
CONSEQUENTIALISM AND COMMITMENT

BY

ALASTAIR NORCROSS

Abstract: It is sometimes claimed that a consequentialist theory such as utilitarianism has problems accommodating the importance of personal commitments to other people. However, by emphasizing the distinction between criteria of rightness and decision procedures, a consequentialist can allow for non-consequentialist decision procedures, such as acting directly on the promptings of natural affection. Furthermore, such non-consequentialist motivational structures can co-exist happily with a commitment to consequentialism. It is possible to be a self-reflective consequentialist who has genuine commitments to individuals and to moral principles, without engaging in self-deception.

Personal commitments, to people and, to a lesser extent, to principles, are a very important part of most people's lives. They give shape and meaning to our lives, and help to constitute our identity. To say that I am Diana's husband and David's father is to give more than mere relational information about me. It says something fundamental about my motivational structure. There are things I would do for Diana and David, risks I would take, sacrifices I would endure, that I wouldn't even consider for a perfect stranger, or even a fairly close friend. For most of us, it is hard to imagine what a life devoid of such commitments would be like. And yet, certain moral theories seem to require that we abandon them, or at least be prepared to abandon them whenever they conflict with the demands of impersonal morality. Both consequentialist theories and Kantian versions of deontology have been accused of being inimical to commitments. My concern in this article is with consequentialist theories, in particular utilitarianism.

In a famous phrase, Bentham required “everybody to count for one, and nobody for more than one”, and Mill said of the utilitarian agent, “As between his own happiness and that of others, utilitarianism requires him to be as strictly impartial as a disinterested and benevolent spectator.”¹ On the face of it, a good utilitarian agent cannot have strong personal commitments to friends or family. A minimum requirement for me to have a personal commitment to an individual is that, at least sometimes, the welfare of that individual is more important to me than that of a randomly selected stranger. Perhaps some consequentialists will declare “so much the worse for commitments”, but most of us are loath to give up so easily our conviction that commitments play a vital role in our moral lives. I argue in this article that consequentialism can accommodate the importance of commitments without incurring further problems.

In section 1, I present the basic utilitarian defense of commitments – that they are beneficial from the point of view of utility – and defend the utilitarian against the charge that this renders her endorsement of commitments unacceptably contingent. In section 2, I consider the criticism, most notably urged by Bernard Williams and Michael Stocker, that this defense of commitments leads to alienation when a utilitarian justifies *her own* commitments this way. It appears to be incompatible with the notion of acting on a commitment that an agent reassure herself, while she is acting, that having the commitment is for the best. However, this criticism rests on the mistaken assumption that utilitarians must advocate always reasoning from utilitarian principles. An examination of Sidgwick’s distinction between criteria of rightness and decision procedures shows that utilitarians need not, and indeed should not, advocate continual calculation. In section 3, I consider Williams’s and Stocker’s responses to the claims of section 2. Once the utilitarian admits that there are situations in which utilitarian reasoning should not be employed, it is open to the critic to charge that utilitarianism demands that such reasoning never be employed. Both Williams and Stocker endorse the following claims: (i) if utilitarianism is true, it would be best if at least most people didn’t believe it to be true; and (ii) the truth of (i) exposes a defect in utilitarianism. I argue that neither Williams nor Stocker provide good reasons for accepting either (i) or (ii). Even so, (ii), or rather a generalized version that applies to all moral theories, appears to be independently plausible. It seems right that if a moral theory is correct, it should be possible for everyone, or at least most people, to act as it prescribes and to believe that it is correct. In section 4, I argue that such a generalized version of (ii) is, at most, contingently true, and cannot be a necessary feature of an acceptable moral theory. I also argue that the temptation to think otherwise stems from an impoverished view of how reasons for behavior can operate. In section 5, I present Peter

Railton's recent account of what he calls "sophisticated consequentialism", which allows that a "standing commitment to leading an objectively consequentialist life" can co-exist happily with deep personal commitments to other people, that sometimes result in performing a less than optimal action. I also demonstrate how the same approach can apply to moral commitments to non-consequentialist moral principles and can accommodate the view that regret or remorse can be appropriate even when one has knowingly done the right thing. In response to Railton's account, William Willcox has argued that the sophisticated consequentialist agent who follows the promptings of natural affection against the dictates of consequentialism must, if she is truly a consequentialist, be self-deceived. I argue that this charge fails. I also argue, in section 6, that the charge of self-deception is more troubling in the case of an agent who follows a non-consequentialist moral commitment against the dictates of consequentialism. In response to this, I present an account of moral commitments that allows the sophisticated consequentialist to avoid self-deception, even when knowingly failing to optimize.

1. Consequentialism and personal commitments

It is sometimes claimed that utilitarianism cannot account for the importance of personal commitments to other people. Consider the following pair of cases: (i) John is at a party also attended by Jane and Mary. John knows that Jane would suffer terribly if he went home with Mary. John goes home with Mary. Jane discovers that John has gone home with Mary, and suffers terribly. (ii) The same as (i), except that the characters are Bill, Betty and Sue. (Betty is the one who suffers.) There is not a lot we can say about the moral status of these two examples without more information. We can, perhaps, say that it is a bad thing that Jane and Betty suffer. Let us add the following information to our descriptions of the two cases: (i) John is married to Jane. He has always professed love and loyalty to Jane, and she has to him. She has been an excellent spouse and given him no cause for complaint. (ii) Betty and Sue are rivals for Bill's affections. He has given neither of them any reason to think that he preferred her to the other. Clearly, this information enables us to say a lot more about the moral character of John's and Bill's behavior. As far as we can tell, John is behaving very badly indeed. He is betraying his wife. Bill, on the other hand, may be accused of tactlessness in his handling of the situation, which, given the feelings of Betty which are damaged, is not a trivial charge. We may also say that it is unfortunate that Betty is hurt in this situation. We won't be able to accuse Bill of behaving anything like as badly as John. It seems clear, then, that the two situations are morally different. It is sometimes said

that utilitarianism is unable to do justice to the difference between such pairs of situations. After all, it is claimed, the utilitarian looks only to the consequences of actions in terms of the pleasure or pain which is produced. On this count, the two situations seem to be similar: John runs off with Mary, causing Jane pain; Bill runs off with Sue, causing Betty pain.

The utilitarian, it is claimed, cannot account for the special nature of betrayal. What John does to Jane is basically the same as what Bill does to Betty. He causes her pain. How can a utilitarian reply? If a utilitarian wants to make a moral distinction between two actions, she should show a distinction between the consequences of those actions. Can she do that in this case? Jane and Betty are both caused pain, but perhaps there is a significant difference in the pain they are caused. Betty's pain is the pain of jealousy, disappointment, humiliation and unrequited lust. There is no doubt that Betty's pain is severe and highly regrettable from a utilitarian viewpoint. Jane's pain is the pain of betrayal. It is the peculiar pain which can only be caused by someone who has a special relationship with the sufferer. Bill could not betray Betty, because he doesn't have the kind of relationship with her which would put him in a position to betray her. In order for a utilitarian to account for the difference between Jane's pain and Betty's pain, the former would have to be clearly worse than the latter. It is fairly plausible to assume that this requirement would be met. The pain caused by betrayal is usually worse than the pain caused by jealousy, disappointment, humiliation and unrequited lust.

But is this enough to account for the difference between the two cases? I suspect that it isn't. This account relies on the contingent feature of the world that betrayal causes worse pain than certain other types of behavior. But what if Betty were different from other women, in that she felt the same sort of pain as Jane when Bill ran off with Sue, even though Bill wasn't betraying her? This doesn't seem to be conceptually impossible. Wouldn't the utilitarian have to say that the two situations would then be morally alike? And yet, wouldn't we still want to say that there was an important difference between the cases? After all, no matter what Betty's psychological make-up, it is still true that Bill is not betraying her, and that John is betraying Jane. Worse still, what if Jane is unaware of John's behavior, and so doesn't feel any pain at all? Wouldn't the utilitarian have to say that there is nothing at all wrong with John's behavior in this case? In fact, it might even be admirable behavior, if he and Mary enjoy it enough.

A consideration of these sorts of cases may prompt the charge against utilitarianism that it cannot account for the fact that some people stand in special relations to others. The utilitarian is concerned only with the amount of pleasure and pain, or happiness and unhappiness, in the world, and is not concerned with questions about which individuals experience

those states as a result of the actions of which other individuals, except inasmuch as the answers to such questions will enable them to produce more pleasure or happiness.

I think that the utilitarian can give an answer to these charges, though it won't satisfy everyone (some people are never satisfied). First of all, she will recognize the fact that we simply do value certain people more highly than we value others. The praise of a loved one is usually more important than that of a stranger. (There are, of course, exceptions.) A cutting remark from a stranger or from someone whose opinion you despise is usually far less painful than a similar remark from a loved one. These are facts which the utilitarian is just as much in a position to acknowledge as any other moral theorist. Furthermore, she can acknowledge that a life in which one values a few people more highly than others is usually a happier more rewarding life than one in which one values all people equally. (Again, there can be exceptions. Perhaps Mother Theresa valued all people equally and had a far more rewarding life than our miserable self-centered existences.) This is probably just a fact about the way we are.² It may be that most of us are so constituted that a life worth living must involve special relationships. If that is so, the utilitarian can acknowledge that fact. What is especially wrong with betrayal, whether or not it actually causes pain, is that it shows the agent to be deficient in some respect with regard to the values he must have in order to lead a good life and to enable those around him to do likewise.

The utilitarian can admit that John's betrayal of Jane shows him to be morally deficient. Given our psychological makeup, personal commitments, to people and to principles, enable us to be good moral agents. This is only a contingent truth, but it is true nonetheless. This is not to say that such commitments are *always* good or that one cannot be a good moral agent without them. It is a strength of the utilitarian position that it can acknowledge the importance of commitments without making them either necessary or sufficient for good moral character. Someone who is generally lacking in commitments to people and to principles usually experiences and causes less happiness than someone who has such commitments. The utilitarian can readily admit this. She can, that is, heartily approve of someone with plenty of commitments, on the grounds that such a person is likely to produce more good than someone without such commitments.

Someone might object at this point that the consequentialist must regard the goodness of commitments as merely the regrettable result of the limited range of dispositions available to us. In this sense, she might continue, the consequentialist is less than wholehearted in her endorsement of commitments. Their value is merely the product of the regrettable lack of plasticity in human nature.³

Two types of response are appropriate here, both of which admit, or rather assert, the main point of the criticism, which is that, however good commitments are, the world would be a better place if people could do more good without them. The first response is to ask why this should count as a criticism at all. Why *should* consequentialism wholeheartedly endorse commitments? Presumably, because our untutored moral intuitions do so. But untutored (or even tutored) moral intuitions have wholeheartedly endorsed many things (such as slavery, sexism, trickle-down economics) that we have later come to see as not merely not unconditionally good, but downright abhorrent. One of the appeals of consequentialism is that it subjects all our received opinions to rigorous scrutiny. This is not to say that a consequentialist should simply dismiss deeply felt moral intuitions out of hand. It would strengthen her position if she could explain, and perhaps even justify, such intuitions. In the case of the intuition that commitments are to be wholeheartedly endorsed, she can do just that. Given the role that commitments play in producing utility, it is both unsurprising and a good thing, from the consequentialist point of view, that most people wholeheartedly endorse them.

The second response to the present criticism is to question the sense in which the consequentialist is less than wholehearted in her endorsement of commitments. Given a psychologically possible choice between being a person with commitments and being one without, she will wholeheartedly opt for the former. It is true that there is a logically, but not psychologically, possible character type, which avoids the weaknesses of both her options. But the consequentialist doesn't waste time lamenting the unavailability of the merely logically possible. Imagine a gourmand, Michael, who is offered his choice of any meal he wants at any restaurant in the world. Is Michael's endorsement of chef Henri's thirteen-course *Spécialité de la Maison* at the *Quatre Saisons* in Paris less than wholehearted, just because he can imagine a logically possible alternative meal that would give even more gustatory pleasure, on condition that human taste buds and digestive systems functioned otherwise than they do? Such a possibility might keep Michael amused in his idle moments between meals, but it certainly doesn't render his enthusiasm for Henri's creations any less than wholehearted.

2. Williams and Sidgwick on commitments

It has been claimed, most notably by Bernard Williams and Michael Stocker, that the utilitarian justification of commitments outlined in the previous section runs into problems in a case in which a utilitarian seeks to justify her own commitments. Williams describes a situation in which a man chooses to rescue his wife rather than some other person: "It might

have been hoped by some (for instance, by his wife) that his motivating thought, fully spelled out, would be the thought that it was his wife, not that it was his wife and that in situations of this kind it is permissible to save one's wife." What is being criticized is not the agent's action, but his decision procedure: the conscious application of a principle of morality in a situation in which the agent should act simply on his natural affection. If the agent is a utilitarian, he doesn't simply act on his commitment to his wife, but he also justifies such a commitment on utilitarian grounds. This, according to Williams, leaves him with "one thought too many". The problem here is not just that the wife might have hoped that her husband's commitment to her would have provided the sole motivation for his action. Rather, it seems that the kind of conscious evaluation of his commitment that appears in Williams's example is actually inconsistent with having the commitment in the first place. What it is to have a strong commitment to someone is, in part, to be disposed to perform certain actions without subjecting either the actions or the disposition to independent moral scrutiny.

The sort of criticism of utilitarianism that appears in the previous paragraph has been so widespread in recent years that it is somewhat surprising to discover that the groundwork for an effective utilitarian response was laid nearly a hundred years ago by one of the best known classical utilitarians, Sidgwick, in *The Methods of Ethics*. Williams is, I presume, familiar with Sidgwick's work. Does Sidgwick leave his agent with one thought too many? Let us consider whether Williams's charges had already been answered before he leveled them.

Sidgwick argues that agents should be especially concerned with benefiting their 'near and dear'. Let us look at what he has to say on the matter:

In the first place, generally speaking, each man is better able to provide for his own happiness than for that of other persons, from his more intimate knowledge of his own desires and needs, and his greater opportunities of gratifying them. And besides, it is *under the stimulus of self-interest* that the active energies of most men are most easily and thoroughly drawn out.⁴

The second claim is the more interesting in this context, since it is about motivation. Sidgwick's claim is that a particular form of motivation, self-interest, is most effective in 'drawing out' energies, and thus in producing results. He also defends a bias towards one's near and dear, that is, a claim that "each individual should distribute his beneficence in the channels marked out by commonly recognized ties and claims".⁵ Such a bias will be "conducive to the general happiness" for three reasons: (i) Acts prompted by natural affections are highly pleasurable and tend to sustain the natural affections which are in themselves highly pleasurable. (ii) Spontaneous acts

of genuine affection do not tend to diminish the desirable self-reliance of the recipients, "they have less tendency to weaken the springs of activity in the person benefited; and may even strengthen them by exciting other sources of energy than the egoistic".⁶ (iii) We are motivated to know how to benefit our near and dear far more than others. "On these grounds", says Sidgwick, "the Utilitarian will evidently approve of the *cultivation* of affection and the performance of affectionate services."⁷

Is Sidgwick recommending that, on every occasion, or even most occasions, on which an agent does or might feel the promptings of self-love or natural affections, he should apply the sort of reasons given above in calculating what to do? I don't think so. He does not say so, and seems to imply the contrary. He defends "spontaneous beneficence", which could hardly exist if the agent were to subject his feeling to utilitarian assessment every time it prompted him to some action. He claims that "we feel that the charm of Friendship is lost if the flow of emotion is not spontaneous and unforced."⁸ He stresses the connection between natural affection and the activities which it prompts in his claim that such activities are pleasurable because they are prompted by affection. Perhaps the clearest indication that Sidgwick didn't think that the utilitarian agent must always apply utilitarian reasoning is to be found in the following passage:

[T]he doctrine that Universal Happiness is the ultimate *standard* must not be understood to imply that Universal Benevolence is the only right or always best *motive* of action ... it is not necessary that the end which gives the criterion of rightness should always be the end at which we consciously aim: and if experience shows that the general happiness will be more satisfactorily attained if men frequently act from other motives than pure universal philanthropy, it is obvious that these other motives are reasonably to be preferred on Utilitarian principles.⁹

If, on a particular occasion, we do not consciously aim at Universal Happiness, we can hardly justify our actions *on that occasion* with reference to Universal Happiness. That is, we cannot, on that occasion, employ a conscious reasoning process which involves recognizing that certain justifications for what we are doing are utilitarian justifications and therefore to be heeded. If we take this view of Sidgwick's, together with a recognition that certain forms of motivation are valuable as *direct* motivations, we can see that a good utilitarian agent need not have one thought too many.

3. Williams and Stocker on the esotericism of consequentialism

Williams is aware of the claim that utilitarianism might demand that moral agents not employ utilitarian reasoning on all occasions. Indeed,

he regards this as a grave defect in the theory. He writes: “[I]t is reasonable to suppose that maximal total utility actually requires that few, if any, accept utilitarianism.”¹⁰ In which case, he claims, it is a matter for discussion whether “utilitarianism is unacceptable, or . . . no one ought to accept it.”¹¹ Thus we are confronted with two different claims:

FALSE: If utilitarianism is true, it would be best if at least most people didn’t believe it to be true;

DEFECT: The truth of FALSE exposes a defect in utilitarianism.

Williams claims that there are “some powerful reasons” for thinking that the prevalence of “utilitarian thinking as a personal and social phenomenon . . . could be a disaster.”¹² He suggests two such reasons and claims that Smart occasionally hints at others “at those points at which he wishes (as I have suggested, inconsistently) to keep direct utilitarianism and at the same time spirit away utilitarian calculation.”¹³

Smart does endorse the use by the act-utilitarian of ‘rules of thumb’ to save time, counteract personal bias, or even encourage spontaneity. There may be many cases in which it would be counter productive, from a utilitarian point of view, for an agent to perform a conscious utilitarian calculation prior to acting. An act-utilitarian may recognize this and see that it would in fact maximize utility in the long run if she were to habituate herself to follow certain rules of thumb, though she would be prepared to break them on those instances when she knew that obeying them would have bad consequences.

This is, of course, a long way from saying that utilitarianism demands that no-one believe it. Smart is simply saying that there are occasions when it is better that a conscious utilitarian calculation not be performed. It is important to see that this suggestion doesn’t require anybody not to accept or believe utilitarianism, even while they are performing acts from habituation to a rule or from other seemingly nonutilitarian motives. It *may* require that they don’t consciously entertain a belief in the truth of utilitarianism *on some occasions*, but this doesn’t mean that they will not believe, on those occasions, that utilitarianism is true. On other occasions, such as those involving the use of a rule of thumb to save time, it may be perfectly possible, and involve no inconsistency, for an agent to entertain a conscious belief that utilitarianism is true and at the same time act according to the rule of thumb without performing a utilitarian calculation.

Williams recognizes that Smart displays some caution in “licensing non-utilitarian states of mind”,¹⁴ but he seems to think that this is because he is committed to act-utilitarianism as a theory which gives an answer, and an unvarying one at that, to the question of what people should think about in deciding what to do. I presume that Williams thinks that Smart’s

reasons for endorsing the use of rules of thumb are, in fact, powerful reasons for thinking that the prevalence of utilitarian thinking could be a disaster. As far as I can see, however, they are just reasons for thinking that utility would not be maximized over all if utilitarian calculations were to be employed to determine every action. When Williams speaks of a "disaster", it is clear that he has something much stronger in mind.

Williams presents two more reasons why "utilitarianism's fate is to usher itself from the scene"¹⁵, the second of which seems to be a confused attempt to adapt the Prisoner's Dilemma to utilitarianism.¹⁶ I will say no more about it. The first is the familiar criticism that certain non-utilitarian qualities in people and society are prized and do, in fact, affect people's happiness. Williams has in mind, I assume, such qualities as spontaneity and a disposition to feel and act on natural affections. Williams calls these qualities non-utilitarian "both in the cast of mind that they involve and in the actions they are disposed to produce."¹⁷ However, if this is to provide a utilitarian reason for abandoning belief in utilitarianism, it must be the case that the actions produced by such qualities result in more utility than actions done in the absence of them. We may assume, then, for the purposes of the argument, that such qualities are utilitarian in the actions they are disposed to produce. They are not, however, utilitarian in the cast of mind that they involve. Williams says that it is not clear how utilitarianism co-exists with such qualities, claiming of spontaneity that "you cannot both genuinely possess this kind of quality and also reassure yourself that while it is free and creative and uncalculative, it is also acting for the best."¹⁸

Perhaps an agent cannot consciously reassure herself, while acting spontaneously, that her spontaneous action is for the best. But in a reflective moment, one can recognize that one's spontaneous actions do tend to maximize utility, and thus believe, even when acting spontaneously, that such action is for the best. It may be difficult, or even impossible, to plan to behave spontaneously. But the utilitarian agent is also a human agent, and probably disposed to spontaneous behavior anyway. This argument, then, shows only that utility will probably not be maximized if every action is motivated by the conscious desire to maximize utility; it certainly doesn't show that utility can only be maximized if no-one believes utilitarianism.

Michael Stocker, in a discussion of maximization, also seems to endorse both FALSE and DEFECT:

Maximizers hold that the absence of any attainable good is, as such, bad, and that a life that lacks such a good is therefore lacking. The basic moral psychological reason for denying this is that regret over the absence or lack is a central characterizing feature of narcissistic, grandiose, and other defective selves.¹⁹

Stocker admits that it may be argued that this shows, not that maximization is wrong, but that being a maximizer is itself not maximizing. His reply to this is that

the fact that being a maximizer may not be maximizing is not why I said maximization all too easily can be part of a defective character. The defect is shown, and indeed is constituted, by the sort of self that this involves – e.g., one that is narcissistic or too driven.²⁰

The maximizer, however, could reply to this that what is wrong with the narcissist is that he is not a maximizer. In any case, she could maintain that no argument to show the maximizing character to be defective could, by itself, show maximization to be wrong. Stocker has one more thing to say on this subject:

I see it as a severe problem for a theory if, by its own lights, it cannot be embraced and followed What I have trouble in understanding is why we should be expected to think that a theory which is so esoteric as the one now in question is worth serious consideration as our ethical theory.²¹

What is Stocker claiming here? He moves from the possibility that being a maximizer may not be maximizing to the claim that maximization cannot be embraced and followed by its own lights. The tone of his remarks (“so esoteric as the one now in question”) suggests a claim as extreme as Williams’s claim that utilitarianism entails that no one believe it. But the evidence for this claim, that being a maximizer may not be maximizing, supports, at best, the claim that utility will probably not be maximized if every action is motivated by the conscious desire to maximize utility.

Although Stocker clearly seems to embrace DEFECT, he offers no argument for it. Even if it is true that utilitarianism requires that most people don’t believe it, why should that count against the theory? The most we get from Stocker is the rather cryptic claim that such a theory would be “esoteric”. Williams offers a little more, but his arguments for DEFECT seem to center on the mistaken impression that act-utilitarianism is solely concerned with “the situation of decision”.²² This is illustrated by the following piece of reasoning:

[The direct utilitarian] tells us that the answer to the question ‘what is the right thing to do?’ is to be found in that act which has the best consequences. But it seems difficult to put that to any use in this connection, except by taking it to imply the following: that the correct question to ask, if asking what is the right thing to do, is what act will have the best consequences.²³

Thus Williams claims that the distinction between act-utilitarianism and other forms of utilitarianism is a matter of motivation. If this were the

case, then FALSE, if true, would be a problem for act-utilitarianism. But it is obviously not true that act-utilitarianism is only concerned with motivation. Traditional utilitarian theory has, as its basic goal, the maximization of utility, and if this can only be achieved by no-one believing the theory, then that is what the theory demands.

Williams predicts that utilitarianism will

retire to the totally transcendental standpoint from which all it demands is that the world should be ordered for the best, and that those dispositions and habits of thought should exist in the world which are for the best, leaving it entirely open whether those are themselves of a distinctively utilitarian kind or not."²⁴

If we allow the utilitarian to specify what she means by 'for the best', I don't think she would object to this formulation as a basic characterization of utilitarianism. Utilitarianism must leave open the question of which dispositions and habits of thought are 'for the best'. That doesn't mean that it cannot take a stand on the issue, just that it must always be prepared to change its view in the light of new psychological evidence. Not only may our knowledge of and opinions about human nature change, but human nature itself may change. I would have thought that the fact that utilitarianism can accommodate itself to such change should be seen as an asset of the theory, not a drawback.

4. Moral theories and publicity

Williams may have failed to provide good reasons for believing DEFECT, but isn't it nonetheless plausible? Isn't there something strange in the suggestion that a particular moral theory require that no-one believe it? This certainly seems to violate what Rawls calls the 'publicity condition' for a moral theory. Doesn't it seem right that if a moral theory is correct, it should be possible for everyone, or at least most people, to act as it prescribes and to believe that it is correct? Perhaps this is at least contingently true, but it is hard to see how it could be a necessary feature of the correct moral theory. Imagine a world, DC, in which, in addition to the people, there are two deities, the Donkey, who is good, and the Elephant, who is bad. Imagine, further, that the correct moral theory is the following version of divine command theory:

DONC An act is wrong iff it is forbidden by the Donkey, otherwise it is permissible.

Many people in DC believe DONC and act on it. Many other people in DC believe a false moral theory, ELEPH, which has the same structure as DONC, but which centers on the commands of the Elephant. Both the

Elephant and the Donkey regularly appear to the people of DC and issue their very different commands. There is nothing in the story so far to suggest that DONC could not be the correct moral theory. Suppose now that the people of DC are getting better and better. In fact, most people now believe DONC and very few believe ELEPH. This annoys the hell out of the Elephant, who desperately wants people to do what he says, so he works the following piece of trickery on the minds of the people. Every time the Elephant appears to the people, they believe they are seeing the Donkey, and vice versa. The Elephant, who, though evil is also more powerful than the Donkey, also fixes the Donkey so that she is not aware of the people's reversed perceptions. Now most, if not all, people who believe in DONC will actually fail to act as it prescribes. Conversely, those who believe ELEPH will usually do what DONC requires. Has this exercise of evil power by the Elephant rendered the previously true DONC false? This would be a very strange conclusion. If DONC was true before, it is still true now, it's just that now it's better if people believe ELEPH instead.

So far, I have been speaking of the "truth" of utilitarianism (or DONC), and asking whether it would matter if the truth of a moral theory required that people not believe it. However, my argument is not aimed only at moral realists, who believe that moral theories are objectively true. Suppose, for example, that I regard morality as fundamentally *chosen*, rather than discovered. Perhaps my moral commitments express something deeply rooted in my character. Wouldn't such a moral anti-realist have good reasons for embracing DEFECT?²⁵

It will be easier to see both the appeal and the failing of DEFECT, if we pause briefly to consider the role of moral theories, or at least one central aspect of their role. Both moral realists and anti-realists (of various kinds) agree that moral theories are action-guiding in the following sense: they provide reasons for acting. If my moral theory contains a prohibition on coveting my neighbor's ass, I have a reason not to covet my neighbor's ass (I'm not sure whether coveting is a kind of action, but bear with me). But if, according to my moral theory, I shouldn't even *believe* my moral theory, how is it supposed to supply me with reasons? And if it can't supply me with reasons, how can it be a moral theory? The obvious answer to this is to point out that reasons don't have to be embodied in consciously held beliefs, or even unconscious beliefs, in order to apply. The smoker who doesn't believe that smoking is bad for her has the same reason to quit as the better informed (or less self-deceived) smoker. At this point, the moral anti-realist will probably point out that the harmful effects of smoking are a matter of objective fact, whereas moral theories inhabit (according to him) an entirely different realm. The reasons supplied by moral theories are more like the reason I have for benefiting someone I care deeply about than the reason I have for quitting smoking.

If I care deeply about Smith and you don't, I have a reason for benefiting Smith that simply doesn't apply to you. But this example can be modified to illustrate how moral theories can provide reasons for acting to those who don't accept them, even given moral anti-realism. Suppose I care deeply about Smith and want her to be happy, above all else. However, I also know, from bitter past experience, that when I care deeply about someone, I become irrational, possessive, violent, and obsessive. In fact, everyone for whom I have cared deeply has suffered terribly as a result. Given that I really do want Smith to be happy, I judge that it would be better if I could get myself not to care about her at all. Perhaps I succeed in this endeavor, and no longer care about Smith. As a result, she is a lot happier than she would have been. My emotional commitment to Smith provided the reason for me to change my feelings, and continues to provide reasons for my behavior, even though such reasons are now inaccessible to me. This suggests that the assumption that moral anti-realists must be committed to DEFECT stems from an impoverished view of how reasons for behavior can operate.

Now consider how a moral anti-realist might view my example. Given the kind of being the Donkey is, and the kind of person I am, DONC is my chosen theory. My acceptance and advocacy of DONC express something deeply rooted in my character. But what if the Elephant were tricking me in the manner described above? Even though I embrace DONC, I judge that, were the Elephant to be tricking me, it would be better, according to my chosen theory, if I were to embrace ELEPH instead. Perhaps I am told that the Elephant will begin his trickery tomorrow (the trickery will, of course, include erasing my memory of being told this). There is a rigorous course of drug and behavioristic treatment, that I can undergo today. This treatment has a 95 percent chance of changing my character in such a way that I will embrace ELEPH. Given that I currently embrace DONC, I have a very good reason to submit myself to the treatment. The reason is supplied by DONC itself. I might regard it as regrettable that the Elephant's power has driven me to this, but I don't consider DONC any less appropriate as a moral theory because of it. If the treatment is successful, and I come to embrace ELEPH (and therefore act as DONC requires), there is a very clear sense in which DONC is still providing me with reasons for acting, even though I would then believe otherwise.

There are good reasons, then, for believing both that FALSE is false and that DEFECT is false. Neither does the truth (or acceptability) of utilitarianism require that most people not believe it, nor would it matter if it did.

5. *Self-deception and consequentialism*

The consequentialist treatment of non-consequentialist motivations has recently been subjected to a different line of attack by William Willcox. Willcox focuses on Peter Railton's recent elaboration of Sidgwick's utilitarian account of motivation. So I will first present a sketch of Railton's position, before examining Willcox's criticism. Railton argues that a good moral agent will have a "standing commitment to leading an objectively consequentialist life . . . *Objective consequentialism* is the view that the criterion of the rightness of an act or course of action is whether it in fact would most promote the good of those acts available to the agent."²⁶ You can espouse values which can affect decision-making in a non-consequentialist way, just so long as you are committed to removing those values which tend to produce bad results. This doesn't mean that you cannot be committed to non-consequentialist values, even though you would attempt to remove them if you discovered that they were reducing overall utility (or some broader conception of good). You should probably be committed to removing, say, racial prejudice from your value system, but could be committed to a moral aversion to chemical and biological warfare, even though on some occasions such an aversion could result in a decision which was wrong from an objective consequentialist point of view. Railton calls this view "sophisticated consequentialism".

Return to John, Jane and Mary. Let's suppose John doesn't betray Jane, even though she would never find out, and he and Mary would experience a great deal of pleasure and no pain from the infidelity. John could be a sophisticated consequentialist who has both a standing commitment to leading an objectively consequentialist life, and a non-consequentialist commitment of fidelity to Jane. If he had been able on this occasion to perform the objectively right action, he would have been less devoted to Jane. If he had been less devoted to Jane, perhaps he would have done less good in the long run.

In my example, it is not clear whether John's natural affection for his wife overrides his commitment to act for the best on this occasion, or whether he simply doesn't employ consequentialist considerations. Railton's account of sophisticated consequentialism allows for both possibilities. Consider the following passage:

. . . individuals may be more likely to act rightly if they possess certain enduring motivational patterns, character traits, or prima-facie commitments to rules in addition to whatever commitment they have to act for the best. Because such individuals *would not consider consequences in all cases*, they would miss a number of opportunities to maximize the good.²⁷

Railton is suggesting that a good sophisticated consequentialist will sometimes act directly on the promptings of certain commitments or

character traits, without consulting consequentialism. This makes room for spontaneous or affectionate behavior. But Railton is not solely concerned with motivations that are either non-moral or not opposed to moral motivations. He is not just arguing for the place of 'moral holidays' in the life of a good consequentialist agent:

Surely part of the attraction of these indirect consequentialisms is the idea that one should have certain traits of character, or commitments to persons and principles, that are sturdy enough that one would at least sometimes refuse to forsake them even when this refusal is known to conflict with making some gain – perhaps small – in total utility.²⁸

To return to the example of John, Jane and Mary, it is consistent with the description of John as a sophisticated consequentialist that he knows that he could do more good by betraying Jane with Mary. Perhaps he knows this, but doesn't consciously entertain the belief. A consequentialist can, in any case, allow for the possibility of conscious conflict between commitments, or other motivations that are justified on consequentialist grounds, and the belief that the right action is the action that produces the best results. Railton's primary focus is the charge that consequentialists must be alienated from their natural affections. His central example involves a conflict between the dictates of consequentialism and affection for a spouse. However, as the last two quotes suggest, the same approach can be extended to moral commitments to non-consequentialist principles. Railton doesn't illustrate how his approach applies to non-consequentialist moral commitments, or the rationality of regret arising from a conflict between such commitments and the dictates of consequentialism. I will do so, partly because such an account is important in its own right, and partly because the charge of self-deception leveled by Willcox against Railton's account is far more troubling in the context of non-consequentialist moral commitments than in the context of natural affections. Consider the following example.

A German army officer, Schmitt, opposed to the Nazi regime, is, at the start of 1945, offered command of a concentration camp where Jews are being killed at an ever-increasing pace. He is fairly certain that the Allies will force the German surrender within about six months, and also that he would be able to slow the rate of killings in the camp and get away with it if he were to accept the job. He also knows that if he doesn't accept the job, Kurtz, who is a fanatical Nazi, will get it and will continue to speed up the killings. Now I think there is little doubt that we would not blame Schmitt for refusing the job, indeed it might even be psychologically impossible for him to take it, but I also think that he ought to take it. In fact, the more deep-seated his aversion to killing innocent people, the stronger his reasons for taking the job. We may well feel that there would be something wrong with Schmitt if he could accept the job readily

without any soul searching, but this could be because we took such behavior as a sign of a disposition which was not opposed to the duties involved in the job. We think that, as a matter of psychological fact, it would be difficult for someone with such aversions to accept such jobs, but it would nonetheless be right for them to do so. Likewise, if Schmitt carried out his duties without any regret or remorse, we would take that as a sign of a defective character, even though we realized that he had done the objectively right thing. The fact that Schmitt himself realizes that he has done the right thing does not diminish the appropriateness of regret or remorse on his part.

Schmitt's moral aversion to killing innocent people need not be qualified in any way in order for him to qualify as a good sophisticated consequentialist agent on Railton's account. In fact, if it is qualified, he may be less good. That is, Schmitt's aversion is to killing innocent people, not to killing innocent people unless more good can be done by killing innocent people than by doing something else. It seems reasonable to suppose that most of us will be better people, from a consequentialist standpoint, if we have an unqualified aversion to killing innocent people than if we have a qualified aversion. If this supposition is true, it is, of course, only contingently true. There is no inconsistency in the notion of a moral agent who has a qualified aversion to killing innocent people and whose character is in no way defective. We may not like such a character, but our judgments are based on human nature as we know it. Given our psychological makeup, it would be difficult for someone with only a qualified aversion to killing innocent people to have sufficient moral sensitivity. It is just as well that we wouldn't like such a person, since it is almost certain that a qualified aversion to killing innocent people would be a sign of a defective moral character in anyone we were likely to meet. To have an unqualified aversion to killing innocent people is not, of course, to be unable to kill innocent people. It is, rather, to be unable to kill innocent people without experiencing something like guilt, regret or remorse.

Schmitt can have a standing commitment to leading an objectively consequentialist life and have an unqualified moral aversion to killing innocent people, even though the objectively right thing to do may, on occasion, involve killing innocent people. Schmitt may even, in a calm hour, reflect on his aversion to killing innocent people and realize that there may be situations in which this will make it difficult, or even impossible, for him to do the right thing. Yet he may judge that the costs of trying to change his attitude to killing innocent people will outweigh the benefits.

Railton has ably demonstrated how Sidgwick's consequentialist account of motivation provides a consequentialist with a framework for answering the charge that utilitarianism involves alienation from natural

affections. I have illustrated how the same approach avoids alienation from moral commitments and accommodates the view that regret or remorse can be appropriate even when one has knowingly done the right thing. The account has been challenged, though, on the grounds that it involves the sophisticated consequentialist in self-deception.

William Willcox examines Railton's central example, in which Juan, who has a commuting marriage to Linda, decides to visit her rather than give the money to charity, which would have produced more objective good. He claims that, while a consequentialist may approve of Juan's character, Juan himself is either not a consequentialist or a self-deceived one:

While a consequentialist looking at Juan as a third party might well approve of Juan's character, that character does not include the acceptance of an overriding commitment to impersonal value – even if we suppose that commitment to take a counterfactual form. Railton's example is one where a counterfactual motive would come into play if it really existed. It does not come into play.²⁹

Railton's counterfactual condition on the sophisticated consequentialist is central to his account. What makes a sophisticated consequentialist a consequentialist is that she is committed to removing those values that tend to produce bad results. Railton says of Juan that "his motivational structure meets a counterfactual condition: while he ordinarily does not do what he does simply for the sake of doing what's right, he would seek to lead a different sort of life if he did not think his were morally defensible."³⁰ Willcox has the following to say about Railton's counterfactual condition:

We must understand this counterfactual condition as being normative rather than merely psychological if the agent is to be a consequentialist rather than merely someone who acts as a consequentialist would require. Hence, the consequentialist judges that he *ought* to sacrifice his "friends" whenever doing so promotes impersonal value, and he must form an intention to act according to this judgment.³¹

A merely psychological counterfactual condition would have the same effect as a normative one in terms of behavior. An agent whose motivational structure meets a merely psychological counterfactual condition would alter his motivational structure, and therefore his behavior, under the same conditions as one who is subject to a normative condition. However, the mechanism by which he altered his motivational structure would not include a belief that consequentialism required him to alter it. Such an agent would be no more a consequentialist than someone who obeys the dictates of a god who, unbeknownst to him, is a consequentialist. Willcox is correct to claim that a sophisticated

consequentialist must be subject to a normative counterfactual condition, and therefore must explicitly endorse consequentialism, on at least those occasions when the condition comes into play. If he is also correct in his claim that the condition should come into play to block Juan from taking the extra trip to Linda, it would seem that Juan is indeed either not a consequentialist or self-deceived. But should the counterfactual condition come into play in this example? In order for the counterfactual condition to require an agent to form an intention to sacrifice his friends whenever doing so promotes impersonal value, it must be a condition on individual actions. Willcox seems to think that the sophisticated consequentialist is committed always to doing the objectively right thing. Thus the consequentialist agent is able to act directly on non-consequentialist motivations only on those occasions when they don't conflict with the promotion of impersonal value. Although affection for one's wife may often motivate one to do the objectively right thing, the cases of Juan and Linda, and John and Jane are examples of natural affection conflicting with the promotion of impersonal value. Hence, claims Willcox, the counterfactual condition will require that Juan not act on his affection on this occasion, and, presumably, that John betray Jane.

But Willcox has misunderstood the counterfactual condition. The sophisticated consequentialist is committed to leading an objectively consequentialist life, that is a life that involves, of all the possible alternative lives, the greatest amount of promotion of impersonal value. But such a life may well include performing many acts which are not objectively right. The sophisticated consequentialist is committed to remove those of her values and other motivational traits that do not tend to the overall promotion of impersonal utility. If John were to conclude that he would do more good overall if he attempted to remove or diminish his affection for Jane, then he ought to do so. It is perfectly plausible to suppose that John could recognize that he would not do more good overall if he attempted to remove his affection for Jane, and that he could also recognize that on a particular occasion his affection for Jane leads him to do something other than the objectively right thing. The same considerations apply to Juan and Linda.

Willcox suggests that Juan may be a consequentialist after all, but a self-deceived one.

Juan could be a self-deceived consequentialist with a moral blind spot where Linda's welfare is concerned. He would convince himself, probably by muttering something about having the sort of character of which a consequentialist could approve, that he does have an overriding commitment to impersonal value and that even though such a commitment would seem to require him to sacrifice Linda's welfare, the sacrifice is not really required.³²

Willcox is claiming that Juan, if he is a consequentialist, must deceive himself into believing that he really is doing the objectively right thing in

visiting Linda. He would also claim that John suffers from a similar case of self-deception. But, as we have seen, the counterfactual condition is not on actions but on characters. John doesn't have to believe that he is doing the objectively right thing. He may act without considering morality at all. If he does consider morality, he will realize that he could have done more good by sleeping with Mary. He will also believe that his decision not to do so is prompted by a part of his character that is justified on consequentialist grounds. It is this belief that is required for him to satisfy the counterfactual condition. There is nothing in Railton's example or Willcox's discussion to suggest that John (or Juan) can only acquire this belief by self-deception.

6. Self-deception and conflicting commitments

Willcox has failed to show that a sophisticated consequentialist can only avoid alienation at the expense of self-deception. However, maybe the charge can be urged from a different angle. The reason why John is not self-deceived, it might be argued, is that his case involves a conflict between a belief in consequentialism and his natural affection for Jane. John may be fully aware that he is acting on a motivation which, on this occasion, does not lead him to do the objectively right thing. But what of a conflict between belief in consequentialism and a deeply held moral conviction, such as an aversion to killing innocent people? What are we to say of a situation in which John refuses to kill an innocent person because of his moral commitment to the principle that one shouldn't kill innocent people, even though he is aware that he will do more good by killing the innocent person than by doing anything else? If his choice to spare the innocent person is motivated by a *moral* commitment to not killing innocent people, must he not believe that he is doing the morally right thing? But in that case, how can he at the same time believe that the morally right thing is what does most good and that killing the innocent person would do most good? Must he not deceive himself into believing that not killing the innocent person on this occasion is the morally right thing to do?

One possible reply to the modified charge of self-deception leveled in the previous paragraph is to claim that John doesn't even consider the consequences of killing the innocent person. What it is to have a commitment to not killing innocent people is, at least in part, to be disposed to choose a course of action that doesn't involve such conduct, without considering the consequences of this. The sort of person who has a commitment to not killing innocent people is the sort of person for whom the question of the consequences of such behavior doesn't arise.

This reply may be acceptable for a number of cases, but I think it is clear that it won't do as a general answer to the charge of self-deception. Even if John doesn't consider the consequences of killing the innocent person at the time of making his decision, he may well reflect on his choice later. He may realize that he did the wrong thing. It is reasonable to suppose that he will at least consider the consequences if another choice involving killing an innocent person should arise. More important, though, is the possibility of conflicting commitments. An admirable moral agent may have a commitment to not killing innocent people, to saving the lives of innocent people, to increasing the level of welfare of suffering people, and so on. Given the variety of different commitments that John may possess, and the obvious possibilities for conflicts between them, it is hardly plausible to maintain that having a commitment to ϕ ing involves being disposed not to consider the consequences of not ϕ ing. Some commitments may be stronger than others. An admirable moral agent may be more likely to consider the consequences of not increasing the level of welfare of suffering people than to consider the consequences of killing an innocent person. To adapt an example of Elizabeth Anscombe's, a judge may be so strongly committed to not procuring the judicial execution of innocent people that she never actually considers the consequences of doing this in any particular case. Some people may have some commitments which never lead to a conscious conflict with the commitment to produce the best results, but it seems likely that most of the commitments which a consequentialist would endorse may sometimes lead to such a conscious conflict.

Despite the possibility that John's commitment to not killing innocent people will lead him to do something that he knows does not have the best results of all his available alternatives, I do not think that the charge of self-deception is justified. Consider John's commitment to not killing innocent people. Does this include the belief that killing innocent people is always wrong, or even the belief that it is usually wrong? It is not clear that John's commitment needs to include any beliefs about what is right or wrong. It may include the belief that killing innocent people is bad, either in itself or because of its consequences, but it is not even clear that this is needed. John's commitment to Jane doesn't include any moral beliefs, so why should his commitment to not killing innocent people? John may have many moral beliefs about killing innocent people, but these need not be part of his commitment to not killing innocent people. Similarly, John may have moral beliefs about his behavior towards Jane, but these need not be part of his commitment to Jane.

John has a deep aversion to killing innocent people. If he kills an innocent person, he feels bad. Not killing innocent people is one of the many projects around which he builds his life. John also has a deep desire to watch Italian films. If he misses out on an opportunity to watch Italian

films, he feels bad. Watching Italian films is one of the many projects around which he builds his life. John's commitment to watching Italian films is not a *moral* commitment. So what makes his commitment to not killing innocent people a *moral* commitment? Part of the answer to this lies in the etiology of the commitment. John was probably taught that it was morally wrong to kill innocent people. More importantly, John universalizes his commitment to not killing innocent people and tries to get other people to share it. John's aversion is not just to John killing innocent people, but to anybody killing innocent people. If John meets someone who is not committed to not killing innocent people, he will probably try to get them to share his commitment. On the other hand, John can have a deep commitment to watching Italian films, but not care in the least whether anybody else watches them. He may not even be inclined to spread his passion for Italian films.

We have seen that a sophisticated consequentialist needn't engage in self-deception, even when she knowingly follows a commitment at the expense of producing the best outcome. Nevertheless, isn't there something strange about the state of mind of someone who acknowledges that she is morally required, on consequentialist grounds, to do *x*, but admits that she is not going to do *x*? At least, isn't it strange to describe such a person as having a standing commitment to leading an objectively consequentialist life? Recall, however, that such a commitment involves the disposition to attempt to change character traits that are overall harmful, not those that merely prevent one from maximizing utility on some occasions. At the very least, though, there seems to be a tension between the belief that *x* is morally required and a commitment to doing something other than *x*. The tension may be diminished somewhat by rejecting the maximizing requirement in consequentialism. If John does not kill an innocent person, he has not done the morally best thing, but he hasn't failed a moral requirement. Such a scalar approach to morality, which has ample independent motivation,³³ can be combined with the approach of this paper to give a more convincing answer to the problem of alienation than is found in Railton's work alone. We should not, however, expect a moral theory to remove *all* the tension between our moral beliefs and commitments. Such tension is a fact of the moral life, not a problem for an account of morality. It is part of the richness of our moral lives that an admirable character involves commitments that conflict. A good life, even the best life, may involve often doing what is less than best.³⁴

Southern Methodist University

NOTES

¹ Mill, J. S., *Utilitarianism* (1861, widely reprinted, ch. 2).

² Perhaps the world would be a far better place, if we could all be like Mother Theresa. Some may recoil in horror from the idea of so much impersonal benevolence. (By 'impersonal' I simply mean 'valuing all equally'.) The ability to form special relationships is one of the most important ingredients in a good life, a life worth living, it may be claimed. There is a lot of force to this objection, but I suspect that that is because we are trying to imagine what *our* lives would be like without the sort of special relationships which make them worth living.

³ I am grateful to an anonymous referee for this journal for suggesting this criticism.

⁴ Henry Sidgwick, *The Methods of Ethics* (7th edition) (Hackett, 1981, p. 431, *emphasis added*).

⁵ *Ibid.*, p. 433.

⁶ *Ibid.*

⁷ *Ibid.*, *emphasis added*.

⁸ *Ibid.*, p. 437.

⁹ *Ibid.*, p. 413.

¹⁰ Williams, *op. cit.*, p. 135.

¹¹ *Ibid.*

¹² *Ibid.* p. 130.

¹³ *Ibid.*

¹⁴ *Ibid.*, p. 128.

¹⁵ *Ibid.*, p. 134.

¹⁶ It's of the essence of the Prisoner's Dilemma that it involves two or more people each of whose decisions have no causal effect on the decisions of the other, while it's of the essence of Williams's example that it involves two people each of whose decisions have a lot of effect on the other's.

¹⁷ Williams, *op. cit.*, p. 131.

¹⁸ *Ibid.*

¹⁹ Michael Stocker, *Plural and Conflicting Values* (Oxford, 1990), p. 321.

²⁰ *Ibid.*, p. 322.

²¹ *Ibid.*

²² Williams, *op. cit.*, p. 125.

²³ *Ibid.*

²⁴ *Ibid.*, p. 135.

²⁵ A notable moral anti-realist, Jonathan Bennett, explicitly rejects DEFECT in his recent book *The Act Itself* (Oxford, 1995; see especially p. 22).

²⁶ Peter Railton, "Alienation, Consequentialism, and the Demands of Morality", *Philosophy and Public Affairs* Vol. 13, No. 2 (Spring 1984), reprinted in Samuel Scheffler (ed.), *Consequentialism and its Critics* (OUP 1988), pp. 93–133, at p. 114.

²⁷ *Ibid.*, p. 120, *emphasis added*.

²⁸ *Ibid.*, pp. 118 and 120.

²⁹ William H. Willcox, "Egoists, Consequentialists, and Their Friends", *Philosophy and Public Affairs* 16 (1987), pp. 73–84, at p. 81.

³⁰ Railton, *op. cit.*, pp. 111–12.

³¹ Willcox, *op. cit.*, p. 79.

³² *Ibid.*, p. 83.

³³ See Michael Slote, *Common-sense Morality and Consequentialism* (Boston: Routledge and Kegan Paul, 1985, chapter 5); Frances Howard-Snyder and Alastair Norcross, "A Consequentialist Case for Rejecting the Right", *The Journal of Philosophical Research*, 1993; Alastair Norcross, "Good and Bad Actions", *The Philosophical Review*, Jan 1997.

³⁴ I am grateful to many people for help with this paper. I presented earlier versions at Texas Tech University and at the Ethics Colloquy at SMU, at both of which events I received many helpful comments and criticisms from people whose names I have, unfortunately, forgotten. Both Michael Stocker and Jonathan Bennett supplied me with detailed comments on an early version, and, more recently, Steve Sverdlik and an anonymous referee for this journal provided me with very helpful comments on the penultimate version. Any remaining defects are entirely my own doing.