

IN DEEPEST CONSEQUENCE:
What Makes Effects Moral

Dr Peter J. King

Pembroke College
Oxford
OX1 1DW
U.K.

tel.: +1608 659174
e-mail: peter.king@philosophy.oxford.ac.uk

ABSTRACT

It is often argued that Utilitarianism is impractical because our actions have extremely complex and far-reaching sets of effects, thus making moral calculation impossible. My response to this criticism uncovers a rôle for the knowledge and intentions of the individual moral agent which goes some way towards answering the complaint that Utilitarianism makes too little allowance for such factors. I look at the notions of causation and of intentions and motives, and draw a distinction between the moral and non-moral outcomes of actions, which I dub *consequences* and *repercussions*. This leads to the conclusion that the degree to which an agent is morally responsible for the effects of her actions depends broadly upon three factors: her intentions, her knowledge, and the care she takes in relation to her situation and to the context of her actions.

KEY WORDS

action, causation, consequence, intention, Mill, motive, repercussion, utilitarianism.

IN DEEPEST CONSEQUENCE: *What Makes Effects Moral*

Two criticisms often levelled at Utilitarianism, and especially at J.S. Mill's version of it, are: first, that our actions frequently if not always have extremely complex and far-reaching sets of effects, making moral calculation impossible, and secondly, that the ideal Utilitarian agent is a cold, inhuman calculating machine. In examining and developing Mill's own responses to these criticisms I hope to uncover a connection that will support and strengthen them. To do this, I shall be led into a brief investigation of Mill's account of causation, in an attempt to discover how it relates to his view of the morally relevant consequences of actions, but more importantly I shall investigate the notions of intention and motive, and their rôles in the analysis of actions and agents.

i. Complex effects & chronic criticisms

The first sort of criticism can be found in the form of three similar but different types of claim. What they have in common is the view that Utilitarianism demands too much of the agent, in that the Utilitarian calculation is too complex to be carried out on most if not all occasions. This approach is taken, not only by those who openly reject Utilitarianism, but also by those who argue for Rule Utilitarianism.¹

Of the three claims, the first is the most familiar; it is that our actions change the world in various ways, these changes producing further changes, etc., so that in trying to calculate the effects of our actions we'll find ourselves locked into a never-ending calculation; I'll call this the *diachronic criticism*. Robert Spaemann provides a clear case of it:

Utilitarianism fails first of all because of the complexity and unpredictability of the long-term consequences of our actions. If we did have to take into account all the consequences of our actions, we would spend so much time trying to work these out that we would never get round to doing anything at all.²

He offers the following example:

The lowering of infant mortality rates in poorer countries often has catastrophic consequences in the long term, but then these in turn lead to further pressure to improve overall living conditions, though it is not clear whether or not this is possible. Who can judge what, finally, will turn out to have been the most important consideration? No one would do anything at all if they had, beforehand, to take all that into account.³

The second claim, which in practice is difficult to disentangle from the first, is that the world is a tremendously complex place, so that if we have to take into account the effects of our actions on everybody else we'll find ourselves faced with an impossible calculation; I'll call this the *synchronic criticism*. The emphasis here is not on the ramification of effects over time, but on the present complexity of the world within which we act. The difference between the diachronic and the synchronic criticisms isn't that one concerns future effects and the other doesn't – most if not all effects will occur *after* our actions – but that one concerns the potential infinity of ramifying effects while the other concerns the potentially infinite complexity of effects. (The synchronic criticism is related to complaints that the Utilitarian calculation ignores principles of distributive justice.)

The third and least commonly expressed claim is that, on any given occasion of action-choice, there's an indefinitely large number of possible actions from which to choose; how can

I consider them all? This claim differs from the other two in that it makes no explicit mention of the consequences of my action. I'll call this the *spoilt-for-choice criticism*. So far as I'm aware it has never been expressed in this explicit, pure form, but the other two criticisms can be seen as special cases of it.

All three claims are designed to show that Utilitarianism asks too much of us in a practical sense; not only is it unreasonable to expect us to perform the Utilitarian calculation before we act, but it is actually impossible for us so to do.⁴ My main concern here is with the synchronic and diachronic criticisms, but much of what I'll have to say will apply to the pure spoilt-for-choice criticism too.

Now, Mill of course offered his own defence against the diachronic and synchronic criticisms, and implied a defence against the spoilt-for-choice criticism. In *Utilitarianism* itself he tackled the notion that we must think of humanity as a whole when we decide how to act.⁵ His response to the synchronic claim was simple and practical:

The great majority of good actions are intended, not for the benefit of the world, but for that of individuals, of which the good of the world is made up; and the thoughts of the most virtuous man need not on these occasions travel beyond the particular persons concerned, except so far as is necessary to assure himself that in benefiting them he is not violating [...] the legitimate and authorised expectations — of any one else.⁶

I shall come back to this point later. Mill's response to the diachronic claim had as its target Adam Sedgwick, who had written:

Feeble as man may be, he forms a link in a chain of moral causes, ascending to the throne of God; and trifling as his individual acts may seem, he tries in vain

to follow out their consequences as they go down into the countless ages of coming time. [...] we are, at least, certain, that man has not foreknowledge to trace the consequences of a single action of his own; and hence that utility [...] is, as a test of right and wrong, unfitted to his understanding, and therefore worthless in its application.⁷

This is clearly the sort of criticism that Spaemann had in mind in the quotation above (page 1). In reply to Sedgwick Mill wrote, with more than a touch of asperity:

Mr Sedgwick appears to be one of that numerous class who never take the trouble to set before themselves fairly any opinion which they have an aversion to. Who ever said that it was necessary to foresee all the consequences of each individual action, “as they go down into the countless ages of coming time”? Some of the consequences of an action are accidental; others are its natural result, according to the known laws of the universe. The former, for the most part, cannot be foreseen; but the whole course of human life is founded upon the fact that the latter can.⁸

In fact I shall be arguing that matters are more complex than this, and that we must go beyond the question of which effects are foreseen or foreseeable. Nevertheless, Mill clearly pre-empted the sort of complaint made by Spaemann and others, so why did they ignore his defence?

If we dismiss the possibility that they share the distressing fault that Mill attributes to Sedgwick, one reason might be that Mill’s distinction between the accidental and the natural results of actions, at least if taken metaphysically, is unclear. Indeed, it seems to go against his own acceptance of causal determinism, according to which all effects are straightforwardly

and equally determined by series of preceding causes; that surely doesn't allow for any distinction between accidental and natural effects.

Turning for help to Mill's account of free will and determinism in his *System of Logic*, one passage might at first seem to meet the case:

There are physical sequences which we call necessary, as death for want of food or air; there are others which, though as much cases of causation as the former, are not said to be necessary, as death from poison, which an antidote, or the use of the stomach-pump, will sometimes avert.⁹

Even if this is found to be at all useful, however, it seems to be intended as a distinction merely between effects that follow immediately upon their causes and effects that take time during which the causal chain can be diverted.¹⁰ We still don't have the required distinction between accidental and natural effects.

The answer to this would seem to be simply that Mill's defence can and should be interpreted epistemologically rather than metaphysically, and thus that it is consistent with his deterministic views. That is, if we take him to have been arguing that in some cases we are certain of the causal laws while in other cases we're not (perhaps because the causal factors involved are too complex), his distinction between accidental and natural effects is compatible with his acceptance of causal determinism. And indeed, as we have seen, Mill defines the natural result of an action in terms of the *known* laws of the universe.¹¹

ii. Agents & actions

The second problem with which I am concerned is the rôle allotted by Mill to the agent's intentions. Many commentators have felt that Mill is guilty of allowing too little room for the

agent, thus giving Utilitarianism a cold and inhuman heart inappropriate to a moral theory. In particular, the agent's intentions, goals, projects, etc., are thought to be missing from Mill's account. Sometimes the claim is that Utilitarianism threatens the agent's integrity, usually illustrated by examples in which non-Utilitarians are made to reason and act according to Utilitarian principles. It seems to me that little is needed to defend Utilitarianism against such criticism. My current concern, however, is with claims that Utilitarianism focuses on the calculation of utility, turning the human moral agent into a sort of fleshy Difference Engine.

Mill addresses this sort of criticism in chapter 2 of *Utilitarianism*:

It is often affirmed that utilitarianism renders men cold and unsympathising; that it chills their moral feelings towards individuals; that it makes them regard only the dry and hard consideration of the consequences of actions, not taking into their moral estimate the qualities from which those actions emanate¹²

His answer is that some Utilitarians might well fall into this category, but that the same can be said of those who hold to other moral theories. As he says :

there is nothing in the utilitarian theory inconsistent with the fact that there are other things which interest us in persons besides the rightness and wrongness of their actions.¹³

(He goes on to admit, though, that "in the long run the best proof of a good character is good actions".¹⁴)

What things do interest us in persons? Well, in the moral arena at least, what are most important are their *intentions* and *motives*. Before going any further, therefore, it might be as well to clarify the distinction between these two notions, for it lies at the heart of much of what I want to say in this paper, and hasn't always been acknowledged or taken account of

by Mill's critics.

Mill's rather brief account of the distinction comes in the course of a famous footnote in chapter 2 of *Utilitarianism* (henceforward 'FN'), where he says that the intention is "what the agent *wills to do*", while the motive is "the feeling which makes him will so to do".¹⁵ This is a start, but more explanation is needed. Roger Crisp cites Mill's 1869 work, *James Mill's Analysis of the Phenomena of the Human Mind*, and takes from it the notion that by 'intention' Mill means "the foresight of consequences".¹⁶ Now, it's at first difficult to see how this simple definition could work; if I fall off a cliff, I am perfectly capable of foreseeing the consequences, but my fall surely isn't thereby intentional. That the intention *includes* some foresight of consequences, however, is undeniable, and sits well with a comment that Mill makes in FN when he argues that the distinction between motive and intention has been emphasised by utilitarian thinkers, "and Bentham pre-eminently".¹⁷ This is significant because Bentham distinguishes between what he calls *direct* and *oblique* intentions: a consequence of an action is directly intended when it was a link in the chain of causes that led the agent to act, and obliquely intended when, though the agent foresaw it, it didn't form a link in that chain.¹⁸

Although Mill appeals to Bentham's discussion, can we be sure that he accepted this bipartite account of intention? Michael Ridge argues that we can, appealing to the passage cited by Crisp:

Intention, when we are said to intend the consequences of our actions, means the foresight, or expectation of those consequences, which is a totally different thing from desiring them. The particular consequences in question, though foreseen, may be disagreeable to us: the act may be done for the sake of other consequences. Intention and motive are two very different things.¹⁹

It might at first seem as though Mill is simply equating intention with Bentham's oblique intention, but in fact he means to combine direct and oblique intention into one notion. Foresight is involved in both cases: in direct intention, the foresight plays a causal rôle, in oblique intention it doesn't. In other words: an agent *a* intends an outcome *x* iff: either *a* foresees and wills that *x* be a consequence of her action, or *a* foresees and is willing that *x* be a consequence of her action. This two-clause definition will do for the moment, but I shall argue below that matters are more complex still

Mill writes that "[t]he morality of the action depends entirely upon the intention".²⁰ This is an odd claim, given his usual insistence upon the exclusive rôle of the consequences of an action in determining its moral rightness or wrongness. Ridge suggests that what Mill has in mind is this:

Before we determine whether a given action is right or wrong, we need to know just what action is in question. Hence, determining the rightness of a given action requires having some way of individuating actions."²¹

In fact, though, we need two sorts of individuation: there's Ridge's notion of the individuation of one action from another, but we also need to individuate actions from non-actions. In the latter case, there's surely little question but that actions are individuated in terms of intentions. What distinguishes my falling of a cliff from my jumping off it is in large part the fact that, in the latter case but not the former, the event occurred because I intended it — my mind was directed towards the event in the right way, and brought about the fact of my leaving the edge of the cliff. And before we determine whether a given event is morally right or wrong, we need to know whether an action is in question; when we hear that a fire in a night club has killed hundreds of people, we judge that it was a dreadful event, but we only

dub it *morally* bad if we believe that some intention (whether direct or oblique, in Bentham's terms) was involved in a relevant way.

In the former case, however, there's more of a question; it's not at all obvious that we're committed to individuating between actions on the basis of intentions. Ridge offers three other choices:

We could instead take a more coarse-grained approach, and individuate actions in terms of bodily movements. [...] Or we could take a more fine-grained approach, individuating actions in terms of external bodily movements, intentions, and motives. [...] Or we could individuate actions in terms of external bodily movements and the feelings that give rise to them.²²

But, first, why should choose just one way? And, secondly, are we restricted to one of Ridge's options (or to the simple use of intentions)? There are, after all, very different reasons for us to want to distinguish one action from another, and it's very likely that those different reasons will call for different approaches. For example, when we're trying to decide how best to act, we're concerned primarily if not wholly with the *consequences* of different actions; we could hardly individuate actions on the basis of intentions, because the intention we form just is the choice we finally make. On the other hand, when we're judging other people in terms of their actions, we're concerned with their intentions and motives, and we distinguish between their actions accordingly.

Mill is clear that a difference in intention gives a difference in the morality of an action, even when the consequences of the action are the same. His reason, I must confess, doesn't seem to me to be very strong: it's that the agent's intention is bound up with the nature of the action itself, serving to distinguish a complex action from a series of simple actions. Thus,

in the first of the examples offered by J. Llewellyn Davies (to whom Mill is responding in FN), we are presented with a tyrant who saves his victim from drowning in order to torture him further; Mill's response is that the act of saving the victim is "only the necessary first step of an act far more atrocious than leaving him to drown would have been."²³ That is, the tyrant's intention acts to bind together the act of saving the man and the act of torturing him into one complex action, which must be judged as a whole. This allows Mill to deny that the motive plays any part in the nature of the action or in its morality, though it will certainly be relevant to our moral assessment of the agent:

utilitarian moralists have gone beyond almost all others in affirming that the motive has nothing to do with the morality of the action, though much with the worth of the agent. He who saves a fellow creature from drowning does what is morally right, whether his motive be duty, or the hope of being paid for his trouble: he who betrays the friend that trusts him, is guilty of a crime, even if his object be to serve another friend to whom he is under greater obligations.²⁴

I find it difficult to see the significant difference between these two examples:

- A. m saves n from drowning in order to torture him;
- B: m saves n from drowning in order to receive payment.

If m 's intention in A. serves to combine saving n and torturing n into one (wrong) action: {saving n in order to torture him}, why does it not do the same in B.: {saving n in order to gain reward}? It's true that the former is worse than the latter; in A. the second component action is bad in itself, whereas in B. it's not. Nevertheless, the combination in each case should surely modify the moral judgement that we make.

We can dissent from Mill's approach in two ways: we can agree that certain intentions can bind simple actions together into one complex action but argue that the same is true of the motive, or we can deny the possibility in both cases. My own preference is for the latter; what reasons might there be for preferring the former? Anne Maclean tackles this question using an example taken from James Rachels. Jack and Jill both visit their grandmother, and both comfort her and cheer her up for an afternoon; Jack does it out of love, Jill out of desire for an inheritance. Both have the same intention: to comfort their grandmother; they differ in the motives behind this intention. Maclean attempts to use this example to show that the distinction between the moral rôles of intentions and motives is unfounded, and thus that Mill's clear separation of the moral assessment of agents from that of their actions is untenable.

Maclean has three lines of attack on the Millian distinction, each presenting ways in which the differing motives lead us to describe Jack's and Jill's actions differently. First, she takes issue with the claim that both Jack and Jill should be described as having spent the afternoon comforting and cheering up their grandmother:

Is it so certain that that is what *Jill* did? It might well be argued that given her ulterior motive she should more properly be described as having spent an afternoon *pretending* to comfort her grandmother.²⁵

In the absence of an actual argument for this claim, we must surely simply reject it. Jill doubtless pretended to be fond of her grandmother, pretended to enjoy being with her, and so on — but her grandmother was actually comforted. If Jill didn't do that, who did? The grandmother was comforted by what Jill said and did, and Jill said and did it with the intention of providing comfort, whatever motives underlay that intention. (Moreover, Maclean's analysis

rules out the possibility of saying that someone could have an ulterior motive in comforting someone; on her account, one can only have an ulterior motive in *pretending* to comfort someone.)

Maclean wisely doesn't insist on this point, however. Her second line of attack appeals to a different set of ways of describing Jack's and Jill's actions:

we can say, for example, that Jack showed his grandmother the respect that was due to her and Jill did not do that. Or we can say, for example, that Jill manipulated the old lady for her own purposes, and Jack did not do that. In saying these things, we are talking about what Jack and Jill *did* (or *did not do*); we are describing their *actions*.²⁶

This is a more extreme version of the previous point; now we're being asked, not merely to *add* motive-words like 'pretending' to our description of the action, but to *substitute* them for that description. Imagine that you're asked what Jill did yesterday afternoon, and you reply: 'she manipulated her grandmother for her own purposes.' At the very least your interlocutor would regard this as insufficiently precise; perhaps she'd protest: 'No, I want to know what she did, not (your opinion concerning) her motives.' That's not to say that such an answer mightn't sometimes be an acceptable response to the question: 'What did Jill do?' — but this is because such a question can in fact be a request for a motive-account of an action.

Maclean's third line of attack involves a further stage in this direct mentioning of the motives of the two agents in the descriptions of their actions:

We can say of Jack's actions, but not of Jill's, that they were acts of love, compassion, and concern; it is not simply that *Jack* is loving, compassionate, and concerned but *Jill* is not. Similarly, we can say of Jill's actions, but not of

Jack's, that they were acts of hypocrisy and greed; it is not simply that *Jill* is hypocritical and greedy but Jack is not.²⁷

But this isn't an argument against the distinction between the morality of actions and of agents; rather, again, it's simply a rejection of that distinction. Moreover, it's clear that the relation between motive and action isn't external; as we've already seen, an action is an action precisely because it's performed by an agent, otherwise it would be a mere event. On the other hand, the task of turning a mere event into an action was to be performed by the *intention*; why should we need to bring in the motive here?

The picture we're considering is this: {Motive ➡ Intention ➡ Action}. The motive leads to the intention, which leads to the action. In Jill's case, her motive is to secure her inheritance; this motive leads to her intention to visit and comfort her grandmother; this intention leads to the actions of visiting and comforting her grandmother. The Millian claim, however, is that whereas the motive is relevant to our judgement of the moral status of the intention, it is irrelevant to our judgement of the moral status of the action, to which only the intention is relevant.

Now it's true that, if Jill hadn't wanted to secure her inheritance, she (probably, in the absence of other relevant motives) wouldn't have visited and comforted her grandmother; so her motive led to her action. However, there's an asymmetry here: Jill's intention was to visit and comfort her grandmother, which resulted in those very actions, *and which could have resulted in no others*. Either her intention was successful or not, either it led to the relevant actions or it didn't, but only a different intention would have led to a different action.²⁸ Her motive, on the other hand, was to secure her inheritance, and that could have been achieved in many different ways, via many different intention/action pairs. Thus, at the very least, the

relation between motive and action (and between motive and intention) is a good deal weaker than that between intention and action.

Maclean wants to deny that it's possible to perform the same action with different intentions, or to have the same intention or perform the same action from different motives. I don't think that she's demonstrated the latter point satisfactorily. However, some time ago, on page 7, I said that two questions remained. The first concerned the rôle of motives; the second returns to the rôle of intentions. Grant, if only for the sake of argument, that the agent's motive can be excluded from the moral assessment of her action. Now, the only rôle so far allowed for the agent's intention is that it's what makes her action an *action* rather than merely an *event*; this surely detaches the specific action from the specific intention of that specific agent. That is, it's not the content of the intention that counts, but the mere fact of there being an intention. Perhaps it's this that worries some of those who insist on a moral rôle for motives: if Utilitarians deny that either intentions or motives affect the morality of the action, don't they justify the claim that Utilitarianism is cold and inhuman?

Well, clearly not. After all, much of the discussion so far has concerned the claim that the rôle of intentions and motives covers not only what we say about people but also what we say about actions. Take away the actions, and the people remain. That's central to Mill's position, of course; in what follows, I try to fill in more detail of the rôle of intentions and motives, in order to show just how important it is.

iii. Causes, consequences, & counterfactuals

What might rather roughly be called the *descriptive* part of Utilitarianism tells us that the ultimate good consists in the maximum possible happiness of all sentient beings, but the *prescrip-*

tive part cannot insist that we bring about this desirable state (given that ‘ought’ implies ‘can’), only that we act in such a way as to increase the aggregate happiness of those who are affected by our actions. If this is to be a genuine prescription, rather than the expression of a pious and impossible hope, then it must be possible for us to calculate the effects of our actions. And if it is to be a *moral* prescription, then it must refer to the *moral* effects of our actions, which I shall call *consequences*, as opposed to *non-moral* effects, which I shall call *repercussions*.

I shall argue that these two conditions – the possibility of the Utilitarian calculation and the moral status of the relevant effects – are intimately linked, and that the manner of their linking depends upon certain crucial facts about the agent. Before I draw the crucial distinction between moral and non-moral effects, between consequences and repercussions, I need to set a basic minimum requirement, a necessary condition, for an effect having moral significance.

I start by attempting to devise such a condition based upon Hume’s account of causation. We can adapt from Robert Nozick’s account of knowledge²⁹ a pair of subjunctive conditions for an event’s being an effect of a certain action. An event *E* is a morally relevant effect, a consequence, of an action *A* only if:

C₁ if *A* hadn’t occurred, then *E* wouldn’t have occurred;

and

C₂ if *A* had occurred, then *E* would have occurred.

Together the two conditions are meant to ensure that the effect *tracks*, to borrow Nozick’s term, the cause — that the conjunction of *A* and *E* isn’t merely chance. I don’t want *A* to be irrelevant to the occurrence of *E*, nor do I want *A*’s connection with the occurrence of *E* to

depend upon accidental and morally irrelevant factors. That is, I don't want the agent's moral responsibility for the effects of her actions to be a matter of some sort of moral luck; if an agent attempts to produce good, and produces harm through no fault of her own, she should not be judged immoral (and, of course, if she attempts to produce harm and accidentally produces good, she *should* be judged immoral).

In this I'm attempting to bring the rôle of moral luck in Utilitarianism closer to its rôle in 'common-sense morality'. As Michael Slote puts it, although our moral intuitions make room for moral luck, the idea of it

affronts our common-sense moral intuitions. [...] Of course, consequentialism also involves single judgments which, taken together, entail the possibility of moral luck. [...]

But he goes on quickly to point out the difference between the two:

However, since consequentialism in other respects and in other areas knowingly offends against common-sense intuitions, since consequentialism does not seek to measure its own success or validity by a common-sense moral yardstick, the consequentialist may not be as bothered by his commitment to moral luck as the defender of common-sense morality ought to be.³⁰

Now, although this is true enough (and the consequentialist is surely right not to take conflict with other moral systems – whether labelled 'common-sense' or not – as being decisive), it can't be denied that consequentialism of any sort will be more likely to find wider acceptance if it reduces as far as possible its conflict with our current moral intuitions.

To return to the two conditions, C_1 seems the more straightforward of the two, but closer inspection reveals difficulties. When Nozick introduces a similar condition for know-

ledge, he points out that it differs from a straightforward causal condition in that:

the causal condition will be satisfied in cases of causal overdetermination, where either two sufficient causes of the effect actually operate, or a back-up cause (of the same effect) would operate if the first one didn't.³¹

C_1 would not be satisfied in such cases, and this is surely a problem. For example, a concentration-camp guard G is accused of torturing and killing (under orders) a camp-inmate at time t ; here A is the guard's actions and E is the suffering and death of the inmate. Now, in such cases we don't normally accept as a defence:

D 'If I hadn't done it, someone else would have.'

We want to respond: 'Ah, but someone else *didn't* do it, *you* did.' Is it possible to accommodate this intuition in a consequentialist moral theory? Two possible approaches involve relativising E either to a certain agent or to a certain time. Thus, one might define E as: 'torture and death *at the hands of G*'. That would have the effect of showing D to be untrue rather than morally irrelevant, but at the cost of a certain artificiality (to say the least). However, if we weaken the definition of E to a less artificial: 'torture and death *at time t*', then we fail either to falsify D or to bring out its irrelevance. Besides, both attempts smack of the *ad hoc* in the same way that Maclean's redescriptions did.

A standard consequentialist approach is to appeal to matters relating to the agent and to society as a whole. Thus we might say in this case that E includes the weakening of G 's natural reluctance to torture and kill people, as well as the effects on other people of seeing an ordinarily upright and moral person performing such actions. However G could argue that, as it had already been widely known that he was a sadist, it was better that he should torture and kill a prisoner than that some less hardened guard do it; A couldn't deprave G further, and

it wouldn't weaken the faith of other people in his moral uprightness, because they'd know that he didn't have any.³² This might at first strike us as having a certain persuasiveness; I think that we are more inclined to be shocked by a normally moral person performing an act of gratuitous violence than by a normally immoral person performing the same act. The difference in response in the two cases isn't, though, emotional in the right way; that is, in the former case the increase in our shock is a matter of surprise rather than of outrage.

A brief look at Mill's approach to the notion of causation will be of help here, I think. There are two relevant ways in which Mill extends the classic Humean notion of cause.³³ First, he corrects Hume's failure to allow for the causal rôles of *persistent states* and of the non-occurrence of events or the inaction of agents. Hart and Honoré offer various examples:

diverse ranges of singular causal statements such as 'The icy condition of the road was the cause of the accident', 'The signalman's failure to pull the lever was the cause of the accident', 'Lack of rain was the cause of the failure of the crop'.³⁴

One immediate effect of taking into account the effects of an agent's inaction as well as of her action is that the guard, *G*, can no longer offer the defence that he was known to be an immoral and vicious person, and that his action could therefore neither deprave him nor weaken other people's faith in his moral uprightness. In fact, had *G* disobeyed his orders and refused to torture and kill the inmate, the *positive* effect on his own character and on the witnesses to his refusal would have been immense, so that his failure to disobey is correspondingly more culpable.

Secondly, Mill argues that causes and effects are seldom if ever single events. The causal relation

is usually between a consequent and the sum of several antecedents, the concurrence of all of them being requisite to produce, that is to be certain of being followed by the consequent.³⁵

Mill argues that, “philosophically speaking”, the cause is the whole set of antecedents — that we have no right to pick on one of them as being *the* cause. There are two possibilities here: we might say that causes are complex only in so far as effects are complex, or we might make the stronger claim that even simple effects can have complex causes. I’m inclined to accept the latter possibility, but for my present purposes either will do; when we’re dealing with the question of moral consequences, we are inevitably dealing with complex events, because we’re dealing with people. Moreover, when talk of causes involves *reasons*, matters become even more complex; if we say that Mary’s attempted suicide was caused by examination failure, we’re really saying that the failure was Mary’s reason for acting — and this is to appeal to a certain sort of causal network, which must allow for Mary’s having precisely that reason (part of the causal network) yet *not* attempting suicide.

iv. The Hippocratic paradox

Mill’s acceptance that the failure to act can be a causal element could lead one to wonder if other changes to the simple Humean account of cause might not be in order. For example, Jonathan Bennett has noted an apparent asymmetry between events or actions that hasten and those that delay some outcome; we are normally inclined to treat the former but not the latter as being part of the cause of the outcome. That is, if an event e enters a causal chain $c_1 \dots c_n$ eventuating in result r , and *hastens* the occurrence of r , we are inclined to say that e was a cause of r ; if on the other hand e *delays* the occurrence of r , we are inclined to say that e is not

a cause of *r*:

although you cannot cause a fire by delaying something's burning, you can cause a fire by hastening something's burning³⁶

Bennett introduces this asymmetry in the context of an argument against the sort of counterfactual account of causation that I suggested above for morally relevant causation. His claim is that such a theory is inconsistent with the asymmetry. However, later (in Bennett [1988]) he revises his position, arguing that there is in fact no asymmetry, that delayers are in fact "lengthily, remotely, unsaliently" causal elements in the chain leading to the outcome; their lack of salience, in particular, is what makes us reluctant to *say* that they're causes.³⁷

Much depends upon the notions of lengthiness, remoteness, and salience, but it seems that if Bennett's position is accepted it could turn out to be disastrous for the Utilitarian. It would mean, for example, that if I advise Mary against suicide, and my intervention serves to delay but not to prevent her action, then her death is to some degree a consequence of my advice. Penelope Mackie calls this "the Hippocratic Paradox",³⁸ because it means that a doctor who saves her patient's life causes that patient's later death.

The problem with Bennett's appeal to the lengthiness and remoteness³⁹ of a delayer's causal rôle is that we have no reason to accept that those qualities apply in all cases. If my advice prevents Mary from committing suicide by only the few minutes that it takes her to reconsider and reject my arguments, then that advice is at least no more 'lengthy' or remote from her death than its other, more obviously causal antecedents. If this is right, then the weight falls squarely upon the notion of *salience*. The trouble is that the claim that the causal rôle of delayers lacks salience seems to boil down to little more than the complaint that there is something unacceptable about including delayers as causes at all. Unfortunately, Bennett

doesn't discuss these issues at any length, so it's difficult to be sure what he had in mind. However, Mackie discusses the problem more fully, and it is to her discussion, therefore, that I now turn.⁴⁰

Mackie offers as examples three pairs of cases, of which I shall consider the first two:

A₁ Because of rain in April, lightning in May failed to start a forest fire, though lightning in June succeeded.

A₂ After the events of A₁, lightning in July failed to start a forest fire, there being no forest left to burn.

B₁ On Monday a doctor acts to prevent Jones from dying of a heart attack, but on Thursday Jones dies from a second heart attack.

B₂ Smith suffers unawares from a heart condition, and plans to run a marathon on Saturday, with the certainty of fatal consequences. On the Wednesday before, however, she dies of a heart attack brought on by an argument.

A₁ and B₁ are supposed to be cases in which *e* delayed the occurrence of *r*, so that it would be unnatural to say that *e* was a cause of *r*. A₂ and B₂ are supposed to be cases in which *e* hastened the occurrence of *r*, so that it would be natural to say that *e* was a cause of *r*.⁴¹

Now, in each of the four examples it is pretty easy to decide what counts as *e*: in A₁ it is the April rains, in A₂ the June lightning, in B₁ the doctor's actions, and in B₂ it is the argument. The problem arises, it seems to me, when we come to identifying *r*. Mackie's clear intention is that in A₁ and A₂ it is a forest fire, in B₁ and B₂ it is a heart attack. But let us look at one of the cases in which *e* is supposed to be obviously part of the cause of *r*. In A₂ we have the following sequence of events:

a. April rains	b. May lightning	c. June lightning	d. June fire	e. July lightning
----------------	------------------	-------------------	--------------	-------------------

Mackie’s verdict is: “the June lightning *hastened* the burning of the forest fire from July to June” (pp 483–4). But this surely depends upon two things: first, the simple identifying of the June lightning as *e* and the June fire as *r*; secondly, the counterfactual supposition: ‘If there had been no June lightning then the July lightning would have caused a fire’ (in fact there is an alternative: ‘If there had been rain in May then the July lightning would have caused a fire’, but I think that I can safely ignore this). That would give us the following sequence of events:

a. April rains	b. May lightning	c. No June lightning	d. July lightning	e. July fire
----------------	------------------	----------------------	-------------------	--------------

Finally, we have to keep the identification of the June lightning as *e*, but have to identify the July fire as *r*. But the *June* fire was *r*; in order for Mackie’s claim to make sense, then, we have to assume that the June fire just is the July fire. Are there any grounds for doing so? Surely not. The fire in July would have been different from the fire in June in many ways, not least in terms of the flora and fauna destroyed. Moreover, the fire in June might have been successfully tackled by fire-fighters, leaving enough forest for the fire in July to occur; then there would have been a June fire and a July fire.

Thus in A_2 we have a simple sequence of events: the April rains prevent the May lightning from causing a forest fire; the June lightning causes a forest fire; the July lightning causes no fire. There is no question of hastening or delaying, because there is no reidentifiable event to be hastened or delayed.

We can go back to A_1 and say the same sort of thing. The April rains did not delay the May fire ’til June, they prevented the May fire. This might seem to have entered an extended

causal chain leading to the June fire, but we certainly do not (and I think that I've shown why we ought not to) count every element of a causal chain as being a causal element.

I shall now go on to the distinction that I hope will help to resolve some, at least, of the problems with which I am concerned here.

v. Consequences, repercussions, & foreseeability

The distinction that I want to draw between consequences and repercussions is best approached by means of examples. My first involves the following three cases:

- A. Sidney stops Mabel, a stranger, on the street to ask for directions to the Post Office; as a result, she misses her train, fails to attend a meeting, and a family is unfairly evicted.
- B. Sidney distracts Mabel, his lover, in order to make her miss her train (because he wants to spend the day with her); as a result, she fails to attend a meeting, and the family is unfairly evicted.
- C. Sidney, a landlord, causes Mabel, a housing official, to miss her train, in order to make her miss her meeting, so that he'll be able unfairly to evict the family.

If we're simply considering the morality of the action in each case, then there seems to be little or no difference between A, B, and C; in each case the same result follows from the action, and follows in the same way. If we're considering Sidney's moral status, however, matters are more complicated. In case A, the family's eviction is a mere repercussion of Sidney's action, not a moral consequence, and we should judge him to be morally innocent. In case C, on the other hand, the family's eviction is a moral consequence of Sidney's action,

and we should judge him morally blameworthy. The complication arises with regard to case B. Following the account of intention and motive in section iii, Sidney will be morally accountable if he foresaw the eviction of the family, even though it wasn't part of his reason for acting. That is, he's guilty if he foresaw that the eviction would result from his action, innocent if he didn't.

Matters are more complex than that, though. The question now arises: must Sidney *actually* foresee that his action will have a certain effect in order for us to account it a morally relevant consequence? It seems to me that Sidney's responsibilities might be more extensive, depending upon the exact nature and context of his action.

Let's elect Mabel to the House of Commons, and set up two more examples:

D. Sidney is chatting to a friend in a noisy pub, and refers in passing to a member of a minority community. By chance he's overheard by another member of that community, his remark is misunderstood, angry words turn to blows, and the result escalates into a riot, involving violence and injuries.

E. Mabel stands up in the House and makes a speech in which she refers in passing to a member of a minority community. Her remark is misunderstood by the community in question, and the result is an escalation of disquiet, culminating in violence and injuries.

In case D there's certainly no direct intention to cause a riot; that result had nothing to do with the reasons for Sidney's words. Let's also assume that Sidney didn't *foresee* the riot — it simply didn't occur to him that he'd be overheard — and certainly not by a member of just the wrong community — nor that his words would be misunderstood, nor that such a misunderstanding would lead to a riot. Sidney, then, is morally innocent. In Mabel's case,

however, there's more to be said.

There are various possibilities. It might be that there was absolutely no way in which Mabel could have foreseen the results of her speech; in that case, she surely can't be held to be morally responsible for the violence — it's not a consequence of her actions, merely a repercussion. At the other end of the moral spectrum, Mabel might, for some political or personal purpose, have intended her speech to spark off the tension and violence; in that case, she's clearly morally responsible for the violence — it's undeniably a consequence of her actions.

Between these two extremes things become interesting. The mere fact that Mabel didn't intend the effects of her speech wouldn't be enough to absolve her of moral responsibility; if she foresaw those effects but went ahead anyway, she's surely equally responsible. Again, though, the mere fact that she didn't foresee the effects wouldn't be enough to absolve her. We have to ask if she *should* have foreseen those effects, if she did enough to ensure that her remarks didn't have unpleasant effects. This 'should' is only partly or indirectly moral; the degree to which we expect her to take care will depend partly on the immediate context of the speech (for example: she's an M.P., the speech is public), partly on the wider context (for example: the racial or religious situation is particularly sensitive at the moment). As Mill points out:

the occasions on which any person (except one in a thousand) has it in his power to [multiply happiness] on an extended scale, in other words to be a public benefactor, are but exceptional; and on these occasions alone is he called on to consider public utility; in every other case, private utility, the interest or happiness of some few persons, is all he has to attend to.⁴²

And the same applies to the power to multiply unhappiness. Politicians and other public figures are called upon to consider the general utility far more often than the rest of us, and they have correspondingly more responsibility to think through the consequences of what they say and do.

So the moral notion of intention is more complex than either Bentham or Mill realised; we now have a three-clause definition of intention. An agent *a* intends an outcome *x* iff: either *a* foresees and wills that *x* be a consequence of her action, or *a* foresees and is willing that *x* be a consequence of her action, or *a* should – given her knowledge, intelligence, and the context in which she finds herself – have foreseen that *x* would be a consequence of her action. It must be emphasised that his third clause has nothing to do with moral luck; Sidney was unlucky, and therefore escapes censure.

vi. Interference

What exactly is it that Mabel should have foreseen? It's mainly the intervention of other causal factors into the chain of consequences. In my description of E, I said that Mabel's reference to a member of a certain community (let's call them the Lemonmen) is misunderstood, and "the result is an escalation of disquiet, culminating in violence and injuries." My implication was that the misunderstanding was unintentional, so that the only significantly moral agent was Mabel. The Lemonmen would have to bear some moral responsibility, of course — not only because their violence couldn't be justified by reference even to what they *thought* Mabel had said, but because they had a moral responsibility to make every attempt to understand her properly, given the probable effects of any misunderstanding.

What, though, if the misunderstanding had been deliberate? What, for example, if

some rabble-rouser, Herringbone, had interpreted Mabel's words to the Lemon community, distorting her meaning for his own political purposes? Mabel isn't let off the hook, of course; we have to ask whether she should have foreseen Herringbone's action. But, even if the answer is that she should have done so, the matter doesn't end there. We then have to ask whether she could reasonably have acted so as to avoid Herringbone's actions and the resulting violence. One possibility is that Herringbone would have distorted whatever she said or did; she still has to do her best to avoid this, but her being able to foresee it is not the same as her being able to prevent it. In such a case, the violence is not a consequence but only a repercussion of her actions.

Another possibility is that Mabel could only have prevented Herringbone's distortions by backing down on a certain policy decision. Here we have to know the details before we can judge; we have to (as Mabel has to) weigh the effects of dropping the policy against the effects of Herringbone's distortions. This involves not only the specific effects of dropping that policy, but the general effects of backing down in the face of a threat, etc.

This sort of problem is not restricted to movers and shakers, but often faces us as ordinary citizens. I shall take an example concerning which, as a principled non-driver, I am a disinterested observer: the problem of the slow car-driver. It is often said that drivers who drive too slowly are at least as dangerous as those who travel too quickly. The reasoning behind this claim is that the slow driver creates frustration in others, which can result in accidents. A first response is of course to point out that the frustrated speeder is responsible for any accident, and that to blame the slow driver is no more sensible than to blame foggy conditions for multiple pile-ups involving madly speeding cars. Should this be revised, though, in the light of what I've been arguing? The slow driver can surely foresee that others

will be frustrated by being forced to drive sensibly and within the speed limit, and she should therefore speed up. The would-be speeders must still take some (even most) of the blame for resulting accidents, but the slow driver cannot escape responsibility altogether. After all, although I haven't mentioned it for some time, Act Utilitarianism is the moral theory with which my discussion is primarily concerned; the slow driver can't argue that the need for obedience to the law automatically overrides any consideration of the consequences of her actions.

Leaving aside the observation that this (as in the Mabel and Herringbone example) is very much like moral blackmail – 'drive faster or I might cause an accident, and it would be your fault' – the key point here is again the weighing of the slow driver's reasons and intentions against the probable outcome. This is unlikely to be easy or straightforward. Apart from anything else, I suspect that one cause of slow driving is fear engendered by the sight of other cars going much too fast; thus an accident would have as part of its cause the poor driving of the would-be speeder caused by frustration caused by slow driving caused by fear caused by speeders and would-be speeders..

The different rôles of intentions, motives, and consequences are of obvious relevance to the question of social and legal sanctions, and especially when we come to the motives for punishment. If, for example, the purpose of punishment is seen as being to deter others from performing similar actions, then consequences will be thought to be of much more relevance than intentions and motives. In attempting to deter people from stealing, say, the legislator might well consider that failing to punish a thief because her motives were good would weaken the deterrent effect of the law. What counts is the way in which society is seen to respond to certain sorts of actions. If, however, the purpose of punishment is seen as being

primarily to reform the agent, then motive and intention will be likely to outweigh the actual consequences of her actions; an attempted but unsuccessful murder will warrant punishment, while a case of genuine self-defence against an attacker might not, even though the consequence was the death of the attacker. If the purpose of punishment is seen as being retribution, the issue is more complex; strictly speaking, I think, intentions and motives should be more important than consequences (after all, it's the agent against whom retribution is sought), but all too often people in fact desire retribution even when there's no agent, or when it's impossible to establish the identity of the agent; in such cases a scapegoat of some kind is generally found.

vii. Conclusions

One conclusion that can be drawn from much of my discussion here is that, whatever can be said about the nature of agents in general, a *moral* agent must have knowledge of the possible consequences of her actions. It has been argued that being an agent cannot depend upon having the faculty of choice:

a volume of acid cannot choose whether or not to exercise its power to dissolve a lump of zinc [...] But the possession and exercise of the power to make something happen – such as to make a lump of zinc dissolve – is sufficient to make the volume of acid an agent.⁴³

My concern in this paper has been, of course, with the narrower question of moral agency, and whatever the cogency of the general claim, the volume of acid cannot be said to be a *moral* agent. In fact I'm unconvinced even by the general claim, partly because, although I have no objection to the notion of inanimate objects possessing causal powers, I have grave

worries about what is at best a metaphorical use of the verb ‘exercise’. At the root of my discontent is the feeling that to be an agent, to perform actions, just is to be a moral agent — it is to be accountable for one’s actions, to be potentially blameworthy or praiseworthy. What makes an agent a moral agent is the context of action, not some extra property of the agent herself, and to be an agent is to be, not merely causally potent, but rational.

This response to the question of consequences, then, places an emphasis on the rôle of the agent that goes some way, at least, towards answering the claim that Mill’s moral theory is cold and inhuman. The degree to which an agent is morally responsible for the effects of her actions, the degree to which those effects count as consequences, depends broadly upon three factors: her intentions, her knowledge, and the care she takes in relation to her situation and to the context of her actions. The relevantly moral notion of a consequence is thus a complex one – not surprisingly, for morality is a complex business – and in that complexity lies much opportunity for honest mistake; we can behave with all the moral integrity that could be wished for, and still find that our actions give rise to appalling results:

And oftentimes, to win us to our harm,
The instruments of darkness tell us truths,
Win us with honest trifles, to betray’s
In deepest consequence.

(Macbeth, 1.123)

Acknowledgements

I should like to thank Andrea Christofidou, Andrew Wright, and Roger Crisp for their comments on earlier versions of this paper.

Notes

1. see, for example, Bernard Williams (1993, pp 90–91)
2. Spaemann (1989, p.51)
3. *loc. cit.*
4. It's interesting to compare these sorts of criticism of the practical application of Utilitarianism with Bradley's very different complaint that (hedonistic) Utilitarianism offers us as an end of conduct only "an infinite, perishing series" of "this, that, and the other feelings", rather than – as we should demand of a practical end – "some definite unity, some concrete whole that we can realize in our acts, and carry out in our life" (Bradley (1962, p.95)).
5. Perhaps this notion arises for some people when they misapply the slogan 'the greatest happiness of the greatest number' to Mill's Utilitarianism, despite his rarely mentioning it, and then only when referring to other writers.
6. Mill (1861, 2:19:25–31)
7. Sedgwick, p.64, quoted in Mill (1835, p.63)
8. Mill (1835, p.63). J.J.C. Smart's defence against this sort of complaint is also interesting: as the effects of one's actions ramify into the future, it becomes more and more likely that the good and bad effects will cancel each other out for the purposes of moral calculation (see Smart & Williams (1987, pp 64–65).
9. Mill (1872, Book VI, Chapter II, section 3, p.839)
10. In fact even such a modest intention surely fails to hit its mark; one can as easily avert

death by asphyxiation or by starvation (by providing air or food at the last minute) as one can avert death by poison.

11. Mill (1835, p.63)

12. Mill (1861, 2:20:4–8)

13. Mill (1861, 2:20:16–18)

14. Mill (1861, 2:20:30–31)

15. Mill (1861, 2:19:21n.)

16. Mill (1861, p.121)

17. Mill (1861, 2:19:21n.)

18. This talk of causes could of course be translated into talk of reasons without violence to the central thought.

19. Mill (1869, pp 252–3), quoted in Ridge (2002, pp 57–58)

20. Mill (1861, 2:19:21n.)

21. Ridge (2002, p.59)

22. Ridge (2002, p.60)

23. Mill (1861, 2:19:21n.)

24. Mill (1861, 2:19:14–21)

25. Maclean (1993, pp 92–93)

26. Maclean (1993, p.92)

27. *loc. cit.*

28. Leaving aside the possibility that she bungle the action, and the morally irrelevant differences between what it takes to comfort and cheer up different people in different circumstances. In the former case, of course, if her failure to comfort her grandmother is because of the discernible lack of genuine warmth in her actions, then her motive becomes directly morally relevant.

29. see Nozick (1981)

30. Slote (1985, pp 108–109)

31. Nozick (1981, p.173)

32. This echoes a decision made in a 1971 trial that has become known as ‘the dirty old men’ case. This was brought by the Southampton police against the owners of a bookshop found to have been selling books and magazines held to be obscene — ‘obscene material’ being defined in the Obscene Publications Act of 1959 as tending to deprave and corrupt persons who were likely to read or see it.

The justices decided that, as the regular customers of the bookshop, and therefore those who were likely to read or see the material in question, were “inadequate, pathetic, dirty-minded men seeking cheap thrills [...] whose morals were already in a state of depravity and corruption”, they could not be further depraved and corrupted by the material in question, which was therefore not obscene within the meaning of the Act. (This verdict was upheld on appeal, but overturned by the House of Lords on a majority decision.) See *Director of Public Prosecutions v Whyte and another* [1972] 3 All ER 12 for details of the House of Lords appeal and of the original case.

See Peter J. King, ‘No plaything: ethical issues concerning child-pornography’ [unpublished] for a critical discussion of this sort of argument.

33. A third of Mill’s extensions of the notion of cause allows for causal overdetermination, which had a rôle in creating the problem with C_1 in the first place.

34. Hart and Honoré (1985, p. 16). In light of Mill’s other extension of the notion of causation, the phrase “the cause of” in these examples should, of course, be altered to read ‘part of the cause of’ or ‘a causal element of the cause chain leading to’ (for an explanation of this

somewhat cumbersome formulation see §iv below).

35. Mill (1872, Book III, chapter V, section 3)

36. Bennett (1987, p.373)

37. Bennett (1988, p.71)

38. Mackie (1992, p.490)

39. I follow Mackie in interpreting 'lengthiness' as referring to the length of time between e and r , and 'remoteness' as referring to the length of (the number of links in) the causal chain between e and r .

40. In what follows I shall comment on only a small part of her discussion, and only as it relates to my present concerns.

41. Mackie (1992, pp 483–84)

42. Mill (1861, 2:19:32–37)

43. Alvarez & Hyman (1998, p.245)

References

Apart from *Utilitarianism*, all references to Mill's works are to the *Collected Works of John Stuart Mill*, general editor John M. Robson. Toronto & London: University of Toronto Press and Routledge and Kegan Paul, 1961–69):

[1835] 'Professor Sedgwick's Discourse on the Studies of the University of Cambridge' (*Volume X*; 1969)

[1869] *James Mill's Analysis of the Phenomena of the Human Mind* (*Volume XXXI*)

[1872] *A System of Logic* [8th ed.] (*Volumes VII & VIII*; 1974)

[1873] *Autobiography* (*Volume I*; 1980)

In the case of *Utilitarianism* I've preferred to use Roger Crisp's *Oxford Philosophical Texts* edition, partly for the usefulness of its clear numbering of paragraphs and lines:

[1861] *Utilitarianism* — ed. Roger Crisp. Oxford: Oxford University Press, 1998.

Alvarez, Maria & John Hyman, 'Agents and their actions', *Philosophy* 73 (1998), pp 219–245

Anderson, Brian A., 'Mill on Bentham: from ideology to humanised utilitarianism', *History of Political Thought* 4 (1983), pp

Bennett, Jonathan, 'Event causation: the counterfactual analysis', in James E. Tomberlin [ed.] *Philosophical Perspectives, 1: Metaphysics*. Atascadero, California: Ridgeview Publishing Co., 1987, pp

Bennett, Jonathan, *Events and their Names*. Indianapolis: Hackett Publishing Co., 1988.

Bentham, Jeremy, *An Introduction to the Principles of Morals and Legislation* (1789). Edd J.H.

Burns & H.L.A. Hart, Oxford: Oxford University Press, 1996.

Bradley, F.H., *Ethical Studies* (2nd ed.). Oxford: Oxford University Press, 1962.

Hart, H.L.A. & Tony Honoré, *Causation in the Law* (2nd ed.). Oxford: Clarendon Press, 1985.

Mackie, Penelope, 'Causing, delaying, and hastening: do rains cause fires?' *Mind* 101 (1992), pp 483–500.

Maclean, Anne, *The Elimination of Morality: Reflections on Utilitarianism and Bioethics*. London: Routledge, 1993.

Nozick, Robert, *Philosophical Explanations*. Oxford: Clarendon Press, 1981.

Oakley, Justin & Dean Cocking, 'Consequentialism, moral responsibility, and the intention/foresight distinction', *Utilitas* 6 (1994), pp 201–216.

Ridge, Michael, 'Mill's intentions and motives', *Utilitas* 14 (2002), pp 54–70.

Slote, Michael, *Common-Sense Morality and Consequentialism*. London: Routledge and Kegan Paul, 1985.

Smart, J.J.C. & Bernard Williams, *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 1987.

Spaemann, Robert, *Basic Moral Concepts*. London: Routledge, 1989. (Translated by T.J. Armstrong from *Moralische Grundbegriffe*. Munich: C.H. Beck'sche Verlagsbuchhandlung, 1982.)

Unger, Peter, 'The uniqueness of causation', *American Philosophical Quarterly* 14 (1977), pp 177–188

Williams, Bernard, *Morality*. Cambridge: Cambridge University Press, 1993.
