

Why I am not a Consequentialist

David S. Oderberg

This is an introductory talk on why I am not a consequentialist. I am not going to go into the details of consequentialist theory, or to compare and contrast different versions of consequentialism. Nor am I going to present all the reasons I am not a consequentialist, let alone all the reasons why you should not be one. All I want to do is focus on some key problems that in my view, and the view of many others, make consequentialism a totally unacceptable moral theory – a theory about what is right, what is good, or obligatory, or forbidden, or permissible, or praiseworthy.

So let me begin by giving a basic definition of consequentialism, one all supporters of the view can agree on. Consequentialism is the theory that the fundamental aim of morality is to maximize value. Now I was tempted to say ‘sole aim’, but some consequentialists will disagree with that. They might hold, for instance, that one of the aims of morality is to abide by certain rules, or to cultivate certain virtues. But for them, what gives obedience to a rule or the cultivation of a virtue its point is that, ultimately, such behaviour maximizes value. So although maximizing value might not be the sole aim of morality – the sole answer to the question ‘What should I do to be good?’ and similar questions – still it is the fundamental aim of morality, and all other kinds of decision, action, and so on, derive their justification by reference to it. For my purposes, then, the difference between ‘sole’ and ‘fundamental’ is merely terminological.

Now the first thing that might occur to someone is a pair of simple questions. Why should anyone believe that the fundamental aim of morality is to maximize value? What intuitive force does the idea even have in the first place? These are good questions. I was a consequentialist once, and I don’t think I ever posed them to myself. I just took it as understood that since so many philosophers were consequentialists, and since so many of my fellow students were as well, then even if it were ultimately shown not to be true, the maximization thought (as I will call it) was at least the obvious place to *start* when one did ethics. Consequentialists rarely try to justify the thought themselves, but I suppose it goes something like this. Value is, tautologically, a good thing; now, to put it colloquially, you can’t have too much of a good thing; so you should want to produce as much of it as possible; but so should everybody else; so everybody should produce as much of it as they can.

We are talking, of course, about moral value, or intrinsic goodness, or some such – not about money or fast cars. The idea is that whatever it is that is in and of itself the ultimate good – or the ultimate goods – we should aim to make as much of it as we can; not for ourselves, but for the world at large. That’s what it is to be moral. Moreover, consequentialists often think, isn’t this just what God does? Doesn’t God (or wouldn’t He, if He existed, as the non-believing consequentialist will say) act to produce maximum value? In other words, doesn’t He act to produce the best possible state of affairs at every time of the universe, and across the whole history of the

universe as a whole? Furthermore, don't we believe that the aim of morality is to do good and avoid evil? And isn't doing good just the same as producing maximum value, and avoiding evil producing minimum disvalue?

To take the last point first, it is not at all obvious that doing good just is maximizing value. Now I think the whole idea of maximizing value in the consequentialist sense is incoherent, but suppose we consider a case where, in an intuitive and non-technical sense, it is better to produce more rather than less. Suppose I see someone begging on the street for money to buy food. I have in my pocket two 5 euro notes. I don't need them for anything urgent, and I can be sure that if I give the beggar money he will indeed spend it on food. If I give him one note he will have enough food for one day, and if I give him both then he will have enough for two days. In such circumstances, assuming no other countervailing factors, and assuming it is a good thing to give alms – on which we can all agree – I should give him 10 euros. In such a case the mere quantity of good that I do counts, moreover it counts for everything. I should do more, rather than less, and so produce the better of two possible states of affairs that I can produce.

I now ask: how does it follow from this sort of case that doing good just *equates* to maximizing value? Even if we give 'maximization' a wholly non-technical sense, and factor out all the issues that make the technical concept of consequentialist maximization so problematic – a serious problem I cannot examine here – it simply does not follow that doing good is to be identified with producing the best state of affairs, or the maximum value. Now you might think that the reason I say this is that I have also had to factor out all sorts of other considerations, such as whether I have already promised the money to someone else, whether the beggar will use the money to buy drugs, whether by giving him the money I will make myself late for an important meeting at work, and so on. But this is not my reason, since of course the consequentialist will interpret all of these factors in maximizing terms as well. If, for instance, I have promised the money to someone else, then all things being equal I should keep my promise since that will produce more value, or promise-keeping in general produces more value than always giving money to beggars, and so on.

No, my objection is more fundamental. It is simply that the mere existence of – admittedly highly artificial – cases where what I should do is produce the most value does not entail that doing good just is the same as producing the most value. Even if the consequentialist were able – which he is not – to interpret all the factors that might have been relevant to the beggar case in maximizing terms, this would give no intuitive support to the maximization thought. But what if he could interpret *all* possible cases of doing good as cases of maximizing value? That would give not just intuitive but theoretical support to consequentialism, on the assumption that we all agreed on what the examples of doing good are. It is well known, however, that consequentialists and non-consequentialists do *not* so agree. Or when they do, the friendly consequentialist has to interpret the interesting cases, where unfriendly consequentialists disagree with him and with non-consequentialists, in terms that are so distorted, misleading, and unfaithful to non-consequentialist intuitions that the

project is a failure. I will give a typical – and one of the better – examples of this below.

From this I conclude that there is at least no obvious intuitive support for the idea that doing good just is maximizing value. Which brings me to the second point I raised above on behalf of the consequentialist, about producing as much good as possible – *promoting* the good, as it is sometimes called. This is really a point about our *attitudes* to the good. The consequentialist thinks that our sole attitude to the good must be one of promotion: everything subject to moral evaluation must, directly or indirectly, be justified or condemned in terms of whether the good is promoted. My questions are: why should anyone think that this is the only attitude one should have toward the good? Why should anyone take it to be a priori plausible that since, tautologically, value is good, the only right and proper attitude to it must be to promote it? Being kind to people is tautologically good, but why would anyone think that the only attitude I should have to people is one of kindness? What about attitudes of sternness, inflexibility, intolerance, resistance, dislike? A parent should sometimes be stern with their children; a teacher should sometimes be inflexible with a student; we should be intolerant of racists; we should resist drug dealers; and so on. We are not always being kind in such cases, unless ‘kind’ is interpreted so loosely and broadly as to mean something like just ‘doing good’.

Suppose my friend comes to me and says: ‘If you will give me advice on how I can ingratiate myself with my boss so I can get a promotion, I’ll be eternally grateful.’ Suppose the ingratiating behaviour he has in mind is not minor, but a truly obsequious and demeaning kind of behaviour that is undignified and shameful. I ought to say to my friend: ‘I respect you too much to advise you on how to do that sort of thing’ – even if by doing so my friend will get the promotion of his dreams, a huge salary rise, and bliss for the rest of his life. I ought to have too much respect for my friend to help or let him do *that*. On the contrary, I should advise him strongly against doing it. By analogy, our attitude to the good itself is like this, and so it should be. The good does not demand just promotion from us, whether we think of this in terms of consequentialist maximization or in some looser sense such as ‘producing good states of affairs’. It also demands from us *respect*, an attitude of non-hostility or non-aggression, as we might call it. And respect, as non-hostility and non-aggression suggest, has a negative aspect as well as a positive one. When you think about the good, it is intuitively clear, and so quite contrary to consequentialist ‘intuition’, that one should respect it as well as promote it, encourage it, bring it about, even maximize it. The principle of doing good and avoiding evil is as much about respect as about other attitudes. And respect means, as I intimated, sometimes *forbearing* from doing certain things that undermine or attack the good. If the consequentialist thinks respect is *not* a proper attitude to the good, every bit as proper as promotion, it seems to me the onus is on him to explain why.

As to whether God is a consequentialist, there is a lot one could say about this, but I will briefly respond that the very idea is nonsensical. God is no more a consequentialist than He is a virtue theorist, a deontologist, or a natural law theorist. He does not abide by some moral theory, nor does He believe in one. But doesn’t the

consequentialist simply mean that God maximizes value? After all, don't theists believe that God created the best of all possible worlds, and that all of His interventions in the world are aimed at making sure the best possible state of affairs obtains in the end, even if it does not obtain right now or at some arbitrary period of the history of the universe?

Leaving aside the difficult question of whether God created the best possible world, let's simply assume He did. Let us assume, further, that the world He created does indeed maximize value. How does it follow that *we*, in order to act morally, must maximize value? What is the argument for saying that because God maximizes value, so should we? Well, replies the consequentialist, shouldn't we imitate God? But why does imitating God mean doing everything God does? If you have good parents then you should imitate them, but it does not follow that you should do everything they do. But, rejoinders the consequentialists, shouldn't you imitate them by doing good? That's true enough: they do good, and you should imitate them by doing good. God does good, and I should imitate Him by doing good. Then, says the consequentialist, since God does good by maximizing value, so should you. This, I reply, is wrong. God does not do good *by* maximizing value. Rather, He does good, *and* it is the case that He has arranged the world – so we are assuming – in such a way that it maximizes value. It does not follow that every (or any) time He does good it is *in order* to maximize value, whether directly or indirectly. So, for instance, He punished Cain for murdering Abel. Why did He do that? Because Cain deserved it. There is no further reason; rather, we have no reason at all to think that God punished Cain *because* this would help to produce the best possible world. Rather, He punished Cain for doing wrong *and* He has arranged the world so that this punishment is one part of the best of all possible worlds.

Yet, answers the consequentialist, isn't creating the best of all possible worlds one part of the good that God does? Well, maybe it is, but part of the good your father does might be to manage the household finances. It does not follow that you too should manage the household finances, or help to do this, first because you and your father can't *both* do it and secondly because it's none of your business anyway, so you shouldn't even help. But I am not aware that God had any trouble creating the best of all possible worlds, nor that He is experiencing difficulty maintaining it. So none of us has any reason to help Him along, and good reason not to.

Furthermore, it is not at all clear that, even if God created the best of all possible worlds, such a world is one that maximizes value. Let us suppose that all God is interested in is rewarding the righteous and punishing the wicked. (In fact this is one part of what He is interested in, but I want to keep it simple.) The best of all possible worlds, then, will be one in which everyone righteous is rewarded and everyone wicked is punished. Such a world will not necessarily be one that maximizes value, since there could have been a world with one more righteous person in it and one more wicked person, each of whom was rewarded and punished respectively. Or there might have been a world in which one of the actual righteous people who was rewarded did a bit more good and had a greater reward, and one of the actual people who was punished did a bit more evil and had a more severe punishment. In terms of

pure value – assuming, as I am, that we can make sense of the idea – such worlds would contain more of it than the actual world. But the failure of God to create such worlds does not mean that God has not created the best of all possible worlds. Such worlds, it is reasonable to say, would have been equally as good as ours (there can be equal best worlds) even though God did not maximize value in ours. To insist that our world is not, then, the best of all worlds would merely be to object to the fact that God is not a maximizer, and to maintain that the best of all worlds could *only* be one that maximized value. But the consequentialist needs an argument for this and must not beg the question in favour of his own theory.

I have given a number of fairly abstract reasons why consequentialism is on the face of it unintuitive and unmotivated. But I also think it is straight out false, and not only false but an evil and dangerous theory – a view I am not alone in holding. There are a number of ways in which I could defend the view, but I want to focus on one in particular, the one that has always seemed to me – at least ever since I stopped being a consequentialist – the most damaging. It is a very familiar objection, but no less persuasive for being well known. This is the charge that consequentialism allows, indeed requires, certain kinds of action that are *obviously* wrong and so not to be done. In particular, consequentialism permits and requires actions that are horrendous evils, as evil as anything can be. The typical example often given is of the judge who condemns an innocent man to death in order to satisfy a rioting mob that will murder hundreds of people if the judge lets the innocent man go free. Another is the doctor who kills patients for their organs so he can transplant those organs into many other patients who need them. In general, according to consequentialism, it is at least permitted, often obligatory, for a person to commit what looks to any sane observer like a blatant and serious violation of someone else's rights, and hence to commit an act of grave injustice, in order to maximize value, or at least to do what he thinks is likely to maximize value. Now, for the non-consequentialist, no intuition his opponent can bring to bear in support of the consequentialist position on this matter is as strong as the intuition that such apparent injustices are indeed injustices, and so to be forbidden on all occasions, no matter *what* the consequences. According to Elizabeth Anscombe, even to *entertain* the supposition that the judge is allowed, let alone required, to condemn the innocent man to placate the angry mob is to show evidence of a morally corrupt mind. Someone who thinks the issue debatable, she says, is not someone with whom you should enter into debate.

Consequentialists have various responses to what I will call the injustice objection. A well-known one is to appeal to rule consequentialism. Act consequentialists believe the judge should condemn the innocent man if it promotes the good, since people should maximize value on each and every occasion in which they act. Rule consequentialists, sensitive to the many problems of act consequentialism, including the injustice objection, hold that we are obliged to follow moral rules that are part of a system of rules that itself maximizes value. Doing so does not entail that each and every action in accord with such rules will itself maximize value. So, they will argue, a rule which allowed or required judges to condemn innocent people, whether to maximize the good or for any other reason,

would itself not be a rule that maximized value. Following such a rule would lead to a breakdown of trust in the legal system, corruption, social disorder, and so on. Alternatively, such a rule could not coherently form part of a system of rules that maximized value, for similar reasons.

The rule consequentialist is generally right about the consequences of such a rule. But is he always right? Could there not be a society ordered by a system of rules that included rules allowing or requiring punishment of the innocent, torture of the innocent, oppression of minorities, slavery, theft, bribery, and so on? Haven't such societies actually existed throughout history? But, says the rule consequentialist, value is never maximized in societies abiding by such systems. Yet how can he be so sure? On the one hand, he could say that the injustices in such societies are themselves so horrendous that no good consequences could outweigh them. But then the rule consequentialist would himself be operating with a *non*-consequentialist conception of what is good and bad: for that is just what the non-consequentialist holds, namely that certain acts are so wrong that no amount of good consequences can outweigh them.

On the other hand, the rule consequentialist could take such a society at face value: it could be a well ordered, highly stable society offering the best situation for the vast majority of its citizens, and as such the system of moral rules they obeyed would maximize value even though a minority of innocent people – maybe a tiny minority – were subject to the harshest, most unjust treatment imaginable, treated perhaps as nothing more than human cattle by the majority. Now such societies are clearly realistic, very arguably historical, and at least easily conceivable. This is why some rule consequentialists insist that you have to build into rule consequentialism extra principles protecting, say, liberty, or fairness, or justice. To do so, in my view, is to give the game away: it is effectively to abandon consequentialism. Furthermore it is ad hoc. Why stop at justice or liberty? Why not introduce many more non-consequentialist rules, rules against certain kinds of action that, if performed, might nevertheless maximize value?

A rule consequentialist might object that even if some societies have existed in which evil systems of rules governed the citizens, or even if such societies are conceivable, such systems could never govern humanity as a whole. But why should we be so confident that this is so? Could we not easily imagine a world in which people behaved just as they did in the particular unjust society? We certainly do not know a priori that an unjust system governing the entire world would not maximize value even if it might in individual societies. For all we know, such a system could be globally workable and value-maximizing. Yet it is wrong to say that for all we know, torturing the innocent might turn out to be permissible. Metaphysically speaking, it is contingent whether such a system is globally workable and value-maximizing. But it is not contingent whether torturing the innocent is permissible.

Leaving the specifics of rule consequentialism to one side, and looking at more general consequentialist responses to the injustice objection, there are a number of things they can say (and have said). One is that it is question-begging to call, say, torturing the innocent unjust, or condemning the innocent man to appease the angry

mob. We should, instead, focus on the overall picture, in which the good is promoted by what appears to be an unjust act. What would *really* be unjust, according to the consequentialist, would be for the judge to fail to save the hundreds of innocent people who would be killed if he did not condemn the innocent man. It would be a culpable omission. There are various non-consequentialist replies to this, but one is: why should anyone believe that condemning the innocent man is not unjust *because* it maximizes value unless one *already* believes that maximization is the fundamental aim of morality? But as I argued earlier, there is no reason why one should believe this, hence no reason to think that maximization can convert an otherwise unjust act into a just one, or an evil act into a good one. The consequentialist has to tell us why this is so without simply presupposing that promoting the good is all that matters in morality – and not also respecting it.

Another consequentialist response is to argue that the judge is in some way responsible for the death of hundreds of innocent people if he does not condemn the single innocent man. He could have saved them, but he chose not to. Moreover, would he not regret their deaths? This sort of reply is superficial and unconvincing. The judge might regret the death of the hundreds, but not (at least if he is a non-consequentialist!) regret having brought about the death himself. He would regret having been put in such a situation, having had to make such a terrible decision, and so on: but it doesn't follow that he would, or should, consider himself responsible for the death of the hundreds. It is the angry mob who should regret that. Why should anyone consider the judge responsible? Because, it might be replied, he could have saved them but he didn't. Had he condemned the individual, he would have saved the many. Well, there are numerous things we could do but do not, yet we are not responsible for them all. You could have been on a plane to the Arctic right now and on the way to saving a starving Eskimo. Maybe I could even name the Eskimo for you. But you are scarcely responsible for that individual's death. To assign responsibility based on what one *could* have done, as opposed to what one *should* have done, is a serious mistake; to think otherwise is to have a deficient understanding of responsibility, and maybe also to presuppose that condemning the innocent man is what the judge should do. But the consequentialist cannot justify the condemnation of the innocent man by appealing to the judge's responsibility for not doing so, and then justify the claim of responsibility by appealing to the judge's obligation to condemn the innocent man. Yet why else should we regard the judge as morally responsible?

Another consequentialist response is somehow to incorporate respect for rights and justice into the theory, so as to accommodate non-consequentialist intuitions. This is not an easy task. One well-known attempt comes from Philip Pettit, who argues that the consequentialist should want to maximize *dignity*, but that dignity is such a fragile and, as he terms it, 'elusive' feature of human relations, that infringement of the rights of another, even on a single occasion, would not maximize dignity. The reason is that it would put people in fear of being similarly treated by the infringer, and place the whole host of intimate human relationships at risk of sacrifice on the altar of maximization. A person might try to hide his infringement of rights, so as not to endanger his other relationships, but any hesitation or reservation by the person

thinking as a consequentialist would easily be found out by others. If he plays games with one person, infringing their rights and trying to hide it, others will soon know about it and expect the same treatment. Human relations would be undermined.

The non-consequentialist has several responses to Pettit. First, why should he have chosen 'dignity' as the value which consequentialists should promote? More precisely, what could he mean by 'dignity' in this context? In Pettit's sense, it seems to mean simply a respect for others that does not place stronger claims on the agent than consequentialism itself allows: in other words, the agent must not play games in his relationships with others because of the high likelihood of detection and the subsequent weakening of relationships detection would bring about. This is clearly a narrow sense of what dignity is, since it merely reflects the consequentialist idea that human relationships are never good *for their own sake*, but for what overall good they produce. Even if detection is not merely highly likely, but absolutely certain, the fact remains that, qua consequentialist, the agent does not see friendship, or love, or loyalty, or justice, as intrinsically good, but as good for what they can produce in the world as a whole. The notion of dignity Pettit is employing, then, is consequentialist in its very essence, and so his proposal amounts to no more than trying to take account of rights by presupposing a consequentialist understanding of them. What the consequentialist must maximize, namely 'dignity', is simply that order of human relationships which is consistent with consequentialist thinking in the first place.

Suppose, on the other hand, Pettit were to say he was appealing to a stronger notion of dignity, the non-consequentialist one, according to which the moral claims one person has on others are not grounded in what respect for those claims can procure as an overall benefit, but in the inherent and unassailable dignity of the claim-holder himself as a human being, apart from all other considerations. Now if *that* is the notion Pettit is relying on – and some of what he says suggests that it might be – then his use of it will not work. The non-consequentialist concept of dignity does not ground the value of human relationships involving love, friendship, loyalty, justice, and the like, in what overall state of the world they can produce. It therefore does not demand condemnation of injustice because of the fear or instability this will create in others with whom one is (or is not) intimately related, even if that fear or instability are a *certain* consequence. The non-consequentialist sense of 'dignity', then, involves a recognition *in the first place* that rights are not a mere ingredient in a general consequentialist decision procedure, and so appeal to it by Pettit would undermine his very project of reconciling consequentialism and rights.

Secondly, Pettit is far too sanguine about the likely disruption to human relationships or social order from a policy of consistent infringements of the rights of some for the benefit of others. To echo my earlier objection to rule consequentialism, it is not hard to imagine a society in which the vast majority were secure in their 'dignity', in Pettit's sense, but in which a minority were subject to routine injustice. The majority would have no reason at all to expect their control over their lives, or the stability of their relationships, to be vulnerable to a state policy of infringing the rights of, say, a well-defined minority of the population; maybe an ethnic minority, a religious minority, a minority of vulnerable, defenceless individuals. If the majority

benefited tangibly from such a policy, they might even approve of it; but even if they did not approve, for all their condemnation they could be convinced that *their* rights were not vulnerable. And neither scenario depends on the policy's being *hidden* from anyone, so the problem of hiding the injustice, raised by Pettit, does not arise. Since maximization is what consequentialism requires, 'dignity' could well be maximized in such a society, contrary to Pettit's supposition that it never could. Again, it is arguably the case that such societies have existed as a matter of historical record; but even if they have not, given the fact that whether a policy maximizes value is always a *contingent* matter, whereas the inviolability of rights is not, consequentialism and rights are conceptually irreconcilable. It is enough that we can *conceive* of a world containing beings who are recognizably human, with human needs and wants, and in which the consequentialist justification of rights no longer guarantees the integrity of such people in all cases.

There is more that can be said about Pettit's and other attempts to give a consequentialist defence of rights and justice. All such attempts will fail, simply because the structure of rights and justice and the structure of consequentialist thinking are contrary to each other. The only way the consequentialist can reconcile them is by *reinterpreting* rights and justice in terms that are explicitly or implicitly consequentialist. For the non-consequentialist, this would not merely be a case of begging the question, but of discarding at least some of the non-negotiable intuitions we have about rights and justice, such as the inherent and absolute injustice of knowingly condemning an innocent man to death. On the other hand, if the consequentialist wants to preserve such intuitions, he must incorporate into this theory principles that are themselves non-consequentialist. But if he does this, there is no logical stopping point short of abandoning consequentialism altogether.

I have discussed only some of the many problems with consequentialism. They should be sufficient to convince any rational, clear-thinking person who wants to know what morality requires, how he and everyone else should act, and who has not already been corrupted by consequentialist thinking itself, that it is not only lacking in justification and intuitively implausible, but that consequentialism is also downright false and dangerous, an evil doctrine that should be avoided by all right-thinking people.

This is the text of a talk delivered at the University of Lisbon in 2007.