has come to play a dominant role in ethical theory. Ironically, the more this is perceived to be the case, the greater the anxiety becomes to “escape” from its “spell”.¹ Proposed escape routes come in many varieties. However, there is a common core to the main arguments against consequentialism: the charge that consequentialism either demands too much of people or demands the wrong things of them. This charge is motivated by consideration of cases involving (what are usually taken to be) agent-relative reasons, including, on one hand, the pursuit of one’s own projects, interests, commitments and personal relationships, and on the other, rights and deontological constraints. These cases are understood in very different ways by the different critics, but the conclusion is the same: consequentialism must be wrong because it cannot accommodate our strong moral intuitions about these sorts of cases.

Consequentialism is indeed a powerful moral theory. I find nothing surprising nor scandalous about its dominant role in contemporary ethics. However, I agree in the end that it must be rejected, though for very different reasons. Nevertheless, the reasons for my rejection can be appreciated by considering a consequentialist reply to these criticisms. The core of the reply is that consequentialism can accommodate the intuitions lying behind the troubling cases by suggesting that while we can certainly imagine cases in which acting on one’s projects or on constraints will fail to maximize value, it also maximizes value for individuals to develop dispositions and habits to act with these things as motivations, and so people who act on such dispositions even in prob-

lematic cases cannot be held to be blameworthy (even if, strictly speaking, their actions are wrong). This is an extremely important reply. Indeed, within it, we find one of the most important realizations to be included in any plausible ethical theory: that moral evaluation cannot be limited to the examination of a particular act isolated from considerations of the dispositions, habits, patterns of behavior and causal history leading to it. But I will also argue that if we take this insight seriously, when combined with certain plausible assumptions regarding the causes of human action and behavior and the nature of agency—and in particular, what I call *psychological determinism*—what we have is no longer a defense of consequentialism against its more deontological critics. Rather, we have argument for the *incoherence* of consequentialism itself. But this should not be taken as vindication of the critics of consequentialism. Both camps are guilty of a shared misunderstanding.

This misunderstanding infects the way in which the very debate is set up. Both sides attempt to explain the rightness, wrongness, permissibility and impermissibility of individual actions. But if the realization above is correct, this is not how ethical theory should be organized. Rather than studying and evaluating individual acts, we need rather to study causal relationships between whole patterns of behavior. My radical suggestion is that ethics should not primarily concern itself with the notions of *right* and *wrong*, and insofar as consequentialism is one theory with this concern, it must be replaced. While I cannot fully elaborate on the sort of alternative moral theorizing I have in mind, it would not likely coincide with the largely agent-relative theories offered by so many contemporary critics of consequentialism. These theories are as much, and probably *more*, guilty of the same misplaced concern.

2. Options and Constraints: The Attack on Consequentialism

Consequentialism—limited in this debate to act-consequentialism—is here understood as the theory that an act is right if and only if there are no alternatives to that act that would have better overall consequences. Consequentialism is also standardly wedded to a purely agent-neutral conception of “better

consequences", i.e. the goodness or badness of the consequences of an act is not relative to the standpoint of the evaluator. Critics charge that this theory cannot accommodate some of our most basic moral intuitions, and they usually trace this failure back to a poor understanding of moral agency and the full range of morally relevant reasons we have for acting. In particular, consequentialists cannot make room in their theory for *options* or "agent-centered permissions" and *constraints* or "agent-centered restrictions".

To understand agent-centered permissions, consider the following hypothetical scenario. It has been my lifelong project to become a professional academic in philosophy. I have earned my Ph.D. and am now faced with searching for a job in a tight job market, one in which I have very little hope of receiving a position. Suppose I do manage to become a finalist for a certain position, although I am relatively certain that if I do not receive this appointment, another offer will not be forthcoming. The decision has come down to two candidates. My rival, who has the same lifelong project as I, has rather similar philosophical credentials and either of us would do an equally good job at the position. One of us will be appointed, and the other will be unable to fulfill his or her project. Now suppose also that I have certain computer related skills that my opponent lacks, skills having no bearing on my ability to fulfill my potential role as an academic philosopher, but which could secure me a relatively stable and happy future at a different job. My rival has no such skills, and indeed, if she or he were to fail to obtain an academic job, while certainly she or he would find work, it would not be as secure or enjoyable as the sort I could receive. Now, the question becomes whether I should morally withdraw myself from consideration. It seems that, impersonally considered, consequences will be better if it is I who ends up working outside academia. According to consequentialism, not only is it right, but I am *morally obligated* to withdraw myself from consideration.

Critics of consequentialism point out that this seems wildly at odds with our ordinary moral intuitions. While it may seem permissible, perhaps even morally praiseworthy, for me to withdraw, I am not obligated to, nor would I be doing anything wrong if I chose to continue my candidacy. Consequentialism is thus missing

something, but how is this to be explained? According to Samuel Scheffler, we must recognize that, while we might be *able* to adopt an impersonal perspective for evaluating the aggregate good, we also have a quite *independent* personal point of view in which our own projects and commitments loom larger than those of others. This personal point of view is an important feature about us as moral agents, and it is rational for moral theory to do justice to the nature of persons by somehow taking this into account. Scheffler argues that one reasonable way to do this would be to include agent-centered permissions, or allowances to agents to weigh their own concerns and projects somewhat more in moral deliberation. Thus, they would be given a limited protective sphere in which to pursue their own projects even if this leads to failure to maximize the overall good.² Other theorists tend to put the point in a stronger way. Bernard Williams argues that consequentialism forces agents to see their own projects as dispensable, and therefore represents an attack on the *integrity* of individuals, and can only lead to alienation.³ The core of the argument, however, is the same. Without the inclusion of *options* to pursue one's own interests and projects out of strict proportion to how doing so maximizes the good, consequentialism ignores that people have their own subjective standpoints, with their own concerns, projects and commitments, and no theory that ignores this very important feature of moral agency can be adequate.⁴

Often included under the general rubric of options or agent-centered permissions, but sometimes treated separately, are those "special obligations" and commitments stemming from our personal relationships with family members, friends and loved ones. Standard consequentialism demands that we weigh the good of all people equally, and that just as we are not permitted to put our own interests, projects and livelihood ahead of that of others, we are not permitted to give extra weight to the interests, projects and livelihood of our friends and relatives. This, too, seems at odd with our normal moral intuitions. Diane Jeske and Richard Fumerton ask us to consider a case in which we must decide between saving the life of one's own child or the lives of two unknown children.⁵ They suggest that it seems as if one should have the option (and perhaps even the duty) to exhibit greater concern for one's own child in

such a case. But this poses a problem for consequentialism, with its purely nonrelativistic conception of value. Consequentialism demands that we treat everyone impartially, but it is a morally significant feature of persons that we do develop relationships that seem to *require* us to treat individuals partially, giving extra weight in our moral decision making to individuals because of the relationships they bear to us. Here, too, consequentialism fails to take adequate account of the nature of persons as moral agents.

Deontological constraints are more controversial. Even the critics of consequentialism cannot agree on whether they exist and can be given the same defense as options.⁶ There are at least some theorists, however, who argue that they do exist and that their existence reveals a similar flaw in consequentialist thought. Still, much of the literature deals not with unambiguous examples in which we *clearly* must not engage in some morally objectionable but optimific acts, but rather with cases for which the answer provided by the consequentialist is not clearly the right act. These examples are meant to prove that there must be *something* more to the moral assessment of acts than consideration of their consequences, because otherwise these situations would not be the difficult puzzles that they obviously are. Thomas Nagel provides us with this example:

You have an auto accident one winter night on a lonely road. The other passengers are badly injured, the car is out of commission, and the road is deserted, so you run along it till you find an isolated house. The house turns out to be occupied by an older woman who is looking after her small grandchild. There is no phone, but there is a car in the garage, and you ask desperately to borrow it, and explain the situation. She doesn't believe you. Terrified by your desperation she runs upstairs and locks herself in the bathroom, leaving you alone with the child. You pound ineffectively on the door and search without success for the car keys. Then it occurs to you that she might be persuaded to tell you where they are if you were to twist the child's arm outside the bathroom door. Should you do it?⁷

The consequentialist answer to this sort of answer is clear: the badness of the child's arm being twisted momentarily is far outweighed by the need to get your injured friends to safety. But common sense morality seems to recognize that certain types of acts are objectionable even if their consequences are for the overall good. Many would conclude that twisting the child's arm is such a restricted act. There are limits to what one may do in pursuit of good consequences. Even if we do conclude of this particular case that it is permissible to twist the child's arm, certainly the decision is not as easy as the consequentialist makes it out to be. Especially if we are to imagine that the relative advantages in terms of overall consequences are small, the permissibility of acts involving lying, breaking promises, betrayal, torture, intentional injury or harm can be called into question. A moral theory that claims that there are cases in which we are morally barred from performing some act that maximizes the good is one that incorporates *constraints*, or agent-centered *restrictions*.

How are such restrictions to be justified? Nagel suggests that the most promising avenue is to also understand them as agent-relative reasons, or reasons stemming from our subjective standpoints as moral agents. I can recognize, from an impersonal perspective, that twisting the child's arm (or killing one innocent to save the lives of five others, etc.) will lead to better consequences. But I also recognize that it requires me to *knowingly* commit an evil or objectionable act as a means towards some good end. But as a moral agent, I have a personal reason to object to allowing my actions to be directed by evil (even if I have the best overall consequences in mind.) Williams suggests something similar: as agents, we care not only about what states of affairs come about, but about what we directly *do* as opposed to what we only *allow* to happen. Insofar as consequentialism incorporates a purely agent-neutral conception of the good, and a theory of right based entirely on the consequences of the act in question, consequentialism cannot accommodate these sorts of considerations.

So far, I have only been engaged in describing uncritically how it is that many critics of consequentialism argue by invoking options, constraints and special obligations. Criticism will come later. But it is worth here summarizing some of the major themes of this

approach and bringing to light some of the more implicit assumptions about how the task of ethical theory is understood within this tradition. In each case, we are asked to imagine a certain case in which a person has to make a moral decision. It is then argued in each case that consequentialism fails to make sense of the perspectives or situations of those moral agents who are forced to make such decisions. It seems that underlying this strategy is the idea that ethical theory is theory about how individuals on individual occasions should go about *reasoning* about what acts in that situation are right, wrong, permissible or obligatory *in that situation*. The focus is thus on the nature of moral agency, and it is presupposed that what is important about us as moral agents is that we are sometimes faced with moral quandaries in which we need to decide between different alternatives, and that the task of ethical theory is to provide us with guidelines about what sort of reasoning is appropriate in these situations. This will be important later.

3. Dispositions and Character: A Consequentialist Reply

I now turn to what I take to be the most important consequentialist reply to the objection that consequentialism is flawed insofar as it cannot accommodate options, constraints and the value of preferential treatment of friends and loved ones. The reply has perhaps been given fullest attention in the work of Peter Railton,⁸ although I shall be taking it somewhat beyond what he explicitly has to say.

The reply begins by noting that consequentialism typically divorces the theory concerning under what conditions an act is right or wrong from assessment of the motives leading to an act or of the agent herself. There is nothing inherent in consequentialism to the effect that a person must act with maximizing the good *in mind* in order to be acting rightly. Thus, it need not be a dictum of consequentialism that people ought to explicitly conduct their moral deliberation according to consequentialist principles. Certainly, some important historical consequentialists, Sidgwick included, hoped that people would conduct their moral reasoning by attempting to determine what act would maximize the good and then so act. Railton labels the sort of consequentialism that includes the stronger thesis that people ought always to try to determine what

act will lead to the best consequences whenever faced with a moral decision *subjective consequentialism*. But Railton distinguishes this from consequentialism understood only as a theory about under what conditions an act is right or wrong. *Objective consequentialism* holds only that whether an act maximizes the good is the criterion making it right or wrong. One need not accept the former in order to accept the latter. The objective consequentialist can hold that there are cases in which a person ought *not* to try to govern her deliberations according to consequentialist principles. In many cases, the impersonal good will be better served if people act without performing anything resembling a calculation of the potential good and bad to be had by the alternative choices. To complete Railton's vocabulary, a *sophisticated consequentialist* is a person "who has a standing commitment to leading an objectively consequentialist life",⁹ but who acknowledges that acting as a subjective consequentialist is not necessarily the best, only, or most appropriate way to accomplish this.

One should already see here a deviation in how the sophisticated consequentialist understands moral theorizing from how others, particularly the sorts of deontological critics referred to in the previous section, understand it. Here, consequentialism is understood not as a theory about how people ought to go about moral deliberation, but rather only as an explication of the truth conditions regarding what makes actions right or wrong. This may seem paradoxical at first, but is not. One could make a parallel for the rules of logic. Logic tells us what inferences are valid and what are invalid, but logic need not thereby be understood as dictating to us how we ought rationally to go about making inferences. In fact, the sorts of failsafe procedures favored by the logician (including exponentially complex procedures such as truth tables) might for various reasons *not* be the sorts of strategies people should rationally adopt as inference-makers. In fact, we might be prone to make more errors and manage our cognitive feats much less well if we tried to always infer according to the complex rules of logic. Morality may turn out to be similar. Ethical theory can tell us what makes an act right or wrong, but we need not infer from this that we ought to use the teachings of ethical theory in our moral reasoning. In fact, we might end up acting rightly less of the time if we did so.

What does this have to do with responding to the objections based on options and constraints? The sophisticated consequentialist still holds that an act is right only if it maximizes the impersonal good, and so, strictly speaking, cannot hold that it is actually permissible for an agent to disproportionately favor his or her interests, projects or those of families and friends. But the sophisticated consequentialist might hold that people ought to develop dispositions to act and to deliberate in which they favor their own interests or treat people impartially based on their personal relationships, not because it *actually is* permissible for them to act in this way when they have alternatives, but because acting with these motivations and having these sorts of relationships brings about better overall consequences than attempting to live as subjective consequentialists. Railton gives the following example of a person who acts based upon a personal commitment to another person:

Consider, then, Juan, who . . . has always seemed a model husband. When a friend remarks on the extraordinary concern he shows for his wife, Juan characteristically responds: "I love Linda. I even *like* her. So it means a lot to me to do things for her. After all we've been through, it's almost a part of me to do it." But his friend knows that Juan is a principled individual, and asks Juan how his marriage fits into that larger scheme. After all, he asks, it's fine for Juan and his wife to have such a close relationship, but what about all the other, needier people Juan could help if he broadened his horizon still further? Juan replies, "Look it's a better world when people can have a relationship like ours—and nobody could if everyone were always asking themselves who's got the most need. It's not easy to make things work in this world, and one of the best things that happens to people is to have a close relationship like ours. You'd make things worse in a hurry if you broke up those close relationships for the sake of some higher goal. . . ."¹⁰

Juan is a sophisticated consequentialist. He has a standing commitment to maximizing the amount of good in the world, but doesn't believe that this entails that he guide his every action by a consequentialist decision procedure. Rather, he thinks that things will be better off overall if he develops relationships and patterns of behavior that directly cause him to sometimes, perhaps even knowingly, perform acts which do not seem to be in the interest of the impersonal good. Railton goes on to give the following example of how this might play out in a specific case:

Consider again Juan and Linda, whom we imagine to have a commuting marriage. They normally get together only every other week, but one week she seems a bit depressed and harried, and so he decided to take an extra trip to see her. If he did not travel, he would save a fairly large sum that he could send Oxfam to dig a well in a drought-stricken village. Even reckoning in Linda's uninterrupted malaise, Juan's guilt, and any ill effects on their relationship, it may be that for Juan to contribute the fare to Oxfam would produce better consequences overall than the unscheduled trip. Let us suppose that Juan knows this, and that he could stay home and write the cheque if he tried. Still, given Juan's character, he in fact will not try to perform this more beneficial act but will travel to see Linda instead. The objective act-consequentialist will say that Juan performed the wrong act on this occasion. Yet he may also say that Juan had had a character that would have led him to perform the better act (or made him more inclined to do so), he would have had to have been less devoted to Linda. Given the ways Juan can affect the world, it may be that if he were less devoted to Linda his overall contribution to human well-being would be less in the end, perhaps he would become more cynical and self-centred. Thus it may be that Juan should have (should develop,

encourage and so on) a character such that he sometimes knowingly and deliberately acts contrary to his objective consequentialist duty. Any other character, of those actually available to him, would lead him to depart still further from an objectively consequentialist life. The issue is... whether he would in fact decide to stay home if he had that character, of those available, that would lead him to perform the most beneficial overall sequence of acts.¹¹

I will later argue that there are serious problems with Railton's account here, but for the moment it should be clear at least how this is meant as a way of capturing the intuitions many critics of consequentialism have with regard to options and special commitments to friends and family members. Railton's strategy is not to admit that that it is ever *right* for a person to act so as to disproportionately favor a loved one, or one's own interests, but rather that what *is* right is to develop dispositions, habits and the sort of character that in general lead to the best sorts of consequences but, unfortunately, also makes it necessary that a person will on occasion commit acts which—considered in isolation—do not maximize the good. So in this case, even under consequentialist guidelines, it is right for Juan to develop the devotedness he does to Linda, even if it leads him away from acting as a subjective consequentialist. Here we can see how projects can be incorporated within a consequentialist framework, not with the suggestion that it is right for us to disproportionately weigh our own interests and commitments *in all cases*, but that it is right for us to develop dispositions and characters that will lead us to do so wrongly perhaps in some cases, but in general are for the overall good.

Railton doesn't explicitly discuss constraints, but it is clear that a similar argument could be given for them. Again, even the sophisticated consequentialist is committed to the idea that one ought always maximize the good, and so cannot hold that there are times in which we are not even permitted to do so, as constraints entail. Nevertheless, the sophisticated consequentialist might hold that it is right for us to develop certain dispositions to avoid or

refrain from participating in certain sorts of acts even when one has good reason to think that not refraining would have better consequences. Consider lying. Certainly there are cases in which telling a lie would lead to better consequences than not lying, and this excludes the possibility of a general constraint against lying in consequentialist theory. However, this does not mean that the sophisticated consequentialist need hold that in each case when confronted with a decision of whether or not to lie, we ought to perform a calculation of the goodness and badness of the effects and act accordingly. The sophisticated consequentialist might hold instead that we ought to develop and encourage within ourselves a general disposition to be honest, even in conditions in which it appears that lying would have better consequences. Having such a disposition, although it might lead us to be honest in a case in which lying would actually have been our consequentialist duty, might *overall* lead to better consequences. If we instead took the subjective consequentialist approach, we might be prone to lie in cases in which it looked as if lying would be appropriate, but in fact only made things worse. Many of us know from our own experience situations in which we thought we could make things turn out all right with a lie, but had it backfire on us when the lie ended up leading to more lies and eventually to a much worse situation. ("Oh the webs we weave. . .") Rather than approaching such quandaries from the standpoint of the subjective consequentialist, sophisticated consequentialism might hold it to be better to instill a general penchant for honesty in people, albeit one which will lead us to do the wrong thing on occasion. Similarly, we can imagine that it is best to encourage in people a strong disposition to avoid doing such things as twisting children's arms, even to the extent that a person might not in fact do her duty and twist the child's arm when faced with the situation sketched above. I myself have argued in a previous paper precisely for the rationality of developing dispositions against inflicting intentional harm on others for roughly these reasons.¹² Thus, while it is true that consequentialism cannot include constraints as the deontologist understands them, it can explain why cases such as twisting the arm or killing one innocent to save multiple others are not as simple as the subjective consequentialist would have them.

There are interesting questions to be answered here in terms of exactly what sort of dispositions someone like Railton or like-minded thinkers (including myself, in this instance) have in mind. Do we really want people to develop dispositions so strong such that one would be unable to lie in any conceivable circumstances? This is parallel to a problem facing deontologists. Many contemporary deontologists are dissatisfied with "absolutist" constraints, which would hold that it is always wrong to lie or always wrong to use physical force on an "innocent" child. If the deontologist denies that constraints hold unconditionally, he or she must give some principled way to discriminate between situations in which a constraint applies and those in which they do not (without merely resorting to an appeal to consequences.) However, the sophisticated consequentialist actually has the upper hand when it comes to offering a solution to this sort of problem, because the sophisticated consequentialist can hold that such questions are empirical ones. They have to do with human psychology and the overall consequences of different types of dispositions. Only psychological research can tell us how behavioral dispositions are formed, what sorts are possible, how strong they can be, what sorts of dispositions are stable, what the psychological side effects of acquiring "absolutist" dispositions are, etc. What the consequentialist is really interested in, of course, is the ultimate *overall* consequences of the various types of dispositions and characters that are possibly acquired by people. It may be that people are able to develop both dispositions so strong that they would be unable to hurt a child in any circumstances, as well as more flexible dispositions that promote general avoidance of doing so. Which are to be preferred is not something that can be known *a priori*.

Consequentialists such as Railton thus seem to have a reply to the criticisms that consequentialism is lacking insofar as it cannot accommodate options and constraints. Of course, many critics of consequentialism might not agree that this reply is adequate. I myself believe there to be problems with this approach, which I will indicate shortly, but before I turn to those, let us consider the question of whether this approach is adequate as a reply to those who stress the importance of such things as options and constraints. A critic might charge that insofar as sophisticated consequentialism

only makes room for options, constraints and special obligations and permissions in terms of dispositions and patterns of behavior, which are themselves ultimately to be regarded as right only in terms of how such dispositions contribute to the overall good, it hasn't done full justice to the robust agent-relative importance of such things. In the Juan and Linda case, it seems that Juan's special concern for Linda is to be justified *only* by how it leads to better overall consequences than any alternative character he could have. But the moral worth of such agent-relative concerns, the critic might insist, is not derivative on the impersonal good, but is rather a *special* kind of value. But I think we must resist this suggestion. There are a number of responses that could be made by Railton or a like-minded thinker. Firstly, as Railton himself suggests, the consequentialist always has the option of including in her account of the intrinsic good such things as love or friendship; these themselves can be said to have impersonal value. This step taken, it doesn't seem nearly as implausible to suppose that Juan's concern for Linda is only justifiable only if it leads overall to better impersonal consequences. If we suppose instead that Juan's disproportional concern for Linda somehow makes things worse overall, making them or others miserable, so much so that the badness outweighs any good the love itself might represent, then it seems that Juan's disproportional concern *is* unjustified. One of the advantages of Railton's approach over the approach taken by the critics of consequentialism is that it provides a natural demarcation line as to exactly how far options and constraints should be allowed to go. Deontologists seem to be aware that projects cannot be pursued unconditionally, but they find it notoriously difficult to exactly draw the line as what limitations there are on their pursuit. They also seem to be aware that constraints do not always hold up; if the consequences of not breaking the constraint are bad enough, the right thing to do would be to break it. On Railton's account, we can provide a consequentialist answer to these questions based on the balance of overall value.

But there is one additional consideration that must be raised here. The arguments generally given for options and constraints by critics of consequentialism focus on the nature of persons as moral agents and as moral deliberators. They stress that it is natural for

agents to put extra weight on the concerns arising from their personal points of view, and that our integrity as moral agents demands concern more with what we actually bring about as opposed to what we only allow to happen. But these arguments cannot touch *objective consequentialism*. They may be objections to *subjective consequentialism*, which does make claims about how moral deliberation and moral agency ought to proceed. Objective consequentialism, however, is not meant as a theory of proper moral agency or moral deliberation, and so the attacks focusing on how well consequentialism makes sense of us as moral agents miss the mark. It cannot be that objective consequentialism adopts a model for moral agency that leaves out or fails to do justice to important features of persons as moral agents, because objective consequentialism simply does not include *any* model for moral agency or deliberation. It is meant only as an explication of the *truth conditions* for right action.

Of course, many deontologists will still be dissatisfied, believing that moral theory *must* provide a theory of how to morally deliberate because the whole point of morality is to guide us in our practical decision making. While I am somewhat sympathetic to this line of thought, we must be careful. Again, the analogy to logic is helpful. Logic need not be understood as an inquiry into how we ought to reason, nor provide a model for how people ought to go about making inferences. But of course there is need for such models and strategies, and one would hope that whatever such strategies are developed would at least be informed by logic itself. What we must resist is thinking that somehow consideration of actual processes and procedures of inferential cognition is to inform the very characterization of validity and invalidity. The sophisticated consequentialist can reason similarly about ethics. We certainly do not want morality to be entirely impractical, such that it couldn't help us at all in deciding how to raise our children or live our lives. But we must be careful about what we infer from this; we need not conclude that the actual procedures of ethical decision making should inform the characterization of right and wrong. We should also be careful not to fall prey to a potential equivocation. Certainly, when a teacher teaches her students cognitive strategies, there is some meaning of the word "logic" such that what she is

doing is teaching them “logic”, but this is not the same as logic in philosophical sense. Similarly, when we teach our children how to respond to moral quandaries and difficult life situations, we are teaching them “morality”, but not necessarily morality in the sense ethical theorists are primarily interested in.

Doubtlessly, the critic of consequentialism may still feel that this response on the part of the sophisticated consequentialist is not adequate. But it simply would not be felicitous to pursue this matter any further in this context, especially as I believe there are compelling reasons to think that both sides are ultimately mistaken. While I do think that Railton has given a compelling response to these objections, further exploration of the response reveals problems not only for consequentialism, but *any* moral theory focusing on right and wrong. I will be excused, therefore, for not pursuing this debate any further if I am right in thinking that there are bigger fish to fry.

First we must take a closer look at exactly how Railton’s strategy is supposed to work, and perhaps take it a bit further. As it is, it is rather unclear. In fact, at one point Railton comes very close to contradicting himself. To see this, let us return to the case of Juan and the decision of whether to visit Linda or send the check to Oxfam. Railton suggests that the objective consequentialist sees the act of visiting Linda as, strictly speaking, a *wrong act*. But, of course, it is not left at that. Railton also wants to stress that this wrong act of Juan’s is defensible, because it is made *inevitable* by Juan’s having developed a character which, of those available to him, leads to the best consequences. Railton seems to want to claim that, in visiting Linda, Juan is committing a *blameless wrongdoing*. However, Railton cannot consistently hold this. Railton supposes that, although it is true that Juan could still send the check to Oxfam if he tried, Juan will *not* try to send the check because it is inconsistent with the character he has rightly acquired. Railton seems to be operating with an understanding of what a character trait or disposition is such that to have one is to have limitations placed on what one will actually attempt to do in certain cases. He suggests that in order for Juan to have been able to send the check he would have had to have a different, and less good overall, character. Thus we must suppose that what it means to have a disposi-

tion to act in a certain way is for it to sometimes (we need not suppose always) be the case that one cannot lawfully act except in a certain way. There are at least some occasions, and this is one, in which Juan simply cannot put Oxfam or some other impersonal cause before Linda. However, if this is true, it means that Juan cannot actually send the check to Oxfam because in order to actually try (at least on this occasion) to perform this "better" act he would have had to have developed a different—and less good overall—character. However, according to consequentialism, *ought implies can*. The right act is the one (or one of the set of acts) available to the agent such that no other available act has better consequences. But sending the check is *not* an available act to Juan. Sending the check is impossible given the character he has *in fact* developed. Thus, Juan has *not* committed a wrong-doing by not sending the check.¹³ But this is directly contrary to what Railton himself says. Something is amiss here. I propose to attempt to sort it out, perhaps in the process revising or at least restating Railton's position, for I fear this will be necessary if any sense is to be made of his response.

It seems there are two possible revisions we could make to Railton's theory. The first would be to take seriously his claim that the objective consequentialist would still regard Juan's act of visiting Linda as a wrongdoing. To take this route, however, we will have to deny that it is inconsistent with his character to have stayed home and sent the check instead. The second revision would be to maintain that visiting Linda is inevitable given Juan's character, and that sending the check would have required Juan to have previously developed a different character. In the latter case, visiting Linda would not be claimed to be a wrongdoing at all. According to this strategy, Juan's visiting Linda is a necessary consequence of Juan's having developed the right sorts of dispositions and character traits, a perhaps unfortunate consequence of which is that he inevitably does something on occasion which, while not wrong, has less good consequences than something else we could *imagine* him doing if his character were slightly different. It seems to me that Railton is forced to adopt the second of these two approaches if his response is to have any bite. Imagine instead that he adopts the first revision, admitting that, even given his character, Juan

could still try to send the check on this occasion, and that he *would* send the check if he did try. From this it follows that the objective consequentialist must say that he *should* send the check. So far, there is nothing that Railton would find objectionable. However, the problem is that, understood in this way, the case can no longer be used as a response to the critic of consequentialism. Railton wanted this to be a case of defensible wrongdoing. Understood in this way, it is a wrongdoing. But is it defensible? Railton claimed it was blameless because it is the result of Juan's having acquired a certain character that is best, all things considered. But if we make the supposition that we have—that it is consistent with his character, even on this occasion, that Juan send the check instead—it is no longer clear how Juan's act is a *necessary* result of his having acquired this good character. If the act is not a *necessary* result of his having acquired this character, i.e. if he could both have the disposition and nevertheless *not* act in this way, in what way does his having that character make the act any less objectionable? Thus, Railton seems to lose his ability to use this case as a response to the critic, because he doesn't seem to have a way of defending Juan's actions even within sophisticated consequentialism, and hence his theory seems just as inadequate as subjective consequentialism at capturing the intuitions lying behind the original objections.

What we are left with, then, is the alternative way of revising Railton's account. Here, we insist that Juan *could not* have sent the check, thus admitting that his act was in no way wrong even according to objective consequentialism. This way is still able to stave off the objections of the consequentialist critic. In fact, it does so better, for here the consequentialist is actually able to *agree* with the non-consequentialist that Juan has done nothing wrong by putting his wife first. I think, moreover, that there are additional reasons to prefer this way of understanding the position. It seems to better capture what it *is* to have a deeply held project or relationship to someone. To love someone, arguably, necessarily involves being psychologically constituted to at least sometimes naturally act with disproportionate and special concern for her. Part of what it *is* for Juan to love Linda as much as he does *is* for it *sometimes* to be lawfully impossible for him to do things such as putting Oxfam before her when she has had a hard week. If Railton's "disposi-

tions" or "characters" are to capture such things, then it seems logical to think of them, too, as limitations placed on what sorts of things it is lawfully possible for a person to do. This coheres well with the teachings of modern psychology. Psychologists warn against thinking that any person is equally likely to perform any conceivable action; rather, what action an agent will take in a situation is a function of the way her character has been shaped and conditioned by past experiences, upbringing, environmental features and genetic tendencies (or, as they say, *nature* and *nurture*). Thus, we shouldn't assume that Juan is psychologically capable of sending the check merely because he has the money in the bank, knows how to write a check and is aware of the need. Whether Juan *will* or lawfully *can* send the check is partly a function of what his character traits are.

If this is the route for the consequentialist to take, the response to the critic will be rich indeed. We can imagine the sophisticated consequentialist saying roughly this to the critic. "You have pointed to a number of thought experiments and have insisted that, in these hypothetical scenarios, consequentialism tells us that something which is clearly wrong is right, be it twisting the child's arm, letting one's *own child* drown or forgoing one's own project to further that of some stranger. But it is not so simple. One cannot simply describe a hypothetical scenario without also talking about the character traits and dispositions of the agents in them. Consequentialism is a theory that picks the *right* act out of a list of available actions, but what acts actually are available is a function of the agent's personality, character and conditioning. If we enrich our cases with a consideration of these things, the puzzles dissolve. If I am an ideal sophisticated consequentialist agent, then it would have been right for me to have developed a disposition such that I am often *unable* to twist a child's arm, ignore my spouse or child in distress or abandon my own projects at the drop of a hat. Therefore, these are not, and in fact, *could not*, be the right things for me to do in these hypothetical cases. If we suppose instead that I have not developed these dispositions, then it may be true that my theory entails that the right thing for me to do is twist the arm or ignore my spouse, but this is only because of the antecedent wrongdoing of ignoring my duties with regard to my own moral character. Thus,

consequentialism is not so blind to these things as you suggest.”

This is an extremely powerful response on behalf of the consequentialist. In fact, I think in it lies one of the most profound and interesting insights that has arisen within 20th century ethical theory: that we cannot study acts in isolation. Any action that I make is the result of a long causal process, and was made possible by my having been socialized to have a certain sort of character and personality. In performing ethical assessment, we must not only examine the act in question, but (pardoning the Kantian jargon) *the conditions of its possibility*. Thus, it is not so much that the sophisticated consequentialist—so understood—accepts and approves of Juan’s decision not to send the money to Oxfam, but rather, that the sophisticated consequentialist accepts and approves of what Juan did such that the act of sending the money to Oxfam became *impossible*, that is, the act of acquiring a strong devotion to his wife. At this point, then, consequentialism may seem to have been vindicated. However, there are still problems on the horizon. The response given so far has gone a step in what I take to be the right direction: noticing that acts cannot be considered in isolation, and that we must understand acts as the result of causal processes involving people’s characters, personality traits, patterns of behavior, etc. The problem is that once down this road, there is no turning back until *ethical theory is no longer about right and wrong action at all*. While I find nothing objectionable with this, it tolls the death bell both for consequentialism and for its more deontological rivals. To this I now turn.

4. Psychological Determinism and the Real Failure of Consequentialism

In order to respond to its critics, consequentialism is forced to make some radical moves. What is radical about them? If I am right, they amount to a complete shift in the focus of moral assessment and moral theorizing. What we shall see is that this shift is one that consequentialism itself cannot survive.

What is the shift? Recall that the critics of consequentialism were operating from certain shared assumptions about what moral theory is about. In particular, they believed that moral theory ought

to tell us as individuals how to go about deliberation when faced with moral decisions. In fact, their criticisms of consequentialism stemmed largely from their focus on the features of moral agents as deliberators, and the claim that the model of moral deliberation put forth by consequentialism fails to do justice to these features of us as agents. Railton responded by making a distinction between objective and subjective consequentialism. It is true, one might admit, that subjective consequentialism falls prey to their objections. It actually attempts to provide a theory about how we ought to go about moral reasoning and deliberation, and therefore, if it fails to adequately capture essential features of the nature of moral agency, it fails as a theory. But objective consequentialism does not aim so high. Objective consequentialism does not aim to be a theory at all about how individuals should go about moral deliberation. It is meant only as a theory of the truth conditions of the rightness or wrongness of actions. Here, we have already seen one shift in what moral theory is supposed to provide for us. Rather than provide us with an actual method for moral deliberation, it aims only to provide us with criteria that distinguish right action from wrong action. While moral theory might eventually give us some help in moral deliberation, it is not its primary goal.

But if what I argued in the previous section is also correct, the shift does end there. As Railton conceives it, consequentialism is a theory that tells us not how to reason morally, but at least tells us whether certain acts are right or wrong. However, we must pose the question of *exactly what sort of* acts does (sophisticated) consequentialism tell us are right or wrong? Is it the actual acts taken by the agent in the situation, or the antecedent acts and decisions that lead to the development of the character and dispositions the agent has, which, in turn, lead to the action? It seems that it must be the latter. To see this, let us return to the case of Juan, the trip to see Linda and the letter to Oxfam. I argued that, contrary to what Railton explicitly says, he cannot claim that Juan committed a *wrong* act in visiting Linda. But it also would be misleading to conclude that Juan committed the *right* act on that occasion. We are presupposing that Juan's character development has made it inevitable that he visits Linda. Therefore, while it may be true that Juan had no available act that would have had better consequences,

this is only *trivially* true. It would be less misleading then to say that the sophisticated consequentialist simply withholds judgment about whether Juan's isolated act of visiting Linda is right or wrong. Rather, the important question for the sophisticated consequentialist is whether the *disposition or character* Juan has acquired is right or wrong. In this case, it is supposed that Juan's character is better than the others available, and hence, Juan has the *right* disposition.¹⁴ Therefore, we can see that the shift is indeed a large one. If Railton's approach works, it reorients moral theory not only away from providing a theory about how people ought to go about reasoning in difficult moral situations, but also away from even providing criteria for distinguishing right action from wrong action when what is under consideration is a particular act undertaken in a particular situation. For this sort of consequentialist, the interesting moral questions are questions about the right sort of moral character or dispositions a person should have; they are not questions about how a person ought to act on a particular occasion. How a person will act in a particular situation, it will be insisted, is a *function* of her character and dispositions, and so we ought to focus on not the rightness and wrongness of the act but the character and dispositions leading to it.

My objection to this strategy is not that it has shifted the focus of ethical theory. It is instead a question of whether it has shifted *far enough*. In concluding that what action a person will take in a particular instance is the result of the confluence of the character traits and dispositions the person has antecedently acquired, the approach taken by the consequentialist is one that presupposes the truth of what could be called *psychological determinism*, or at least of something very close to it.¹⁵ This is the thesis that when a person finds herself in a certain situation or environment (i.e. state of the world) in which she is faced with a decision, there is only one lawfully possible action available to her compatible with the character traits and behavioral dispositions she has antecedently acquired. I argued earlier that we cannot make sense of Railton's claims in the scenario about Juan and Linda unless we suppose that Juan's visit was made inevitable by his character. Psychological determinism is assuredly a controversial empirical hypothesis, but it seems to me to be a plausible one given the views of the

majority of contemporary psychologists. Let us, for the moment, grant it as a hypothesis. The question now becomes, what does this mean for moral theory? I suggest that once this hypothesis is made, it not only damns attempting to make assessments about right or wrong when applied to individual actions, but undermines the very coherence of talking in terms of right or wrong at all in moral theory, including the rightness or wrongness of dispositions and characters.

I have argued this in somewhat greater detail in a somewhat different context elsewhere.¹⁶ My argument there was roughly this. The distinction between right and wrong in ethical theory depends on being able to develop a list of alternative actions a person can perform in some set of circumstances and then adducing criteria by which to determine which of those actions are right and which are wrong. Consequentialism does this by ranking the available options according to the sum total of goodness and badness that results from each. Other theories take a different approach. However, if psychological determinism should prove to be true, and in each case there is only one lawfully possible action for an agent given her character, upbringing, dispositions and the environment and situation she finds herself in, then there is no plausible way of generating the requisite list of alternatives. Arguments attempting to show that the rightness and wrongness of an act does not presuppose alternative possibilities as well as attempts at defining the list of alternative possibilities independently from what is actually lawfully possible are found to be unworkable. Because I have argued this in more detail elsewhere, I will not duplicate things here. Rather, we must investigate how this applies to this debate.

I argued earlier that if the consequentialist takes Railton's strategy of replying to the critics, it ends up shifting the focus or theme of ethical theory. In assessing the actions and character of Juan, the sophisticated consequentialist cannot assess Juan's individual act of visiting Linda as either right or wrong, but is instead pushed to assess the broader character Juan has that made the act inevitable. The problem is, how is the consequentialist to make even this latter assessment? Consequentialism is a theory about right and wrong, holding that what is right is that of the available options with maximal consequences. If the assessment of Juan's character is to be at

all consequentialist, it must be that Juan's act of having acquired this character leads to better overall results than any other possible modes of character development Juan might have taken. But the question, of course, is *what other modes of character development?* If psychological determinism is true, then every act Juan ever took and every decision he made—even those regarding what sorts of dispositions or traits he cultivates in himself—were determined. There are no other possible characters Juan might have had or other dispositions he might have developed. How, then, can consequentialists make any sort of determination of whether his character is the *right* one for him any more than they can make a determination about whether his acts are right or wrong?

Recall that the sophisticated consequentialist response to deontological critics was to argue that one cannot simply pose a thought experiment or hypothetical situation without also considering the limitations placed on what sort of acts a person can take in such a situation by her upbringing, conditioning and character. Therefore, one cannot prove consequentialism wrong by showing that in the favored deontological cases, consequentialism gives the wrong answer. Consequentialists won't simply admit that it is right to twist the arm, or ignore your wife in favor of Oxfam. These options are not always available, and are likely not available at all if the agent in question is a good sophisticated consequentialist agent. One cannot naively assume that any act we can conceive of a person doing in a situation is in fact available to that agent. Constraints are placed upon what is even possible by what has happened before and how a person's personality and character have been shaped by past events and past decisions. But notice that this same naiveté is present in the consequentialist assumption that any character traits we can imagine Juan as having are possible for him to have. It's true that every person's personality, behavioral dispositions and character change over time. We make decisions about our lives, strive to become different sorts of people, make and break habits, train ourselves to think certain ways, etc. But all of the decisions we make about our own character and what sorts of dispositions and traits we have are equally determined, if psychological determinism is true. How, then, are we to claim that certain possible characters are right while others are wrong? In the end, the

character a person has is the only one that was lawfully possible for him or her.

I do think there is something right about the sort of sophisticated consequentialism suggested by Railton. What is right about it is that it recognizes that human behavior does not come out of nowhere. What actions we perform on what occasions is a function of the situations we find ourselves in and the ways in which we have trained ourselves to have certain characters and to have certain dispositions as to how to think, reason and act. Sophisticated consequentialism understands humans as part of the causal world, as beings who not only cause things to be the case, but have their personalities and behavior shaped by genetics and environment. The problem is that sophisticated consequentialism does not go far enough. If we truly are to understand not only human behavior but humans themselves as the result of causal processes, a sort of infusion of the input of genetics, environment, upbringing and similar forces, then unless we are apt to see a person's genetics or the state of the world a person finds herself in as *right* or *wrong*, it makes little sense to make such an evaluation of the necessary consequences of these things.

The problem, therefore, with sophisticated consequentialism is that it is still consequentialism, and consequentialism is a theory about right and wrong. Where sophisticated consequentialism goes wrong is not in realizing that isolated actions cannot be naively labeled as right or wrong independently of the consideration of the characters and dispositions leading to them. Rather, where sophisticated consequentialism goes wrong is thinking that we must therefore try instead to apply the same criteria for rightness and wrongness that had previously been applied to acts in isolation to the assessment of the development of a person's character. It is as if it was simply assumed that ethical theory had to somewhere include the concepts of rightness and wrongness and since it was realized that these could not be applied to acts *simpliciter*, they were instead applied to the dispositions and characters leading to them.

If we go down the route I am suggesting, I think we must abandon any moral theory that centers itself around the concepts of right and wrong. If psychological determinism is true, then there is no basis for determining the list of possible alternatives to an ac-

tion such that we can make distinctions about what is right and wrong from that list. Railton's response to the objections to consequentialism, however, seem to wed consequentialism's hopes to psychological determinism or something relevantly similar. Therefore, ultimately, what Railton has provided us is not a defense of consequentialism, but an argument for its very incoherence. This is not to say that consequentialism's more deontological rivals fare any better. They too center around the determination of right and wrong. In fact, if anything, their view is the worse off of the two. They naively assume that a theory of what makes an act right and wrong must also be a theory about how agents ought to go about conducting moral deliberation. This being the case, they insist that moral theory must do justice to the nature of persons as moral agents and thus to the personal point of view with which we enter into moral deliberations. As much as their theory stresses the nature of moral agency, it seems that it leaves out some of the most important facts about moral agents: that what decisions we actually make when faced with moral decisions are shaped by our characters, motivational dispositions and past development. In fact, it doesn't seem to me that any current ethical theory has truly come to grips with this fact about human agency.

Of course, there remains a response that I have not considered, and that is simply to deny psychological determinism, insisting instead that there are always, or at least in the interesting moral cases considered earlier, multiple actions the agent in question could take compatible with the laws of nature, current state of the world and the antecedently acquired character of the agent. This may be a genuine alternative to the sort of revisions to moral theory I am suggesting. I admitted earlier that psychological determinism is a controversial empirical hypothesis. In fact, it may be false. I have been merely assuming its truth because it seemed the only way to make sense of Railton's response. Therefore, it probably is not a viable option for Railton, or for anyone who wishes to use his sort of response, to deny psychological determinism.¹⁷ But it may remain a viable option for a *critic* of consequentialism to make this move. But it seems to me that such a theorist does this at her own risk. For all its controversy, it seems to me that psychological determinism—or something relevantly similar to it—is plausible and

consistent with the teaching of modern psychology. In fact, its denial seems to entail granting some special metaphysical or causal status to human beings that seems out of sorts with the modern trend of regarding humans and their actions as simply part and parcel of the natural world, as amenable to scientific explanation and prediction as chemical compounds, amoebas and planetary bodies. But this leads us to other debates in philosophy, ones it is best not to tackle here. I am content simply to remain at an impasse with those who disagree with psychological determinism. I am more interested in the question of what follows *if* it is true. It may not be as bad for moral theory as one might fear.

5. *Ethics Without Right and Wrong*

Many people assume that ethics *must* be about right and wrong. This attitude is so pervasive that some will conclude that if my conclusion is that ethics is not about right and wrong, this serves as a *reductio ad absurdum* of my entire position. Others may conclude instead from what I have said so far that if determinism is true, ethics becomes meaningless or pointless. I would like to argue instead that even if psychological determinism is true and ethics must abandon right and wrong as its organizing concepts, there can still be a very interesting and fruitful place for moral theory in not only academia but in society.

Sophisticated consequentialism *pace* Railton takes a step towards the right sort of ethical theory. I argued in the previous section that this approach shifts the role of ethics such that it not only does not attempt to provide a model for "correct" moral agency but also does not even attempt to evaluate *individual* acts as right or wrong. It instead focuses on the evaluations of whole patterns of behavior or dispositions or types of characters. I have argued, however, that this theory is mistaken for clinging to the idea that we can sensibly evaluate certain characters or dispositions as being right and others as being wrong for a particular person. However, I do believe that this general model of what ethics ought to concern itself with is correct. Ethics ought to concern itself not with the assessment of individual acts, but with the assessment of whole patterns of behavior and types of acts, and how these pat-

terns relate to personality and character types. Like Railton, and probably more so, I am aware that actions have to be understood as the result of a causal process, and that we cannot examine an act in isolation from how the conditions that made the act possible (i.e. the character and dispositions of the agent performing it) have other effects that also have to be understood and taken into account. What I propose is that ethics is the study of different causal relationships between actions, certain dispositions and their consequences combined with the use of the knowledge gained by this study in the attempt to make people into better moral agents.

I need to be much clearer about exactly what I have in mind. The first thing to note is that nothing here calls into question any theory about value. Consequentialists tend to assume that there is such a thing as good and bad, and that it can be measured. If we assume this to be the case, we have the grounds for developing an interesting moral theory. We can study people's actions and their effects, and even measure how much good is brought about by certain actions. Pooling and analyzing this information, we can gain very important insights. While we cannot say, for example, that Juan's developing a deep devotion to his wife leads to better consequences than the other options available to him (for there were none), we might discover that, statistically speaking, people who develop such deep attachments for their spouses tend, overall, to produce more happiness, desire-satisfaction, or other form of intrinsic good (whatever it happens to be, according to our axiology) than people who attempt instead to treat people impartially. Having this information will be extremely useful, and in a derivative sense, its propagation might help people like Juan make decisions. There is no end to the sort of information we could gather.

It may even be hoped that this sort of ethics could lead to moral progress, via what we could call "moral technology". Moral technology is simply our use of the information we acquire through such studies to improve society. After we've obtained copious information regarding, e.g., whether children who are taught to have a disposition against lying end up producing more good than children who receive no such teaching, we might use this information to set up a society in which children receive a sort of moral training that in general leads to a happier and better society overall than

the sort of society we live in now. (Notice, however, that I do not say that such children would receive the "right" sort of moral training, nor that our present moral training is "wrong".) This sort of technology is just like any other sort of technology. Scientists have studied all sorts of patterns of interactions regarding physical objects and how they respond to various stimuli, and after collecting such information, they have been able to put it to use to build better bridges, faster computers and safer cars. But they do not say that the bridges, computers and cars of the past were "wrong" for not living up to the new inventions. Both the new and the old things operate entirely according to the laws of nature. I don't see any reason to think any differently about humans.

In fact, the pursuit of moral technology might lead us eventually to be able to provide exactly what many people will think is lacking from this sort of moral theory. I mentioned earlier that many theorists, especially critics of consequentialism, believe that ethics must, as a matter of definition, provide us with a way of understanding how people ought to go about reasoning and deliberating about morally relevant choices. I believe ethics can do this, but only at a derivative level. We can perform studies of what the results are of people being trained to approach moral decision making in different ways. We could compare people who are trained to think like subjective consequentialists with people who are trained to abide by the categorical imperative in their deliberations. In all likelihood, neither of these will actually be the sort of deliberation strategy that will end up producing the best overall effects. The rule of "try to maximize consequences" is far too broad a rule to really aid people in any sort of concrete decision making processes. Psychologists generally see people as making use of *heuristics* in decision making, highly specific rules of thumb that are particular to one or more regions of thought. This is simply what works best for people, and I imagine that the best sort of moral education will be some combination of rules of thumb mixed with some more open-ended principles (e.g. the golden rules) to provide some direction in unfamiliar cases. But from my perspective, these are all *a posteriori* questions to be answered only after a long period of empirical research. Nevertheless, one shouldn't conclude that psychological determinism precludes the possibility of the use of ethi-

cal theory to aid people in moral decision making. That may be one of its eventual functions. I simply resist seeing it as the fundamental or organizing function.

I should be clearer about what I mean when I suggest that the use of this sort of "moral technology" might aid people in their decision making or improve society. Of course, if my theory is to be compatible with determinism, I cannot hold that moral technology will help people to select one of numerous options that are actually lawfully possible for them at a certain time, or that moral technology could be used to make society better than we might lawfully have made it instead. It might then be objected that what I am suggesting is self-refuting; if we cannot alter our future actions, how can we make things better through moral technology? If all of our decisions are determined already, how can this sort of ethical theory aid our decision making? Here we must be careful. Firstly, when it comes to individual decisions, one must not confuse determinism with fatalism (understood as the view that one's own decisions do not affect her own fate). Any plausible version of psychological determinism will still hold that people make decisions and that those decisions affect their actions; it will only add that these very decisions are themselves made inevitable by previous events and causal processes already at work. So the psychological determinist does believe that decisions lead to action. What decisions an agent makes in a given situation is a factor of many causal influences, but one important factor is the knowledge the agent has. When I suggest that research into different causal relationships between various types of actions and good and bad consequences might help to aid people in moral decision making, I am merely pointing out that we have good reason to believe that those people who are knowledgeable about such research will make decisions that on average lead them to take actions that have consequences better than actions taken by others who have no such knowledge. I do not mean to suggest that the decisions they make will be better than other lawfully possible decisions they do not make. When it comes to the improvement of society, here too we must be careful. When I say that we can use moral technology to "alter" society for the better, I do not mean to suggest that we can create a society better than others we might have had instead. Rather, I mean

we can change society *over time* and improve it from what it was, in the exact same sense that computer scientists and technicians have improved computers over time. Determinism is not the same thing as complete lack of change. Even if the world is a determined one, people and societies can (and often do) improve and alter themselves for the better, and oftentimes the catalyst for such changes is knowledge.

When I speak of such things as using empirical research to discover what sorts of dispositions and actions and patterns of behavior lead to better overall consequences than others, it may seem to some that right and wrong have now crept back into my theory. In a sense, I agree (although this terminology may be a bit misleading). One way of describing the process described above is as an inquiry into what the *right* sort of moral education is. I do think that we can perhaps devise some derivative uses for "right" and "wrong" in moral theory. For example, I find it likely that study will find that slavery is *wrong*, where what is meant by this is that societies that incorporate slavery produce much less good (all other things being equal, etc.) than those that do not. One could imagine that, therefore, we could similarly investigate the rightness or wrongness of such things as pornography, prostitution and animal rights. But we must be very careful. This sort of statistical analysis of "right" and "wrong" can only be applied to questions about *types* of acts, character traits and institutions. I don't see any way of applying it to individual acts or characters. Thus, for example, we might discover that "being devoted to one's spouse" is a good thing, and in this sense, *right*, because those people who are so devoted, other things being equal, tend statistically to bring about better consequences than those who are not. However, I do not see any similar way of analyzing Juan's devotion in particular. To say that Juan's devotion is right is to compare it to something else that might have happened, but we've already discussed why this will not work.

One might attempt to create a sense of "right" that applies to specific actions by identifying them as members of a class of actions that are "right" in the sense described earlier. Juan is right for being devoted to Linda in the sense that his devotion is one instance of devotion to one's spouse in general, which we have determined to be right by statistical means. However, if we do this,

we will quickly run into problems. Almost any act is amenable to all sorts of different descriptions, and whether it will be right or wrong in this sense will depend on which description is being used. For example, it is probably right for parents to spend a great deal of time touching their children when what we mean by this is that parents who touch their children a lot tend overall to produce better consequences than those who do not. But the physically or abusive parent might also be one who touches her child a lot. Certainly, parents who abuse their children tend overall to produce worse consequences than those who do not. Thus, the particular actions of abusive parents are "right" under one description, and "wrong" under another description. This will only lead to confusion. I think it is best if we simply do not attempt to try to apply this theory to the rightness and wrongness of individual action.

Surely, there will be many who will be disappointed with any moral theory that does not provide a method for determining the rightness or wrongness of particular actions. But I do not share this view. I have already described how moral theory could acquire interesting information about various causal processes relating to the production of good and bad states of affairs, and how the acquisition of such information could lead to moral progress and aid us in our moral decision making. What else is really left to be desired? My own political commitments make me wary particularly of those who attempt to use a theory of moral right or wrong (or of moral responsibility) to justify seeing themselves as superior moral agents (and thereby justifying for themselves a greater amount of resources, etc.) On my own theory, we must understand all people's actions, even the criminal, as ultimately the result of causal processes. We can use our knowledge of such causal processes to make the world a better place. However, the point is not to differentiate people into the morally worthy and the morally unworthy. I see no purpose for this whatsoever.

The reader will have probably noticed that my theory is not really so different from consequentialism after all. In fact, I do think that the spirit, if not the letter, of consequentialism is true. As Rawls would likely put it (although he would disagree), I believe that ethical theory must *begin* with the theory of the good and proceed from there. My difference with the consequentialist is simply

how quickly one can (or whether one can at all) proceed from there to a theory of the right. For these reasons, I am a bit less bothered than many at the dominant role consequentialism has come to play in moral theory. In the end, however, we can do better.¹⁸

Notes

¹ See, e.g. Christine Korsgaard, "The Reasons We Can Share", *Social Philosophy and Policy* 10 (1993), 24-51, Scanlon, T. M. "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. A. Sen and B. Williams. (New York: Cambridge University Press. 1982), and Philippa Foot, "Utilitarianism and the Virtues." *Mind* 94 (1985), reprinted in *Consequentialism and its Critics*, ed. Samuel Scheffler (New York: OUP, 1988), 224-242.

² Samuel Scheffler, *The Rejection of Consequentialism*, rev. ed. (New York: OUP, 1994).

³ Bernard Williams, "A Critique of Utilitarianism," in *Utilitarianism: For and Against*, ed. J. J. C. Smart and B. Williams (New York: Cambridge, 1973).

⁴ Scheffler himself would probably put the point somewhat weaker, i.e. only that consequentialism is not the only way of taking into account these features of persons. There are interesting questions here, but the difference is not so important for my argument here.

⁵ Diane Jeske and Richard Fumerton, "Relatives and Relativism: A Problem for Consequentialism," *Philosophical Studies* 87 (1997), 143-157.

⁶ Scheffler, for example, denies that constraints can be justified in the same way as options.

⁷ Thomas Nagel, *The View From Nowhere* (New York: OUP, 1986), 176.

⁸ Peter Railton, "Alienation, Consequentialism and the Demands of Morality", *Philosophy and Public Affairs* 13 (1984), reprinted in *Consequentialism and its Critics*, 93-133.

⁹ Railton, "Alienation, Consequentialism and the Demands of Morality", 114.

¹⁰ Railton, "Alienation, Consequentialism and the Demands of Morality", 110-111.

¹¹ Railton, "Alienation, Consequentialism and the Demands of Morality", 120-121.

¹² Kevin Klement, "Is Pacifism Irrational?" *Peace Review* 11 (1999), 165-170.

¹³ I should also point out, in all fairness, that it may be that Railton is operating with a slightly different conception of what it means for an act to be available, perhaps adopting Moore's suggestion that available acts include all acts I *would* do if I so choose (even if I cannot lawfully so choose). Elsewhere, I discuss and harshly criticize this method of defining alternative possibilities. See Kevin Klement, "Determinism and Right and Wrong", forthcoming.

¹⁴ Railton or a like-minded thinker might, instead of talking about a disposition as being right or wrong, talk instead of the acts that lead to a certain disposition being right or wrong. Neither is tenable, for reasons I will give shortly.

¹⁵ The prefix "psychological" in "psychological determinism" is meant to distinguish this form of determinism from a more radical thesis regarding all the lawful processes in the world, which holds that given the current state of the world at a particular moment and the laws of nature, only one future is lawfully possible. Psychological determinism need not be so strong. It is concerned only with *human behavior*, and is only concerned with it at the level at which we are likely to describe it as *behavior*. Thus, psychological determinism might be compatible with indeterminism in the motions of quantum particles that makes no difference to human behavior at the macro level. For more on this, see my "Determinism and Right and Wrong". To avoid possible misunderstandings I should also point out that I am not here making any claims about the compatibility of this form of determinism and so-called "free will".

¹⁶ Klement, "Determinism and Right and Wrong".

¹⁷ Moreover, there are additional ways in which consequentialism presupposes determinism. Consequentialism presupposes that there is a fact of the matter about what the consequences of my action will be at the time that I perform the act. Obviously, the effects of an action depend on how others react to it. But if there are multiple lawful possibilities in terms of how others will respond to an action I take, then there is no fact of the matter *at the time that I perform my action* what the consequences will be, and therefore no fact of the matter of whether that action is right or wrong. I think there are interesting issues here, but it would take me too far afield to discuss them in detail here. I will add in passing, however, that I think many alternatives to consequentialism make similar determinist presuppositions.

¹⁸ I would like to thank Diane Jeske for helpful comments on an earlier version of this paper.