

ON DEFENDING DEONTOLOGY*

David McNaughton and Piers Rawling

Abstract

This paper comprises three sections. First, we offer a traditional defence of deontology, in the manner of, for example, W.D. Ross (1965). The leading idea of such a defence is that the right is independent of the good. Second, we modify the now standard account of the distinction, in terms of the agent-relative/agent-neutral divide, between deontology and consequentialism. (This modification is necessary if indirect consequentialism is to count as a form of consequentialism.) Third, we challenge a value-based defence of deontology proposed by Quinn (1993), Kamm (1989, 1992), and Nagel (1995).

Each time we act, we make the world different from what it would otherwise have been. In principle, the various outcomes that would result from the different courses of action open to us can be ranked in terms of their value. According to direct-act consequentialism, our task as moral agents is to increase value. And the more good we can do the better. So the right action is the best one; the one that produces more good and less harm than any other¹ – and it is *solely* the value of the actions open to us that is relevant to which is the right one. We judge an act by the value of its consequences, using the term ‘consequence’ in a sufficiently broad sense to include the performance of the act itself, as well as what flows causally from it. Consequentialism is not, in itself, a complete moral theory – a theory, that is, which tells us which actions are in fact right and why. We need to add to it an account of what things are valuable and what things are bad. Consequentialism itself provides a formal structure within which a family of substantive moral theories can be found.

Deontological theories lack this structure. At the fundamental

* We have benefited greatly from the comments of Jonathan Dancy and Brad Hooker, and from correspondence with John Skorupski. Piers Rawling would like to acknowledge financial support from the NEH and the University of Missouri Research Board.

¹ As we discuss below, this is not true on all forms of consequentialism.

level, we claim, deontology, in opposition to consequentialism, acknowledges moral reasons that do not rest on considerations of value. There is an intriguing argument abroad, however, that attempts to defend deontology by arguing that worlds in which there are deontological reasons are better (i.e., more 'valuable') than worlds in which there are not (see, e.g., Quinn (1993), Kamm (1989, 1992), and Nagel (1995)). We begin by giving what has become, following such authors as Nagel (1986), Parfit (1984), and Scheffler (1982), the generally accepted contrast between (direct-act) consequentialism and deontology. We then outline our favoured defence of the latter, and urge that the standard way of distinguishing it from consequentialism must be modified in order to accommodate indirect consequentialism. Finally, we explore, and reject, the initially appealing value-based defence of deontology.

I

It is central to consequentialism that value is determined *impersonally*:² the real value of any state of affairs does not depend on the point of view of the agent. Features of the particular agent may be morally relevant, but only in so far as they bear either on which state of affairs will be impersonally best, or on what range of actions is open to the agent. Direct-act consequentialism provides an *agent-neutral* account of the right: the right act is that which maximizes the impersonal good. Thus, on this account, all agents share a common aim: that the world go as well as possible. The notion of agent neutrality is somewhat tricky to spell out in terms of reasons; but, roughly speaking, an agent-neutral reason is one that ultimately rests on considerations that make no reference to the agent for whom it is a reason. (In, e.g., McNaughton and Rawling, 1991, we discuss the difficulties of distinguishing between agent-neutral and agent-relative reasons; and we there provide a precise account of the agent-relative/agent-neutral distinction in terms of rules – we omit this account here for reasons of brevity.)

A reason is *agent-relative*, by contrast, if there is an ineliminable reference to the agent in the ultimate statement of the reason. Thus, to take a simple example, Lee's reason for insuring 73

² Some authors (e.g., Dreier, 1993) use the term 'consequentialism' in a wider sense to include theories in which this is not so. We shall steer clear of such usage here.

Hatfield Street is agent-relative if she does so simply because it is *her* house, and she wants to protect *her* property. If her reason for doing so were agent-neutral, on the other hand, it would have ultimately to rest on the thought that everyone should insure their own houses (so that in the unlikely circumstance that by not insuring her own house herself she could thereby increase the number of people insuring their own houses, then she should not, *ceteris paribus*, insure her house).

Reasons that stem from considerations of self-interest would appear to be agent-relative.³ Roughly speaking, if an agent performs a certain action, A, because Aing is (at least in part) in her self-interest, then that factor, with its ineliminable reference to her, will be noted in the ultimate statement of her reason for Aing. Lee, when pressed as to why she insures her house, might declare: 'Because it is in *my* long-term self-interest'.

But there seem also to be *moral* reasons which are agent-relative in form. And an agent-neutral theory, such as consequentialism, apparently (see section II) cannot accommodate them. There are at least three areas of moral thought in which agent-relative reasons figure.

First, there are duties which stem from special relationships: duties which I owe to some specified individuals because of the relationship in which I stand to them. Instances of such relationships include those of parent to child, spouse to spouse, friend to friend, as well as purely contractual relationships, such as those of promiser to promisee, of debtor to lender. Why do these relationships generate agent-relative reasons? Because in each case the agent's reason for acting, her obligation to act, stems from the fact that *she* stands in this particular relation to another. The fact that Alex is *Lee's* child gives her a reason to look after him which is not shared by anyone who is not also Alex's parent. This is not to say that we have *no* duty whatever to be concerned about the welfare of other people's children, only that each of us has a distinct and special reason to be concerned about our own.

Second, there are *constraints*, which proscribe certain types of action, even if their performance, in a particular circumstance, would make the world somewhat 'better'. Thus, it is generally thought that it is wrong to kill the innocent, torture people, lie or cheat, even in pursuit of an otherwise good goal. To have this

³ Not everyone agrees with this – Jonathan Dancy (in correspondence), for example, demurs.

thought is to reject direct-act consequentialism. We need not go so far as to claim that we should *never* do these things, no matter how dire the emergency; to think that we should not do these things just to make the world slightly better overall is already to reject direct-act consequentialism. How do constraints generate agent-relative reasons? To hold, for example, that it is wrong to kill the innocent, is to be committed to the view that *I* should not do or sanction such killings, even if by killing an innocent myself I might thereby reduce the total number of such killings. This thought is enshrined in the common moral intuition that one should not oneself sink to the level of the terrorist or criminal in pursuit of stamping out terrorism or crime.

The third area of ethical thought in which agent-relativity has a place is in the availability of *options*. Direct-act consequentialism is a very demanding theory. It tells us that we should always act so as to maximize the good. As the world is currently organised, this would require us to make huge and continuous sacrifices. We might have to abandon important personal projects if that were the only way to maximize the good. Clearly, *my* projects, *my* family life etc. have an importance for me which they have for no-one else – an importance out of proportion to their agent-neutral value. We do not generally think morality requires the continual sacrifice that direct-act consequentialism demands. Someone can be doing the right thing, even if she gives considerably more weight to *her own* projects than they strictly deserve in the consequentialist calculations.

If a satisfactory moral theory must be consequentialist in structure then we shall have to abandon the thought that there is any place for underivative (see section II) agent-relative reasons in morality. It may be, of course, that there are good consequentialist reasons why each of us should care for our own children, or why we should be governed by constraints, but those consequentialist reasons will ultimately be agent-neutral in form. Lee should care for her child, for example, only because the world is a better place on account of such relationships. And she might be required to desert him if the total amount of such caring would be increased thereby.

Deontology embraces agent-relative reasons of some or all varieties. Thus deontological theories fail to fit the consequentialist structural templet. (It is definitive of deontology that it includes constraints and duties of special relationship; not all deontologists accept options, however – some think that we have a duty to

maximize the good, whenever we would not be in breach of one of our other duties in doing so.)

So long as we hold that the right action is determined by the amount of agent-neutral value our actions will produce, then consequentialism will be the only coherent picture. Many philosophers have failed to see the power of this argument. They have supposed that deontology can be defended within the basic consequentialist structure, by pointing to some value which consequentialism has ignored. Scheffler (1982, chapter four) is a notable exception: he clearly articulates the point that any value a would-be deontologist might put forward can be accommodated by consequentialism. Scheffler shows, in essence, that, as long as we play the game by consequentialist rules, the consequentialist will win. Any value can be sucked up into what we have elsewhere dubbed 'the consequentialist vacuum cleaner' (or CVC – see McNaughton and Rawling (1991)). To give an example: one might think that, in rejecting constraints – in allowing that there may be occasions when we may do really nasty things to innocent people – consequentialism fails to pay proper respect to persons. The consequentialist can respond by agreeing that failing to respect persons has great disvalue. So we must, all else equal, maximize respect for persons. But now what, the consequentialist asks, of the case where I can prevent two people being treated disrespectfully by so treating one myself? I should surely prefer the less bad state of affairs where only one person is treated disrespectfully. So, the consequentialist triumphantly concludes, no matter how great the value of respect, there is no *constraint* against violating respect *oneself*.

Deontology can only be defended if there is a defensible alternative structure to moral theory. The best prospect for such a defence is classically found, we claim, in the pre-war British Intuitionists, particularly W.D. Ross (e.g., 1965). Ross, in criticising Moore's theory of ideal utilitarianism (which is essentially a pluralist consequentialism), identifies a structural difference between consequentialism and intuitionism which enables the latter to accommodate the agent-relative. Whereas the consequentialist thinks of the right as determined by the good, the intuitionist conceives of the right as being, at least in part, independent of the good. Which action is right is not solely determined by the relative values of states of affairs. Other considerations also play a role.

What are these other considerations? We shall look in turn at

obligations of special relationship and at constraints. Ross accepts that there is a (*prima facie*) duty to maximize the good, and hence does not argue for options (although he does argue that there is a *prima facie* duty to self-improvement); we leave them on one side here.

We begin with duties of special relationship. In a telling paragraph, Ross contrasts his view with consequentialism thus:

[Consequentialism] says, in effect, that the only morally significant relation in which my neighbours stand to me is that of being possible beneficiaries by my action. They do stand in this relation to me, and this relation is morally significant. But they may also stand to me in the relation of promisee to promiser, of creditor to debtor, of wife to husband, of child to parent, of friend to friend, of fellow countryman to fellow countryman, and the like; and each of these relations is the foundation of a *prima facie* duty (1965, p.19).

The various relationships in which we stand to other people generate a variety of agent-relative reasons: one can be required, say, to keep a promise, even if by breaking it promise-keeping as a whole would thereby be increased. On the direct-act consequentialist view, on the other hand, if promises should be kept, that can only be because promise-keeping is valuable, and hence should, *ceteris paribus*, be maximized. Accepting Ross's view of duties of special relationship entails that what is right is not solely determined by considerations of agent-neutral value.

At this point, the consequentialist may rightly claim that no account has been given of why certain relationships are morally basic in a way that does not permit their violation whenever such violation would increase value. But it is hard to see how this is a complaint, unless it is also a complaint that the consequentialist has not explained why morality's sole concern is the increase of value – on the consequentialist view, it is apparently taken to be self-evident that all moral reasons ultimately rest on considerations of agent-neutral value. Justification has to stop somewhere, and it is not clear that the intuitionist's stopping place is less defensible than that of the consequentialist. The thought that others have direct moral claims on us that are not explicable in terms of agent-neutral value is a familiar one in everyday moral thinking. Intuitionism accepts this thought at face value and thereby rejects the consequentialist perspective.

In addition to agent-relative reasons generated by duties of

special relationship, of course, Ross also takes as basic those generated by constraints. He is silent, however, on the subject of whether his line on the former can be extended to cover the latter. There is an obvious disanalogy between constraints and duties of special relationship: constraints do not depend on the specific nature of the relationships we have. I should not treat anyone in the way that constraints forbid, whether I have a relationship with them or not. Note, however, that the usual objection that is raised to justifying constraints – how could it *ever* be that we are *required* not to maximize the good? – has already been met. For Ross's account of duties of special relationship has already shown that we can make sense of this once we get away from the consequentialist picture of value determining what is right. Fulfilling the duties which stem from special relationships may require us not to maximize the good. Others can have direct moral claims on us, claims that are not routed via thoughts about the maximization of the good. So we can already make formal sense of the thought that other agents have a claim on us not to be treated in certain ways; a claim we standardly express by saying that they have a *right* not to be harmed, tortured or killed.

Duties of special relationship, then, create conceptual space for constraints by divorcing the right from the good. But that is not enough to show that there *are* constraints, or just what form they take. Since our concerns in this paper are about the formal structure of moral theory, we shall do no more than indicate a familiar account of constraints – an account that builds on thoughts about human relationship. We can seek an account of constraints in either a potential or an actual relationship. In terms of potential relationships, we might claim that there is a morally significant distinction between the relationship in which I *would* stand to someone if I allowed them to be tortured and the relationship in which I would stand if I were to torture them myself. The latter relationship is one that is to be avoided at (almost) all costs. So constraints proscribe a relationship that would exist were I to act in a certain way. Some prefer to make the point in terms of an actual *moral*, rather than social, relationship. In a common moral tradition, all of us are already related, in terms of our common humanity, or as fellow legislators in the kingdom of ends, or as children of God. On this approach, certain kinds of act would breach that relationship in ways that are unacceptable. (We suspect that these two ways of appealing to relationship are just variants on a theme.)

These observations do not constitute, of course, an argument for deontology. It is simply supposed that the prohibition on treating people in certain ways is basic. This is a demand that the existence of other agents makes upon us. The deontological intuition can be spelled out in greater detail and illustrated, but it cannot, on our view, be justified, in the sense of being deduced from or supported by some more basic intuition. Every moral theory, as Mill justly remarked, will have its fundamental principles. Deontology has a greater number of basic principles than consequentialism, and that might be a ground of complaint, but that it has them at all cannot be held to its discredit.

Ross's main contention, then, is remarkably simple. He claims that the deontic cannot, and need not, be justified by appeal to yet more basic considerations. And, in particular, he insists that the deontic is largely independent of the evaluative: thoughts about what is required, forbidden, or permitted, are not identical to, reducible to, or derivable from, thoughts about what will produce the most good.

II

Thus far, we have characterized deontology in terms of agent-relative reasons of certain varieties (those associated with duties of special relationship, constraints and, perhaps, options). And this characterization serves to contrast deontology with direct-act consequentialism. However, there is more to the story when it comes to contrasting deontology with other forms of consequentialism – a story best told by appeal to the notion of an evaluative possible world.

All consequentialists appeal to value, and can be viewed as thinking in terms of a set of possible worlds ranked in accord with their agent-neutral values. The most naive brand of direct-act consequentialist can be seen as advocating that we ought to aim at the evaluatively best possible world (where this world is at least possibly achievable by our actions). But this doctrine is subject to a difficulty brought to the fore by, for example, Regan (1980, pp.264–265, note one). You are faced with a choice between acts *f* and *g*, where *f* has a 50/50 chance (this might be due to an epistemic barrier; but it might be a raw matter of physics – consider quantum effects) of producing zero or ten (objective) utiles, and *g* is sure to produce nine utiles. The best world is one in which you *f* and ten utiles result; but surely you should *g*.

The sensible direct-act consequentialist advocates that the right thing to do is to maximize expected objective utility (where objective utility is to be identified with objective value – not to be confused with the subjective utility of Bayesianism). Thus she advocates that *g* is the right act, even though there is no chance that *g* will result in the best consequence of ten utiles. The utility maximizing direct-act consequentialist, then, abandons the thought that we should always aim at the best, but not the thought that the right is determined by agent-neutral value.

Even the latter form of direct-act consequentialism suffers the difficulty that we cannot know which act will maximize expected objective utility – real agents will both blunder, and devote far too much time to calculating expected utilities. The satisficing consequentialist responds to these difficulties by requiring the agent only to achieve some threshold amount of good (thus, again the right is divorced from the best, but not from value). Alternatively, there are various forms of indirect consequentialism. In the interests of brevity, we mention only two.

Call a possible world ‘rule-bound’ if there is general acceptance of some set of moral rules in that world. Let *R* be that set of moral rules generally accepted in the best *rule-bound* world (for the sake of simplicity, we assume here that there is only one such). According to the rule consequentialist, we are obligated (barring exceptional circumstances) to obey *R* – even if, in so doing, we might have no chance of achieving the best (simpliciter) world. The right is divorced from the best (simpliciter), but not from value.

The indirect-act consequentialist, on the other hand, standardly has the same criterion of right act as the naive direct-act consequentialist – the right act aims at the best consequence. But the indirect-act consequentialist denies that the virtuous agent need be directly guided by thoughts of the right when deciding how to act. The agent should not generally deliberate over which act is right. Rather, she should often act on dispositions that she has cultivated – and cultivated, in part, with a view to ensuring that her actions achieve the best results in the long run. (See, e.g., Railton, 1984.) In not aiming directly at the best, the indirect-act consequentialist claims to avoid not only Regan’s worry, but also the concerns over burdensome calculations and blundering.

To complete our thumb-nail sketch, we need to distinguish intrinsic from instrumental goods. An intrinsic good is a good ‘in

its own right'; an instrumental good is a good because it brings about (or tends to) an intrinsic good. The following trio of theses are in keeping with consequentialism. First, an intrinsic good at a world is intrinsically good wherever it occurs in that world. Second, whether something is an intrinsic good is outside our influence – the list of intrinsic goods is fixed, we can only influence the amount of intrinsic good in the world. Third, which evaluative world we occupy is within our influence. (This is not to say that we can determine which evaluative world we occupy, merely that we have some influence; but the presence of influence is perhaps enough to ensure that the set of evaluative worlds is a *proper* subset of the set of logically possible worlds.) From this it follows that the list of intrinsic goods is invariant across evaluative worlds (otherwise we could influence what occurs on that list). If, say, pleasure is an intrinsic good, this fact is outside our influence. All we can influence is the amount of pleasure in the world.⁴

Instrumental goods are to be pursued only insofar as they bring about intrinsic goods. Consider, for example, the contrast between deontology and rule consequentialism. The latter, like the former, apparently incorporates constraints (the rule consequentialist insists that you should not normally violate any of the rules in R). However, the elements of R are only what we might dub 'ersatz constraints': acceptance of (and obedience to) them is only an instrumental good, and only an instrumental good at certain worlds. The rule consequentialist can acknowledge that in the evaluatively best world, R is not accepted as universally binding. But the chances of achieving this best world are so small that each agent should accept R herself, and encourage others to do so – in this way expected objective utility is maximized. The rules are instituted by us for the achievement of an evaluative goal; and they do not necessarily hold sway in the evaluatively best world. The deontologist, by contrast, insofar as she accepts the notion of an evaluative world, sees each constraint as built into the very fabric of every evaluative world. Just as the list of intrinsic goods does not vary from world to world, neither does the list of constraints – whether or not something is a constraint

⁴ We do not maintain that all authors concur with this account of intrinsic good. One might argue, for example, that intrinsic goods are not necessarily essentially good (cf Dancy's (1993, pp. 22–26) separation of intrinsically and essentially motivating states). But we are not aware of any consequentialists who take this line.

is completely outside our influence. (Note that the consequentialist cannot complain against the deontologist that the latter is committed to some form of undesirable 'moral necessity' by virtue of this property of constraints – the deontologist would have equal grounds for complaint against the consequentialist when it comes to the list of intrinsic goods.)

Furthermore, suppose the list of constraints corresponds to R. Then the rule consequentialist and the deontologist agree that we should not aim at a world in which the acceptance of R is maximized; rather, we should all accept R as our guide. But we contend that the deontologist, unlike the rule consequentialist, cannot give a value-based argument for this. The rule consequentialist's response to the CVC is to argue that anyone who rejects R herself, in an attempt to maximize the intrinsic good, is engaging in a sub-optimal strategy that will not maximize expected objective utility (i.e., the expected amount of intrinsic good) in the long run. The indirect-act consequentialist and the rule consequentialist share, then, a value-based argument against aiming directly at the best – aiming at the best is a sub-optimal strategy for maximizing expected objective utility. The deontologist, on the other hand, does not appeal to considerations of utility in defence of constraints: she simply does not accept that maximizing expected objective utility is the sole determinant of right action.

Our account is only approximate, and far from complete. But it fulfills the purpose of illustrating how consequentialists invoke a ranking of evaluative worlds, and view the right as some function of the good.

In terms of our discussion of agent-relativity, the following has emerged. We know that, because of the CVC, agent-relative reasons cannot rest *directly* on considerations of agent-neutral value – neither acting on nor accepting agent-relative reasons necessarily maximizes the good. Ross is unmoved by this: why should we take increasing the good as the only duty? There are other moral reasons – and, in particular, other basic agent-relative moral reasons that hold sway in all evaluative worlds, and do not rest on considerations of value. The indirect consequentialist (whether indirect-act or rule), however, can also entertain agent-relative reasons. But she claims that such reasons rest on considerations of value, albeit indirectly: agent-relative moral reasons feature in the optimal strategy for maximizing expected objective utility. In contrast to what many authors (ourselves included)

have argued, then, agent-relativity *per se* is not what distinguishes deontology from consequentialism; at the most general level, we now claim, the distinction rests upon whether or not appeal is made to agent-neutral value in justifying the moral system.⁵

III

Just as there are a variety of consequentialist positions, so there are various deontological views. For example, where Ross declares certain agent-relative reasons simply basic, other deontologists seek some underpinning for such reasons. Kantian and contractualist deontologists, for instance, might be seen as attempting to justify agent-relative moral reasons by appeal to more basic considerations of rationality. If these more basic considerations hold in all evaluative worlds, then agent-relative moral reasons are present in all evaluative worlds, and we have a thoroughgoing deontological view; if not, then perhaps the view is merely ersatz deontology. Since our central focus here is the relation between moral reasons and agent-neutral value, however, we shall discuss neither Kantianism nor contractualism. Rather, we devote the sequel to exploring a value-based defence of deontology advocated by Quinn (1993), Kamm (1989, 1992), and Nagel (1995). If this defence succeeds, it vitiates the claim that value-based justification of the moral system is what ultimately distinguishes consequentialism from deontology.

We shall focus on Nagel's account of the defence (he draws on both Quinn and Kamm). Nagel discusses the case of constraints, but much the same conceptual points can be made *vis-à-vis* options, and duties that stem from special relationships.

Nagel effectively accepts the consequentialist view that a system of moral rules can only be defended by showing that their adoption brings about some good that could not otherwise be realized, and then seeks to show that deontology is such a system. The claim is not, of course, that agent-relative reasons rest directly on considerations of value in a manner obviously susceptible to the CVC; rather, the grounding is indirect – the notion is that worlds in which there are agent-relative reasons are better than worlds in which there are not. Nagel argues that an agent-relative morality, qua moral system, is intrinsically valuable. Thus

⁵ We concur with Hooker (1994), then, *pace* Howard-Snyder (1993), that rule consequentialism is not a 'rubber duck'.

rights (the obverse of constraints) have value, and are, therefore, part of the basic structure of moral theory.

'A right is an agent-relative, not an agent-neutral, value', says Nagel (1995, p.88). This is precisely because it is supposed to resist the CVC (one is forbidden to violate a right even to minimize the total number of such violations). So Nagel faces the Scheffler problem: 'How could it be wrong to harm one person to prevent greater harm to others? How are we to understand the *value* that rights assign to certain kinds of human inviolability, which makes this consequence morally intelligible?' (p.89, our emphasis – note the presumption inherent in the question).

The answer 'focuses on the status conferred on all human beings by the *design* of a morality which includes agent-relative constraints' (p.89). That status is one of being inviolable (which is not, of course, to say that one will not be violated, but that one *may* not be violated – even to minimize the total number of such violations). A system of morality that includes inviolability encapsulates a good that its rivals cannot capture. For, 'not only is it an evil for a person to be harmed in certain ways, but for it to be *permissible* to harm the person in those ways is an additional and independent evil' (p.91). So there is a sense in which we are better off if there are rights (they are a 'kind of generally disseminated intrinsic good' (p.93)). Hence there are rights. In short, we are inviolable because inviolability is intrinsically valuable.

Does this purported justification of constraints work? We detect at least three ways of interpreting the proposal. The first two construals are explicitly rejected by Nagel. Amongst other things, on both these construals, inviolability is something we can bring about, hence it fails to be a genuine right. The third construal correctly asserts that we cannot bring about inviolability – whether or not something is a constraint is completely outside our influence. However, on this third interpretation, the purported justification fails to show that we are inviolable.

The first construal offers an instrumental justification of the constraint forbidding you to violate one even in order to prevent the violation of others. On this account, acceptance of rights in a society will serve to give people a sense of their own worth, and it is a sense of one's own worth that is intrinsically valuable. The general acceptance of the proposition that people have rights is a goal that we may have; we can strive, perhaps through institutions like Amnesty International or the UN, to make that acceptance more widespread. But the ultimate goal here is an increase

in the *perception* of self-worth. There are various problems with this construal. First, we need an independent argument to circumvent the CVC. If what matters is the general sense of self-worth that comes from people accepting inviolability, then why not make that acceptance as widespread as possible, even if that means not accepting it oneself? Second, as Nagel rightly points out, this construal does not yield the conclusion that inviolability itself is intrinsically valuable, rather than the increased perception of self-worth engendered by its acceptance.

On the second construal, it is inviolability that is intrinsically valuable. By not recognising inviolability, consequentialism misses a crucial intrinsic value. Deontology incorporates this intrinsic value, and is thus the preferable theory. This is a reason for adopting deontology as our moral system.

However, that inviolability is intrinsically valuable does not entail that we are inviolable (compare: the absence of suffering is a good at all worlds, but suffering is not absent at all (any?) worlds). And, furthermore, on both the first two construals it is supposed that we can simply *choose to adopt* deontology. We cannot: we have no influence over whether deontology or consequentialism is the case – if, say, deontology is false (humanity is violable), we can do nothing toward bringing about its truth (making humanity inviolable). The inviolability of persons is not dependent upon the kind of moral system that is accepted by a society – our choices cannot bring about inviolability, they can only reflect it. Just as we cannot free ourselves from a constraint by choosing not to respect it, neither can we institute a constraint by choosing to accept it.

On neither of the first two construals is inviolability a feature of every evaluative world (although, on the second, inviolability is an intrinsic good at every world). People are inviolable in some evaluative worlds, and not in others. The former worlds are higher up the evaluative scale; and hence it is the former that we should realize in our actions. But this picture is not coherent: if we have the *right* not to be violated, we have this right in all evaluative worlds; we can have no influence over what rights there are. Or perhaps this is better put: if to be inviolable is to have rights, then to be inviolable is to be inviolable at every evaluative world.

And both construals are vulnerable to a further difficulty. As Nagel admits (e.g., p.90), more people might be violated under a moral system in which people are inviolable, than would be

violated under a system in which they are not granted this status. For in a system where inviolability holds, virtuous agents are restricted in what they may do to frustrate transgressors. But now it is not enough, in order to show the superiority of deontology, that we show that it embodies a value which cannot be embodied in consequentialism. We must also show that it does not do this at too great a cost. This argument envisages a trade-off: we weigh the value of having a moral system which enshrines inviolability against the possible disvalue of having more people violated. Not only is it not obvious which way the judgement would go, or even how it is to be made, but it seems entirely against the spirit of deontology, which this argument is supposed to defend, that the value of constraints be considered in a trade-off in which they might lose out.

Nagel is clear in his rejection of the first two construals: '[The value-based defence of rights] is not supposed to be merely an argument for *creating* or *instituting* rights, through laws or conventions. In a sense the argument is supposed to show that the morality which includes rights is *already true* – that this is the morality we ought to follow independently of what the law is, and to which we ought to make the law conform' (1995, p.92). On this third construal, that deontology is true is independent of human action. Rights are a feature of every evaluative world. But it is unclear what this argument amounts to.

Nagel's attempt to elucidate the value-based defence of deontology is, as he notes, open to serious objections. He argues that, 'we would all be worse off if there were no rights – even if we suffered the transgressions which in that case would not count as violations of our rights – *ergo*, there are rights' (1995, p.92). This is, as Nagel remarks, 'a curious type of argument: [it is not generally cogent to argue] that *p* is true because it would be better if it were true' (p.92). However, Nagel suggests, 'it may have a place in ethical theory, where its conclusion is not factual but moral' (p.92). We demur: it would seem better not to lumber ethical theory with arguments lacking in cogency, especially if the attempt to show that they really are cogent involves claiming that 'there are rights' is a moral *rather than* a factual remark.

To sum up: even if inviolability is valuable, it does not follow that we are inviolable – that a state is valuable does not entail its realisation (for the direct-act consequentialist, for example, that an act maximizes value does not entail that it is performed). On certain (consequentialist) views, it does, of course, follow from

the value of a state that we have reason to pursue it (for the direct-act consequentialist, that an act maximizes value does entail that it *should* be performed). However, this is of no help. At best, all that could be established via this route is that we have reason to make deontology the case. Not only is this not the advertised conclusion, it is also otiose: deontology is not something that can be made true by human agency – whether deontology is true is outside our influence.

Moreover, as Nagel points out, the argument, ‘seems in danger of being circular. For what is the value that a morality without rights would fail to recognize and realize? It seems to be nothing more nor less than the existence of rights, for which ‘inviolability’ is just another name’ (p.92). Nagel’s defence is that ‘[a]ny attempt to render more intelligible a fundamental moral idea will inevitably consist in looking at the same thing in a different way’ (p.92). This is a point made in the Rossian deontological defence. However, the value-based account misses the mark *qua* explication.

Suppose that inviolability is a good, and a good that is in fact realised (we are in fact inviolable). How does the fact that inviolability is a good *explain* its realisation? The explanation cannot appeal to inclinations on the part of humanity to strive for the good, since, as we have already pointed out, no amount of human striving can make deontology true if it is false.

Furthermore, such an explanation would have to meet the following objections. We are asked to contemplate a possible world, *w*, in which we have no rights, and compare it to worlds in which we have rights (the claim is that we are better off in the latter). First: in what sense is *w* possible? For the deontologist, beings without rights are simply not human (or perhaps not even sentient). To draw an analogy: consider a consequentialist who thinks pleasure an intrinsic good – can she make sense of the possibility that this is not so?

Second: even if *w* is possible in some sense, how can we compare it to worlds in which we have rights? In *w* we have no rights, so our rights cannot be violated there. Why are we worse off in a world where our rights are not violated? Arguably, we are *better off* – the premiss that we would be worse off in a world where we lack rights is by no means obviously true. On Nagel’s account, we are inviolable because inviolability is intrinsically valuable – there would be no rights only if inviolability were not intrinsically valuable. But the question then presses: if inviolability were not

intrinsically valuable, how could we be worse off for lacking it? (Compare: if pleasure were not an intrinsic good, how could we be worse off for lacking it?) We might (as Nagel perhaps suggests) try looking at matters in terms of behaviours in *w* that would be rights violations if we had rights. But this fails on two counts. If we lacked rights, the nature of the behaviour would be intrinsically different. And, as already noted, there might actually be less of such behaviour in *w* than in worlds where we have rights.

Alternatively, perhaps, one might try an explanatory account on which inviolability is an *instrumental* good that leads to an increase in some intrinsic good – our inviolability is explained by the resultant increase in this intrinsic good. On such an account it seems, perhaps, that we could compare worlds where we have rights to worlds in which we do not in terms of the amount of this intrinsic good. But the same problems resurface. First, the deontologist cannot make sense of the possibility that we lack rights. Second, what guarantees that inviolability increases the intrinsic good in the world. And, third, even if it does, how does this explain our inviolability?

The attempt either to argue for or to explain the existence of rights by appeal to considerations of value fails. The Rossian account of constraints is to be preferred on all fronts.

Philosophy Department
Keele University
Keele Staffordshire ST5 5BG
England

Philosophy Department
University of Missouri-St. Louis
St. Louis MO 63121-4499
USA

References

- Dancy, J. (1993) *Moral Reasons* (Oxford: Blackwell).
- Dreier, J. (1993) 'Structures of Normative Theories' *The Monist* 76(1), 22–40.
- Hooker, B. (1994) 'Is Rule Consequentialism a Rubber Duck?' *Analysis* 54(2), 92–97.
- Howard-Snyder, F. (1993) 'Rule Consequentialism is a Rubber Duck' *American Philosophical Quarterly* 30(3), 271–278.
- Kamm, F.M. (1989) 'Harming Some to Save Others' *Philosophical Studies* 57, 227–260.
- Kamm, F.M. (1992) 'Non-consequentialism, the Person as an End-in-Itself, and the Significance of Status' *Philosophy & Public Affairs* 21, 381–389.
- McNaughton, D., and Rawling, P. (1991) 'Agent-Relativity and the Doing-Happening Distinction' *Philosophical Studies* 63, 167–185.
- Nagel, T. (1986) *The View From Nowhere* (Oxford University Press).
- Nagel, T. (1995) 'Personal Rights and Public Space' *Philosophy & Public Affairs* 24(2), 83–107.
- Parfit, D. (1984) *Reasons and Persons* (Oxford: Clarendon Press).

- Quinn, W. (1993) *Morality and Action* (Cambridge University Press).
- Railton, P. (1984) 'Alienation, Consequentialism, and the Demands of Morality'
Philosophy & Public Affairs 13(2), 134–171.
- Regan, D. (1980) *Utilitarianism and Co-operation* (Oxford: Clarendon Press).
- Ross, W.D. (1965) *The Right and The Good* (Oxford: Clarendon Press).
- Scheffler, S. (1982) *The Rejection of Consequentialism* (Oxford: Clarendon Press).