9

Deontology*

DAVID A. MCNAUGHTON AND J. PIERS RAWLING

INTRODUCTION

Which actions does morality require of us? What does it forbid and what does it permit? In trying to find some general answers to these questions, moral theorists typically start from commonsense morality, from what ordinary people think about moral issues. In deciding how to act, people often think about the consequences of their actions: they try to find the action that leads to the best overall outcome. One moral theory, act-consequentialism, claims that this is the *only* consideration that is relevant to moral choice. The right action – the one we are required to do – is the one that produces the most good; it is wrong to do less good than we could.

Act-consequentialism seems, however, to conflict with commonsense morality. Although we should be concerned to make things go as well as possible for everyone, most people do not think that this exhausts morality, or even identifies some of its most crucial elements. Are we not, for instance, sometimes required to aid our loved ones, even if we do not thereby produce the best overall? And are there no limits on what we may do to produce good, or limits on what we must do to produce it? Deontology contrasts with consequentialism in its answers to these questions, and is, in one of its versions, the theory we favour.

COMMON SENSE MORALITY

Here are three areas of ordinary moral thought in which considerations other than the amount of good our actions would produce are normally taken to be relevant to what we morally ought to do.

OPTIONS

Act-consequentialism appears very demanding. Given the amount of poverty in the world, maximizing the good would require the better off to make enormous sacrifices to help the very poor. Most people believe that, though they should do something to help those less fortunate, there is nothing wrong with devoting a lot of time, effort and money to one's own happiness and the happiness of those one cares about. There is some point, perhaps hard to determine, at which someone has done all that they are *required* to do by way of helping strangers. At that point they are morally permitted, or have an *option*, not to do more. We admire those who make the extra sacrifice, but it is supererogatory – more than morality requires. Act-consequentialism, however, seems to leave no room for supererogation.

Deontologists do not deny that morality can be very demanding. We may be obliged to make significant sacrifices – even of our lives – rather than breach a serious constraint (see below) or betray a friend. But we are not constantly required to be promoting the general good.

DUTIES OF SPECIAL RELATIONSHIP (OR SPECIAL OBLIGATIONS)

Many people believe that not only are we *permitted* to do more for those close to us, but we are often *required* to put their interests first. We *owe* things to those with whom we have special relationships – such as our friends, colleagues and family members – that we do not owe to strangers. Our own children, for example, have a claim on our attention and resources that other people's children do not. It follows that it would be wrong to neglect our own children, even if we could thereby do slightly more good for other children.

CONSTRAINTS

In addition to our special obligations, many people believe that we have a duty to avoid seriously harming *anyone* unless, perhaps, they are a threat or deserve punishment. We should not lie, kill innocent people or torture. These prohibitions *constrain* us in what we may do to *any* person (not just those close to us), even in pursuit of good ends. People differ in how stringent these constraints are. Some think them as absolute or exceptionless. Roman Catholic moral theology has traditionally held that one may never intentionally kill an innocent person, even to prevent, say, two other innocents from being intentionally killed. Kant infamously argued that it would be wrong to lie, even to prevent murder. Others hold that, though constraints are always a significant consideration, they may be overridden, especially to avoid catastrophe. Either way, such constraints would sometimes *require* us not to maximize the good.

DISTINGUISHING DEONTOLOGY FROM ACT-CONSEQUENTIALISM

How does consequentialism differ from deontology? A traditional answer points to their contrasting accounts of the relation between the right and the good. Act-consequentialism holds that the good wholly determines the right – which act is right depends solely on the amount of good it produces. Deontologists maintain, by contrast, that other considerations, of the kind discussed in previous sections, are also relevant and that consequentialism fails to take them into account. How might an act-consequentialist respond to this charge?

She might plead guilty, but claim that ordinary moral thinking is confused and unreliable as a guide to moral thinking. Since consequentialism is the correct theory, we should eliminate elements in our moral thinking that do not conform to it. Let us call this the defiant strategy.

Alternatively, she might deny that her theory puts her at odds with common sense morality. Consider special obligations: there might be good consequentialist reasons for encouraging people to do more for those close to them than for strangers. We are in a better position to benefit our nearest and dearest since we know their needs and we are more motivated to help them. So we will do more good by focusing much of our attention on them. Or consider constraints. It is not very often that killing an innocent will have the best results, so we may want to discourage the thought that it could ever be permissible. If our valuable reluctance to kill the innocent were weakened then people might be tempted to kill when doing so would not produce the best results. In short, from a consequentialist perspective, things might go better if people were guided by common sense morality, rather than directly making decisions using consequentialist criteria. Let us call this the conformist strategy.

We do not have space here to consider whether the conformist strategy is defensible. Even if it were, act-consequentialism would still differ from deontology concerning what makes acts right. The conformist strategy points out that in less than perfect conditions – imperfect knowledge and imperfect motivation to be moral – we may well do better to act in accordance with the dictates of ordinary morality than to try to produce the best results and thereby inadvertently make things worse. But, for the act-consequentialist, the right act remains the one that maximizes value. Under ideal conditions of knowledge and motivation, therefore, the virtuous agent should never produce less good than she could. The deontologist disagrees: she believes that, even under these conditions, we are sometimes permitted, and even required, to do just that.

The three elements of common sense morality to which we have drawn attention are distinct, so that it would be possible, for example, to believe (like Ross) that there are special obligations and constraints, but no options. Or (as we are inclined to do) one might accept that there are special obligations and options, but deny that there are constraints. Or (like some Kantians) one might accept constraints and (perhaps) options, while leaving less room for special obligations. Or (like Scheffler) one might accept only options. Which elements must one accept in order to qualify as a deontologist? A key feature of deontology is the claim that we are sometimes *required* not to maximize the good. So a deontological theory must include this claim. Options only give us *permission* not to maximize the good, but duties of special relationship and constraints can require us not to do so. Thus one or both of these features will typically be part of any deontology.

Can we say more about what characterizes deontology? We could supply a list of the various duties of special obligation and constraints. But is there a unifying theme to such a list? Ross, for example, gives a list that includes – in addition to beneficence (producing good results) – fidelity (to promises and commitments), gratitude, reparation and avoiding harming others. These considerations, he plausibly maintained, were distinct, and none is reducible to any of the others.[†] However, it has been suggested (originally by Nagel 1986) that the considerations that fall

---

[†]For a discussion of whether Ross's theory exhibits sufficient unity see McNaughton (1996).

under our three headings do at least share a common *form*: they are all *agent-relative*.

AGENT-RELATIVITY AND AGENT-NEUTRALITY

The distinction between the agent-neutral and the agent-relative may be introduced by reference to reasons for acting[‡]. Roughly, someone's reason is agent-relative if, at base, there is reference within it to her. For example, rational egoism is an agent-relative theory – it holds that each agent has reason to promote only *her own* good, whereas act-consequentialism is an agent-neutral theory – it holds that each of us has reason to promote *everyone's* good. Another way of making this point (which we owe to Parfit (1987) is that consequentialism gives us the *common* aim of promoting the general or impersonal good, whereas according to egoism each of us has the *distinct* aim of promoting his personal good: I have reason to pursue my good, you yours. In contrast to act-consequentialism, deontology is an agent-relative theory: at its base, there are agent-relative as well as agent-neutral moral reasons. Each of the three elements in deontology incorporates agent-relativity.

Special obligations are obviously agent-relative. I am required to care for *my* family, you for *yours*: we have distinct aims. Act-consequentialism might allow that parental care giving is valuable, but on this view we would have the common aim of promoting parental care-giving in general. That would require that I neglect my own children if I can thereby increase the total amount of parental care giving – a claim that deontology denies.

Constraints are also agent-relative. Suppose I can only prevent you from killing two innocents by killing one myself. If there are constraints, then each of us has strong (or even overriding) moral reason not to kill anyone ourselves. Constraints give each of us distinct aims: I have reason not to kill anyone *myself*; you have reason not to kill anyone *yourself*. Thus although you will do wrong in killing the two, I should not kill the one in order to prevent you. Consequentialism, by contrast, holds that everything else equal, it is right to kill an innocent myself to save two: killing innocents is bad, so I have an agent-neutral moral reason to contribute to the common aim of minimizing such killing.

Options need not be agent-relative in their formulation. They simply permit us not to maximize the good. But the standard rationale for admitting options into a moral theory is agent-relative. Each of us is morally permitted to give special weight to *our own* interests, just because they are ours.

RULE-CONSEQUENTIALISM: A DEONTOLOGY IN DISGUISE?

Act-consequentialism is not the only version of consequentialism. Perhaps its most popular rival is rule-consequentialism, which offers a 'two-stage' account of justification. Rule-consequentialism assesses acts, not in terms of their contribution to the good, but by whether they conform to the best set of rules governing human conduct. Rules, however, are assessed by their contribution to the good. The best set of rules is the one whose general acceptance would produce most good. Thus, according to rule-consequentialism, 'an act is wrong if and only if it is forbidden by the code of rules whose internalization by the overwhelming majority . . . has maximum expected value' (Hooker, 2000). Which rules are best is determined, in part, by the psychological make-up of human beings. They must, for example, be simple enough for ordinary people to learn and sufficiently appealing that the majority of people can be persuaded to follow them. Given these restrictions, acceptable rules will probably be close to the rules of common sense morality. In particular, rule-consequentialism is likely to include constraints, options and special obligations. There will, for example, be a fairly simple rule against killing the innocent, since having a more complicated rule that allowed killing whenever it would do most good might be disastrous. Given our natural concern for our nearest and dearest and the need for companionship and security, there will also be rules permitting, and even requiring, us to give priority to the claims of friends and family.

How are we to classify rule-consequentialism? Despite its name, it might seem to have more in common with deontology. It agrees with deontology that it is often wrong to do the act that will produce the most good and that some of the moral rules we should follow are agent-relative in form. However, we follow Hooker, a leading rule-consequentialist, in classifying this theory as fundamentally agent-neutral and thus consequentialist. Each form of consequentialism assesses something, at its base, in terms of impersonal or agent-neutral value. But what they assess varies: act-consequentialism assesses acts, while rule-consequentialism assesses rules. As Hooker notes, this makes 'the agent-relativity in rule-consequentialism…derivative. Agent-relative rules are justified by their role in promoting agent-neutral value' (Hooker, 2000). Deontology, by contrast, holds that some agent-relative considerations are *underivatively* relevant. They have weight in their own right, not merely in virtue of their serving some further purpose.

[‡]But see McNaughton and Rawling (1991) for discussion of some problems for this approach.

VARIETIES OF DEONTOLOGY

KANT

Kant's moral theory is the best known example of a deontology. As it is discussed in chapter 10 we only mention points that are relevant to our discussion. First, Kant strongly rejects the whole consequentialist approach to moral theory. He denies that the value of its consequences is relevant to the rightness of an act, and he claims that there are kinds of action, such as lying, which are *always* wrong. Second, he gives structure to his theory by offering us the categorical imperative test to determine whether an act is required, permissible or forbidden. Finally, failure to act in accordance with the results of that test is to be, in some strong sense, irrational. Famous though it is, Kant's is not the only nor, in our view, the most plausible deontological theory.

ROSSIAN DEONTOLOGY

How can we find out which moral considerations are fundamentally relevant to how we ought to act? Ross's answer is that we do this by reflecting on our intuitions about contrasting cases. (Ross 1930, chapter 2) Does, for example, having made a promise make a difference? Ross claims that no one would deny that, if we could give an equal amount of help to two people, the fact that I had made a promise to one of them would make it right to help that one. Such reflection, he claims, eventually reveals a number of basic moral considerations, which he formulates as a list of moral principles or duties. As we have seen, his suggested list includes agent-relative duties of promise keeping, gratitude, reparation and not harming others,[§] as well as an agent-neutral requirement to promote the good. These duties are, he says, only *prima facie* (or, as we prefer, *pro tanto*) since, though each is relevant to determining what is right, no duty automatically trumps any of the others in cases where they conflict. Indeed, there is no general method for resolving such conflicts: which duty is the weightier depends on the complexities of the particular circumstance. Where keeping a promise will harm someone, for example, what we ought to do will depend on how serious the promise is, to whom it was made, how much harm we would do in keeping it and so on. To weigh all these factors correctly requires discernment and judgement.

The items on Ross's list are intended to be basic in two ways. First, they are *underivative*, in that they are not instances of some more general principle. So, for example, the duty to pay one's debts is derivative because it is an instance of the more general duty to keep promises. But the latter duty is underivative because it is not itself an instance of some yet more general duty. Second, they are basic in the sense that they are *self-evident* and need no justification. We can see their truth directly, without reasoning from further premises. Indeed, Ross strongly implies that they not only require no further justification but that none is available. He thus rejects the Kantian claim that such basic principles rest on a common foundation. For Ross there is no test that principles must pass to earn their place on the list.

Of the two deontological theories we have looked at so far, Kant's and Ross's, which is the more attractive? This depends, in part, on what one is looking for in a moral theory. Kant's programme is remarkably ambitious; Ross's is very modest. Kant seeks a *grounding* for common-sense morality, whereas Ross holds that none is available and none is needed. Kant supposes that, at bottom, there is something that morality is about; Ross makes no such assumption. Kant offers a test that purports to show that acting immorally is indisputably *irrational*, rather in the way that it is irrational to contradict oneself. Ross holds that there are moral reasons, but failure to appreciate them would be a sign of moral insensitivity or immaturity, rather than a gross failure of reasoning. Kant maintains that certain kinds of action, such as lying, are morally unacceptable in all cases; Ross rules nothing out in advance – it all depends on the facts of the particular case. Ambitious theories may be exciting and stimulating in a way that modest ones are not, but that does not make them correct. Not only may they fail to realize their ambitions, but those ambitions may themselves be misguided. It may, for example, be an error to suppose that our moral thought stands in need of special justification, or that a successful moral theory should provide unequivocal guidance in perplexing cases. In our view, the modesty of Ross's account is a strength, not a weakness.

PARTICULARIST DEONTOLOGY

To what extent can morality be codified into a set of rules or principles? Some moral theorists think that there are *strong* moral principles that tell us that we are always forbidden (or required) to act in certain ways. Thus Kant claimed that lying is always wrong. Others, such as Ross, think there are only *weak* principles that tell us that certain features always matter, morally speaking The features on Ross's list of *prima facie* duties are supposed always to count for (or against) the rightness of an act. Particularists, however, doubt the existence of even weak principles and deny that we need them to engage in moral thinking. On their view,

_____

[§]This last is, we think, agent-relative for Ross: he does not appear to countenance doing harm oneself in order to minimize the total amount of harm.

almost any feature can be morally relevant and none need always be relevant – it all depends on the context. So, for example, the fact that an act brings pleasure often counts in its favour, but not, perhaps, when the pleasure is sadistic. That it would bring sadistic pleasure is *no* reason to perform an act.

Who is right – Ross or the particularist? The matter is hotly disputed. Ross seems to be mistaken about one member of his list – promise-keeping. In normal circumstances the fact that I would be breaking a promise does indeed count against acting in that way. But promises extracted by fraud or force are null and void, as are promises to do something immoral. Suppose I promise to perform a contract killing. It is implausible to hold that, though, all things considered, I ought not to do it, yet the fact that I promised gives me *some* moral reason to do it. However, other considerations on Ross's list seem more plausible candidates. How could an act's being just, for example, not count in its favour?

DEFENDING DEONTOLOGY

There is nothing puzzling about the consequentialist claim that the amount of (agent-neutral) good we could do is morally relevant. Deontology, however, claims that there is also an *underivative* agent-relative component to morality, and this does seem puzzling. How can the *mere* fact that some act will be good or bad for *me*, or for *my* nearest and dearest, itself make a moral difference? Or how can the *mere* fact that *I*, rather than someone else, will kill an innocent be morally relevant? In what follows we try to meet this challenge in the case of special relationships and options, while conceding that it is harder to meet it in the case of constraints.

SPECIAL RELATIONSHIPS
AND THE TIES THAT BIND

On our view, we have underivative, agent-relative, moral ties to certain people in virtue of our relationships with them. Consider the tie of friendship. Such ties are agent-relative: part of my reason for helping Eve rather than Adam is that Eve is *my* friend. These ties are underivative because such reasons do not depend wholly on considerations about the general value of friendship. Friendship is valuable, but that fact does not explain *my* tie to Eve. Finally, the ties of friendship have a moral component. It is not just that friends like and support each other. In addition, your friend has the *right* to expect *your* loyalty and support because she is *your* friend. If you betray her, she has a moral complaint against you that no one else has. The (tacit) acknowledgement of a moral tie between friends appears essential to friendship.

Friends come through for one another; someone who did not even recognize this would not be loyal and so would not be a friend.

Act-consequentialism does get something right: we only have reason to act when it will do some good. But how much good I can do is not the only factor. Who will receive the good also matters. Friends have stronger claims than others to my good offices. Which act would be right depends on two factors: the good I can do someone and the strength of the claim she has on me.

If this is right, then consequentialism has a serious strike against it. Loyalty is essential to friendship. Loyalty involves the recognition of an underivative agent-relative obligation to one's friends. Consequentialism has no place for underivative agent-relative obligations; thus it has no room for friendship. But friendship, as is generally acknowledged by consequentialists, is an important intrinsic good. Consequentialism holds that the good is to be promoted; but here is a good that it apparently cannot accommodate.

How might the consequentialist respond? Act-consequentialists, as we saw, may opt for a strategy of defiance or conformity. Those who favour defiance may deny that the moral point of view has room for friendship. Morality, on this approach, requires giving equal weight to everyone's interests. Since we are not even permitted, still less required, to favour any group of people, the consistent act-consequentialist cannot have friends.

Unsurprisingly, most act-consequentialists reject this stark approach and opt for the conformist strategy. We saw earlier (see Section Distinguishing Deontology From Act-Consequentialism) that an agent might do better in achieving consequentialist goals if she did not think about which act is right on each occasion, but instead had dispositions to favour friends and family. Pursuing this line, the act-consequentialist might concede that friendship requires loyalty, but then maintain that loyalty does not require that there actually be moral ties, only that people *believe* there to be such ties. So act-consequentialism might encourage people falsely to believe that they have special obligations.

A serious objection to this strategy, however, is that it pictures the ideal moral agent as someone who is fundamentally mistaken about what makes acts right or wrong. In effect, such a consequentialist is saying: consequentialism is the theory you have most reason to believe, but it would be better, from the consequentialist perspective, if no one believed it. If not actually incoherent, this position is certainly uncomfortable. Simpler and better to accept that special obligations exist, unless there are compelling reasons to believe that they do not. But no such reasons have been offered.

Rule-consequentialism does rather better. Unlike act-consequentialism, it acknowledges that one can be acting rightly in favouring friends, but it fails to capture friendship because it maintains that preferential treatment of friends can be justified only by appeal to the general good:

> Moral requirements of loyalty are . . . needed . . . when affection isn't up to the job. . .. [S]pecial moral obligations towards family and friends can then be justified on the ground that internalization of these obligations gives people some assurance that some others will consistently take a special interest in them. Such assurance answers a powerful psychological need (Hooker, 2000).

This does not yield genuine loyalty. Friends have moral reason not to let us down, and assurance is engendered in part by a belief that they will respond to this reason. (This is not to say that the only reasons here are moral.) But for rule-consequentialism, the *moral* reason for John not to let Mary down is the general assurance that results from the internalization of a rule requiring the special treatment of 'friends', not anything special about his relationship with Mary.

A final objection to our account. Morality requires us to be impartial, but do not duties of special relationship require us to be *partial* to our friends, and so on? No. In our view, we show partiality in allocating goods only if we give the claims of one person or group more weight than we are *warranted* in doing. To favour those who have a claim on us is not to show an unacceptable favouritism. We do that only when we give *undue* or *inappropriate* weight to the interests of our nearest and dearest.

OPTIONS

Each of us has special personal reason to pursue our own benefit, just by virtue of it being ours. I have personal reason to care about my pain that I cannot have to care about yours, namely that it is mine. This does not mean that I have no reason to care about your pain, nor does it commit me to denying that pain is equally bad whoever has it. But each agent has moral permission not to maximize the good when the cost to *her* would be significant.** An agent is allowed, in determining what she is morally required to do, to accord greater weight to the cost borne by her than to the cost borne by others. The act-consequentialist denies this: whether an act is morally permitted, on her view, depends solely on the amount of impersonal value it would produce. She might try to soften that view by arguing that if too much is required of agents they will become disillusioned or exhausted and so do less good than if they had more modest targets. But that strategy allows us to favour ourselves merely as a

concession to our weakness, not as an acknowledgment of a right.

In our view there has to be some balance between the demands that the needs of others put on us and our right to live our own lives. Determining where that balance lies is notoriously difficult. But this does not entail that there is no balance to be struck.

Moreover, the act-consequentialist denies room for supererogation (acting beyond the call of duty): she maintains that the person who bears great personal cost in maximizing the good, although admirable in the extreme, would be doing something morally wrong if she did otherwise.

Rule-consequentialism appears to do better. An agent who follows the rules does not act wrongly: she does enough good – to do more would be meritorious but is not required. The presence of this personal space, however, stems from impersonal costs: we are psychologically resistant to making significant sacrifices, and this makes it too expensive to inculcate a more demanding rule. But this resistance is understood as a regrettable flaw, not a mark of personal reasons. Rule-consequentialism denies the *moral* significance of personal reasons at the fundamental level. They matter only because of the cost of training people to ignore them.

CONSTRAINTS

Constraints, though often regarded as the most distinctive feature of deontology, seem harder to justify than options or duties of special relationship. Consider an absolute constraint against (intentionally) killing an innocent person. Suppose Anne and several other innocents are about to be shot by Bert, but he agrees to let the others go if you shoot Anne. The constraint forbids you to do it. Yet, as Scheffler (1994) points out, this appears inexplicable: Anne is going to be shot, but at least you can prevent the other shootings.

Some advocates of constraints might concede that you should shoot Anne in this case, but they would object to your shooting some unthreatened innocent in order to save Bert's intended victims. But even this weaker position does not, in the end, evade the standard 'irrationality objection' to constraints: that they forbid their own violation *even to minimize such violation*.

An important feature of constraints, as understood by traditional deontology, is that they are underivative. Rule-consequentialism incorporates prohibitions against, say, killing the innocent, but these are *derivatively* 'justified by their role in promoting agent-neutral value' (Hooker, 2000). In that sense, rule-consequentialism does not include constraints.

We have defended special obligations and options by contending that, in addition to the amount of good we do,

---

**For an extended defence of this approach, see Scheffler (1994), *passim*.

what we might call positional facts – that the good would accrue to *my* friends or to *me* – are also morally relevant. Constraints, however, cannot be similarly defended. Constraints, unlike duties of special relationship, single out no group on the basis of my relationship to its members, thus they cannot rest on my being more closely related to some than others. Could they be justified on the grounds that my violating a constraint, even to prevent worse actions by others, is bad *for me*, and so something I have personal reason to avoid? After all, many of us have an understandable reluctance to get our hands dirty. But that thought, at best, *might* ground a *permission* not to violate the constraint. It cannot ground a *requirement* not to do so.

Constraints, then, embody as fundamental the fact of your agency. It is the bare fact that you would be doing the killing that rules out your killing an innocent as a means to preventing other innocents suffering a like fate. Since constraints are fundamental, we should not expect to find a deeper justification for them. Their being fundamental does not, however, preclude our defending constraints, as we did with special obligations and options, by explicating their nature in ways that make their force clearer. The problem is that we can find no explication of constraints that dispels their air of irrationality.

Given that air of irrationality, it is unsurprising that many deontologists have nevertheless attempted to find a deeper justification, but such attempts to explain why agency matters seem to presuppose the very point at issue (see Scheffler, 1994, ch. 4). Thus it is said (in a Kantian vein) that persons deserve respect in view of their unique importance as rational moral agents. But why does such respect forbid you to harm others rather than requiring you to minimize harm? It may be said that just as we owe particular duties to others in view of our special relationships with them, so we owe to everyone else a duty not to harm them because of our general relationship with them. But what is that relationship? Perhaps that of being fellow humans or fellow persons. Whatever the answer, the problem remains: why does our standing in that relationship ground a *constraint* against harming as opposed to a duty to minimize harm?

Some defenders of constraints[††] have complained that in seeing constraints as *agent*-relative, recent attempts to ground constraints have wrongly focused on agency. Rather, they claim, we should focus on a *patient*-centred justification – on what it is about innocents that entails the existence of constraints against harming them. But this does not seem to help. Innocents do not deserve to be harmed; ideally, then, we should not harm them. But what are we to do in our non-ideal world in which innocents are under threat?

If we deny there are constraints, however, would not we have to abandon the many intuitions that seem to support their existence? Not necessarily. Many intuitions that appear to be support constraints may actually rest on other features. We may, for example, think it wrong to take $10 from one person in order to enrich another by $20. But our grounds for that (depending on the circumstances) might be that the harm of taking $10 *honestly possessed* outweighs the benefit of bestowing $20 *unearned*. Second, we may think it wrong to do considerable harm to one person in order to prevent small harms to a large number. But that may have nothing to do with agency but be explained by the fact that no number of small harms to different people, when added together, is as bad as a serious harm to one person.

We are tentatively proposing, then, a morality devoid of constraints (as traditionally understood) but incorporating underivative duties of special obligation and options. Is such a theory worthy of the name of deontology? Yes. Ours is a theory that acknowledges underivative agent-relative requirements not to maximize the good.

REFERENCES

Brook R. Agency and morality. *J Philos* 1991, **88**: 190–212

Hooker B. *Ideal Code, Real World*. Oxford: Clarendon Press, 2000.

Kamm F. Nonconsequentialism. In: LaFollette H, ed. *The Blackwell Guide to Ethical Theory*. Oxford: Blackwell, 2000; pp. 205–26.

McNaughton D, Rawling P. Agent-relativity and the doing-happening distinction. *Philos Stud* 1991; **63**: 167–85.

McNaughton D, Rawling P. Deontology and agency. *Monist* 1993; **76**: 81–100.

McNaughton D. An unconnected heap of duties? *Philos Q* 1996; **46**: 433–47.

Nagel T. *The View from Nowhere*. Oxford: Oxford University Press, 1986; ch. 9.

Parfit D. *Reasons and Persons*, revised edition. Oxford; Clarendon Press, 1987; p. 27.

Ross WD. *The Right and the Good*. Oxford: Clarendon Press, 1930.

Scheffler S. *The Rejection of Consequentialism*, revised edition. Oxford: Clarendon Press, 1994.

[††]See for instance Brook (1991), Kamm (2000). For discussion see McNaughton and Rawling (1993).