

Introduction

‘Deontology’ is a term that is used in a variety of ways. Some definitions have it that deontology is an account of moral duty set out as a system of rules, as opposed to a theory (consequentialism) that assesses the morality of our actions in accord with the impersonal value of their consequences. Nagel (1986: 165, 175–80), on the other hand, gives an account of deontology in terms of reasons rather than rules. And according to him, if deontological reasons ‘exist, they restrict what we may do in the service of either relative or neutral goals’ (1986: 175), where neutral goals are associated with impersonal value, and relative goals with ‘the desires, projects, commitments, and personal ties of the individual agent, all of which give him reasons to act in the pursuit of ends that are his own’ (1986: 165).

Like both these accounts, I shall characterize deontology in part by its opposition to consequentialism. My focus in this regard will be the contrast between deontology and a simple form of consequentialism according to which we should always act so as to maximize value by bringing about the (impersonally) best consequences. Consequentialism comes in many varieties beyond this simple version (see Portmore, 2020), but I shall only address one other (rule consequentialism: see Section 4.6), and I shall not discuss attempts to ‘consequentialize’ all moral theories (see Dreier, 2011; Portmore, 2011). I should note that, while I am not a consequentialist, I do think that the flexibility of the view is sometimes underappreciated, and I shall frequently deploy consequentialist arguments against deontology – there being a distinction between deploying such arguments and arguing for full-blown consequentialism.

As regards particular deontologists, W. D. Ross and Immanuel Kant both more or less fit the bill, despite major differences in their theories. They are united (anachronistically) in their opposition to consequentialism, and both advocate moral rules, albeit of different kinds. I shall discuss both, but Kant to a far lesser extent (his ethics is the subject of Moran, 2022).

Contractualism (see Scanlon, 1998), which has Kantian elements, is a further ethical theory that might be considered a form of deontology (see McNaughton and Rawling, 2006, for discussion; it is also the subject of Suikkanen, 2020). As in Kantian ethics, and in contrast to Ross’s view, contractualism’s moral principles (such as a directive not to kill) derive from more basic considerations. Given that Kantianism has a clearer claim to being deontological, however, I shall use it as a contrast to Ross’s position.

Deontology and consequentialism are two of the main contenders in normative ethical theory, but virtue ethics is another (the subject of Snow, 2020). Virtue ethics is addressed briefly right at the end of this Element (in Section 6.2), with an attempt to see it, in its most plausible form, as a part of deontology.

1 What Is Deontology?

1.1 Deontology and Consequentialism

In this section, I contrast deontology with the simple version of consequentialism mentioned in the previous section, according to which there is exactly one basic and overriding moral principle: we should act, on each and every occasion, so as to maximize the amount of good in the world – to bring about the best state of affairs that we can. (Unless otherwise noted, I shall use the term ‘consequentialism’ to refer to this simple version.) For the consequentialist, then, unlike the deontologist, the good determines the right. And different (simple) consequentialist theories are generated by differing accounts of what is good – utilitarianism (strictly speaking ‘direct act hedonistic utilitarianism’), for example, famously holds that the one and only good is pleasure, and that we should maximize it with each act.

Here is an example sometimes used to contrast deontology with consequentialism. Imagine that you are in a position to quell an angry mob, and thereby save many innocent people, by killing an innocent person yourself. The consequentialist might argue that the good of the innocents saved would outweigh the bad of the one innocent killed, whereas the deontologist might well decry the appalling injustice of killing an innocent person, and contend that you should not kill them, regardless of the consequences.

As a result of this sort of example, some people see deontology as focused upon acts, and consequentialism as focused upon their consequences, where these are seen as distinct from the acts that give rise to them. But this is not the way that the distinction has been drawn more recently. According to Parfit, for example, consequentialism ‘gives us one substantive moral aim: that history go as well as possible’ (1984: 37), where history includes present, past, and future acts, as well as their consequences in the more ordinary sense.

Deontologists, on the other hand, claim that, even if making the world a better place (i.e., making history go better) is one moral goal (some might not accept even this), it is not the only one: the right is not determined by the good. (But note that its denial of consequentialism is only a necessary feature of deontology; other theories also deny it.) For example, some deontologists claim that there is an *absolute constraint* against intentionally killing an innocent person: doing so is always morally prohibited, even if it would produce more good – even if, indeed, intentionally taking the life of an innocent person yourself is the only way for you to prevent many more innocent people having their lives intentionally taken by other agents. The consequentialist, on the other hand, would contend that intentionally killing innocent people is bad, and thus the performance of such acts should, other things being equal, be minimized

(in order to maximize the good by minimizing the bad) – so you may be morally required to kill an innocent person intentionally yourself if that is the only way to minimize such acts by others.

Other deontologists advocate for merely *threshold* constraints, where a threshold constraint is a proscription against certain kinds of acts, but only up to a certain threshold. Advocates of merely threshold constraints might maintain that, while killing the innocent is ordinarily forbidden, in a tragic case in which your killing an innocent person is the only way to prevent a calamity in which thousands of innocents would die, you may be permitted, or even required, to do it. This still contrasts with consequentialism, since on the latter view the threshold may be set considerably lower: the consequentialist might advocate that you are required to kill an innocent person yourself if that is the only way to save considerably fewer than thousands.

Deontology, then, not only denies that we are always required to maximize the good, it may require us not to, as in the case of the requirement not to violate a constraint even if you could maximize the good thereby. *Duties of special obligation* (sometimes referred to as duties of special relationship, directional duties, or positional duties) are another case in point: deontologists typically contend, for example, that you may be morally required to stick by your loved ones, even at significant personal cost, and even if you could maximize the good by abandoning them. Whereas the requirement not to violate constraints forbids us to harm *anyone* in the proscribed way, duties of special obligation are restricted: they apply only to our treatment of those to whom we bear the relationship in question, such as that of friendship, or promiser to promisee.

In contrast to these requirements not to maximize the good, in other cases deontology may merely give us permission – the *option* – not to do so. Deontology can be demanding, but on most versions many non-moral pursuits are permissible, innocent fun being one possibility. It is this that leaves room for *supererogation*. Supererogatory acts ‘go beyond the call of duty’: they are morally permitted and morally admirable (they typically involve self-sacrifice), but not morally required. Deontology may permit you to laze around next weekend, so that you are not required to spend it doing good deeds. Doing so would be supererogatory.

Consequentialism, on the other hand, leaves no room for supererogation: you are morally required to maximize the good, and you can do no more than this – it is impossible to go beyond the call of duty. This relentless requirement to maximize the good with every act also leads to the charge that consequentialism is both over-demanding and pervasive. For example, although the consequentialist may admit that innocent pleasure is a good, there is then

the issue of how, as it were, to fit it in – your having fun is only allowed on the unlikely possibility that it maximizes the good (and then it is required, which perhaps rather spoils it).

1.2 Morality and Practical Reason

Practical reasons are reasons to do things (in contrast to theoretical reasons, which are reasons to believe things). And one question that arises about morality is its relationship to such reasons.

On one picture, consistent with that painted in the previous section, morality is seen as a system according to which acts fall into exactly one of three categories: the morally required, the morally permissible but not required, and the morally forbidden.¹ Practical reasons need not be part of this system. The question of what reason we have to do morality's bidding arises, rather, from outside the system. This view of morality is implicit in the following passage from Singer (Jamieson, 1999: 308–9):

[Nagel and I] were discussing 'Famine, affluence and morality' [Singer, 1972], and Nagel was unable to accept that morality could be so demanding. But eventually it emerged that he was assuming that if morality did demand that we give so much to famine relief, then there must be overriding reason to do so. I was making no such assumption. On my view, I could recognize that if I were totally committed to doing what I ought to do, I would give away my wealth up to the point indicated in my article; but at the same time I may, without any irrationality, choose to be less than totally committed to doing what I ought to do. My own interests, or those of my family, may counteract the demands of morality to some degree, and I may think it reasonable to give in to them, while recognizing that it is morally wrong for me to do so. Once Nagel and I realized that we held these distinct understandings of morality, the practical difference between Nagel and myself over the demandingness of morality became less acute.

On Singer's view (at least as reflected in this passage; he has since changed his mind), then, morality and practical reasons are, to some extent at least, in different camps.

On an alternative view, moral reasons are a variety of practical reason, and a deontologist who appeals to them (on what I'll call the 'deontological reasons approach') might also countenance non-moral reasons. Even though there may be no sharp divide between the two, examples can be provided that are clearly on one side or the other. My reason to choose a peach over an apple – that the former is sweeter – is non-moral. My reason to give to Oxfam, on the other

¹ For the consequentialist, setting aside cases in which two or more acts are tied for being optimal, the second category is empty.

hand – that doing so will reduce innocent suffering – is moral. Or suppose that you promised to repay a debt on Thursday; this fact is a moral reason to do so. What about your reasons to favour your friends and family? Some object to the idea that any of these are moral on the grounds that there is something less than ideal about doing things for friends and family out of a sense of obligation. But, on one account, that is to confuse the issue of reason (in the sense here under consideration) with that of motivation: it's quite possible to do what your moral reasons favour – to rescue your child, say – out of affection.

On this deontological reasons approach, then, we each have a variety of practical reasons, some moral, some not. But do moral reasons always trump non-moral ones? I think not. Imagine you must break a relatively trivial promise if you are to pursue a great career opportunity. Surely you should break the promise (and apologize later). Even though you have most moral reason to keep the promise, this is outweighed by a reason of self-interest to pursue the career opportunity.

But is keeping the promise nonetheless morally required? If what you are morally required to do is what you have most moral reason to do, regardless of whether it is what you have most reason to do overall, then yes. However, it is now unclear what would count as a supererogatory act.

Alternatively, if an agent is morally required to perform some act just in case she has most reason overall to perform it where her 'winning' reasons are moral, then keeping the promise is not morally required. Is it supererogatory? This depends on how trivial the promise, and how great the career opportunity. The more trivial the former, and the better the latter, the more it looks as if it would be stupid rather than supererogatory to keep the promise. So, on this approach, it seems (roughly) that, in acting supererogatorily, you act on comparatively strong moral reasons, even though you do not do what you have most reason to.

This means, of course, that a supererogatory act is contrary to practical reason, in the sense that you have more reason to do something else. And this may run counter to intuition. However, it is a consequence of going beyond thinking of morality as merely a system of requirements, according to which a supererogatory act is morally admirable and permitted, but not morally required, no mention of reasons being necessary.

How does a reasons-based account of matters sit with consequentialism? If the consequentialist were to separate morality from practical reason, as in the passage from Singer earlier in this section, then, of course, the consequentialist could countenance all sorts of practical reasons, but the moral requirement to maximize the good would remain, and a key question would be whether you have most (or any) reason to do this.

Alternatively, the consequentialist could begin with the thought that what an agent has most reason to do is maximize the good. Consequentialism could still countenance practical reasons with a variety of contents, but, since it ranks acts in accord with the value of performing them, the combined net strength of an agent's reasons to perform an act would be proportionate to the amount of good she would do by performing it. That is, although the content of practical reasons need not reference the good on this consequentialist approach, reason strength would correlate with value, and value alone. And in performing the act for which she has the strongest reasons, the agent would be maximizing the good, and supererogation would be ruled out.

As regards contrasting the other components of deontology with consequentialism on a reasons-based account, constraints correspond to what I'll call 'constraining reasons'. Whereas the consequentialist advocates the maximization of value, the deontologist who claims that there are constraining reasons thinks there can be cases in which, for example, although it would maximize value for you to kill an innocent stranger, you have a stronger (constraining) moral reason not to. On this way of putting things, the advocate of an absolute constraint against killing innocents claims that the strength of your reason not to kill them is always infinite, whereas the advocate of a threshold constraint against such killing contends that there are finite constraining reasons: although the strength of your reason not to kill is finite, it can exceed the reason strength corresponding to the value of your not killing.

Duties of special obligation and options are captured in similar vein: you can have more moral reasons to favour those to whom you bear special relationships, or to engage in non-moral pursuits, than the value of your doing so would warrant.

On this approach, then, deontology is characterized, in part, by denying that reason strength is correlated only with value (see also Section 1.6) – you might have most reason, even most moral reason, to do something that is worse (in value terms) than some alternative.

1.3 Common-Sense Morality

Normative ethical theories can be classified as either accommodating or revisionary. Accommodating theories attempt to render ethical verdicts that accord with our 'common-sense' pre-theoretic moral judgements. Deontology can be seen as more accommodating in this sense than consequentialism – constraints, duties of special obligation, and the possibility of supererogation are all claimed to be components of 'common-sense morality', as is the idea that morality leaves room for non-moral pursuits.

Some consequentialists are radical reformers, and they argue for the overthrow of our common moral intuitions. But others might see these intuitions as playing an evidentiary role. Given this, one issue worth exploring is the extent to which the consequentialist can accommodate them. And, furthermore, for supporters of common-sense morality, it would be useful to know to what extent the consequentialist is forced to diverge from their common-sense views.

The consequentialist can argue that her theory is flexible, and that it can at least mimic much of common-sense morality – people may only believe in constraints, for instance, because they haven't considered unusual cases, and, in ordinary cases, violating constraints fails to maximize the good. In addition, the consequentialist can contend, in revisionary mode, that the deontologist's refusal to violate a constraint herself in order to prevent similar, but more numerous, violations by others can only be explained by a self-serving preference to keep her own hands 'clean'.

For instance, consider again the case in which you face the decision of whether to kill an innocent person to quell a riot, in order to save many other innocents. Some deontologists would prohibit such killing, invoking a constraint, whereas some consequentialists would advocate that you do it, in order to maximize the number of innocents saved. But they need not. Rather, the consequentialist could agree that it would be very bad, and, indeed, wrong. The difference with deontology, on this more sophisticated consequentialist account of the matter, would be evident in far-fetched circumstances in which your killing an innocent person to quell a riot is the only way in which to prevent many other people doing the same (killing innocents to quell riots). A deontologist who sees killing innocents to quell riots as a constraint violation would tell you not to kill here, whereas the consequentialist who thinks that killing innocents to quell riots is bad would tell you to minimize it, even if that requires doing it yourself – keeping your own hands clean will result in much more of the behaviour that you abhor. In ordinary circumstances, the practical upshot of the two proposals is, of course, the same: don't kill innocents to quell riots. But distinct prescriptions can result, as we have just seen, in far-fetched situations. This consequentialist strategy (dubbed 'the consequentialist vacuum cleaner' by myself and my frequent collaborator, David McNaughton) will recur often.

1.4 Consequences and States of Affairs

Some philosophers see deontology as differentiated from consequentialism in part by the latter's appeal to the states of affairs that agents bring about through their choices:

In contrast to consequentialist theories, deontological theories judge the morality of choices by criteria different from the states of affairs those choices bring about. (Alexander and Moore, 2016: 3)

But this need not be. Both theories can be seen as concerned with the states of affairs we bring about through our actions. If there is an absolute constraint against intentionally killing innocent people, then you are prohibited from doing this – that is, you are prohibited from bringing about a state of affairs in which you intentionally kill an innocent person. And the consequentialist, her label notwithstanding, can, as we saw above, include an act within the state of affairs it brings about – indeed, the value of an act can be affected by events in the past. Consequentialists might, for instance, view the breaking of a promise, in itself, as bad, regardless of the future effects of doing so – a past commitment and its current violation can both contribute to the badness of a state in which this violation occurs. States of affairs can be viewed as histories, temporally extended into both the past and the future. Consequentialism can then be seen as imposing the moral requirement that we bring about the best temporally extended state of affairs that we can.

It is a mistake, on this broad account, to structurally assimilate all forms of consequentialism to one of its best-known forms, hedonistic utilitarianism, according to which pleasure is the only good, and we should maximize it. This view is only forward-looking: what you do today cannot influence the amount of happiness in the world yesterday. But consequentialism need not be seen as entailing that an act has a distinct temporal end point, with its consequences being confined to what it causes after this end point, and being in no way backward-looking.

Indeed, if consequentialism's notion of a consequence is narrowly construed in this latter way, consequentialism is placed under various unfortunate burdens. For example, if a consequence must be only forward-looking, then there is the danger of a puzzle arising when it comes to describing one and the same consequence in two different ways: one forward-looking and one backward-looking. Suppose a consequence of your mailing a letter is Fred's receiving a check. And suppose also that this event (Fred's receiving a check) constitutes repayment of your debt. The first description of this event is forward-looking, but the second description of that exact same event (your repaying Fred) is backward-looking. So is the event a consequence of your mailing the letter or not?

This puzzle does not arise on the broader construal of consequence, and this construal also enables the consequentialist to respond to two well-known deontological arguments against her position. Some deontologists argue that the distinction between doing and allowing is morally significant (see Section 2.2), as is that

between the intended effects of our actions and their unintended but known side effects (see Section 2.1), but that the consequentialist cannot account for this. For example, whereas a deontologist might see euthanasia as wrong, since it is an intentional killing, she might view as permissible a doctor's allowing a terminally ill and suffering patient to die. And she might claim also that a doctor's administering a lethal dose of morphine to such a patient is permissible, provided that the doctor's intent is to alleviate suffering, even though the doctor knows that death will ensue as a 'side effect'.

On the narrow construal, consequentialism cannot draw these moral distinctions: the same narrow consequence, death, results in all three cases. So the 'narrow' consequentialist must argue either that these distinctions are not morally significant or that they cannot be adequately drawn. The 'broad' consequentialist, however, has a third way out, namely: she can point out that in doing versus allowing, or intended outcome versus unintended side effect, the histories are different, and thus, potentially, whatever the deontologist sees as accounting for the morally significant differences between the components within these two pairs, she can adopt as contributing differentially to the values of the relevant histories.

One response to either consequentialist approach, however, might be to challenge the very idea that states of affairs can bear value. Geach (1956: 41), for example, might be interpreted as claiming this:

'Event', like 'thing', is too empty a word to convey either a criterion of identity or a standard of goodness; to ask 'Is this a good or bad thing (to happen)?' is as useless as to ask 'Is this the same thing I saw yesterday?' or 'Is the same event still going on?', unless the emptiness of 'thing' or 'event' is filled up by a special context of utterance. Caesar's murder was a bad thing to happen to a living organism, a good fate for a man who wanted divine worship for himself, and . . . a good or bad act on the part of his murderers; to ask whether it was a good or bad event would be senseless.

Geach makes his claim regarding events rather than states of affairs, but perhaps his complaint applies equally to the latter (assuming there is a difference): that dubbing a state of affairs 'good or bad would be senseless', because 'state of affairs' is 'too empty a word to convey . . . a standard of goodness'. Geach may be correct that states of affairs are, strictly speaking, neither good nor bad simpliciter. But that does not preclude there being a scale of value with states of affairs being better or worse than one another: for instance, a state of affairs in which innocent people flourish is better than one in which they suffer.

Another response is to agree that states can be better or worse, but maintain that it is no part of our moral duty to produce better ones. Ethical egoism may be

interpreted in this way. According to this theory, each of us has exactly one ‘duty’: to maximize our own well-being (or happiness). States of affairs can not only be ranked in terms of whether they are better or worse than one another simpliciter, but also in terms of whether they are better or worse *for individuals*. Ethical egoism is concerned only with the latter rankings, and urges each of us to promote states in which we do best for ourselves, regardless of whether this makes the world better or worse. Note that good for individuals has to be distinguished from good simpliciter, lest ethical egoism be incoherent. Moore failed to make this distinction in his (flawed) argument to the effect that ethical egoism is incoherent:

What Egoism holds, therefore, is that *each* man’s happiness is the sole good – that a number of different things are *each* of them the only good thing there is – an absolute contradiction! No more complete and thorough refutation of any theory could be desired. (1966 [1903]: 99)

(Moore did later recant this claim of incoherence (1912: 99–100).)

With the distinction between well-being and the good in place, however, ethical egoism can be seen as at least coherent, and the egoist can acknowledge that some states are better than others, but she does not see this as relevant to what she should do – all that is relevant for that is which state would benefit her the most.

Value and the good, as I shall use these terms, are, then, impersonal – or, more strictly (since, strictly speaking, it makes no sense to declare a state good or bad), states can be impersonally better or worse than one another. This is in contrast to one state being better *for someone or something* than another. Both measures are objective, but they measure different things. My unqualified use of ‘good’, ‘value’, ‘better’, and ‘worse’ will refer to the impersonal measures.

In line with their embrace of options, deontologists typically acknowledge that you have some reason to pursue your own well-being, but that is a far cry from the ethical egoist’s claim that this is the only thing you have reason to pursue. And most deontologists also acknowledge both that states of the world can be better or worse, and that it is our moral duty to make the world a better place. But they deny the consequentialist claim that this is the only duty. As we shall see, for example, Ross thinks that we have a duty to ‘produce as much good as possible’ (1930: 27). But this is not our only duty; it has to be weighed against others (see Section 3.1).

1.5 Kant and Ross

So far I have focused upon deontology as contrasted with consequentialism. But deontology can also be characterized independently. Immanuel Kant and W. D. Ross may be seen as doing this.