

ARTICLE

Altruism and Desert

Sean Clancy

East China Normal University

Corresponding author. Email: clancy.philosophy@gmail.com

(Received 15 January 2019; revised 1 May 2019; accepted 3 May 2019)

Abstract

Suppose that virtue is intrinsically morally good, and that we have a *pro tanto* moral reason to act in ways which promote it. Further suppose that the failure of agents to receive what they deserve is intrinsically morally bad, and that we have a *pro tanto* moral reason *not* to act in ways which frustrate desert. When we are deciding whether to encourage others to make altruistic sacrifices, these two *pro tanto* moral reasons come into conflict. To encourage such sacrifices promotes virtue; it also causes virtuous agents to be worse off, preventing them from receiving their deserts. I argue that these effects on desert can reduce the moral desirability of promoting altruism so significantly as to make it morally wrong. This has implications for public policy, since certain practical questions turn on the extent to which we ought to rely on altruism as a means of solving social problems.

I. Introduction

Suppose that Nell, after carefully considering Peter Singer's arguments in *The Life You Can Save*,¹ resolves to give most of her wealth to Oxfam. As she begins to write her first cheque, it occurs to her that she must be an extraordinarily virtuous person. After all, very few people are willing to make such a large sacrifice in order to alleviate the suffering of others. Virtuous agents deserve to have their lives go better rather than worse, Nell thinks, and it is a shame that by donating she will instead be condemning herself to a life with no luxuries. Ever conscientious, Nell then begins to worry. It seems that it is morally bad when agents fail to receive what they deserve, and that we all have a moral reason not to bring about states of affairs in which this occurs. Yet by giving most of her wealth to Oxfam, Nell will simultaneously demonstrate that she deserves to have her life go *better* while ensuring that her life will actually go *worse*, thus preventing herself from receiving what she deserves. Do considerations of her own desert really provide Nell with a moral reason *not* to donate to Oxfam?

I think that they do. The reason is, of course, a *pro tanto* reason, and one which in this case is presumably overridden by other moral reasons which weigh in favour of donating. Important as Nell's desert may be, it cannot be *more* important than the saving of lives or the alleviation of intense suffering. So in this case, considerations of her own desert do not determine what Nell morally ought to do.

¹Peter Singer, *The Life You Can Save* (New York, 2009).

The above vignette nevertheless powerfully illustrates an important and largely unappreciated phenomenon.² Altruistic behaviour, irrespective of whatever morally desirable effects it may have on its beneficiaries, can generally be expected to have a morally *undesirable* effect on desert. In so far as altruism requires a sacrifice that makes the life of the altruistic agent worse, and in so far as altruistic agents are virtuous and therefore deserve to have their lives go well, a significant enough display of altruism will prevent the altruistic agent from receiving what she deserves.

In many cases, altruism's effects on its beneficiaries are so morally significant as to weigh decisively in its favour. But there are other cases in which this is not so. Suppose, for instance, that Nell has ultimately decided to pursue a career in public service, believing this to be a more effective means of reducing poverty than the donation of her personal wealth. She is now the official in charge of directing the poverty-relief efforts of the United Nations, and she must choose between the following (mutually exclusive, we may suppose) strategies:

Ad Campaign: The United Nations will launch an advertising campaign across the developed world which publicizes the severity of global poverty and the moral arguments in favour of alleviating it. The campaign is expected to move a number of agents in the developed world to start caring about global poverty, and thus to become more virtuous; these newly virtuous agents are then expected to donate a total of \$X to poverty-relief organizations.

Global Tax: The United Nations will convince governments across the developed world to impose an involuntary 'poverty tax' on their citizens. The tax will collect an amount of money from each citizen proportional to his wealth or income, generating a total of \$X for poverty relief.

Since both choices would raise a comparable amount of money for poverty relief, they would have similar effects on their intended beneficiaries. But there is at least one morally important difference: Ad Campaign prevents many virtuous agents from receiving what they deserve, while Global Tax does not. That effect on desert, I contend, makes it the case that Nell morally ought to choose Global Tax.

The choice described above is unrealistically simple. But this *kind* of choice – between policies that aim to achieve social goals through the promotion of altruism and those that aim to do so by other means – is surely realistic. It is the kind of choice we face when we evaluate real-world poverty-relief schemes, which might rely on altruism to a greater or lesser degree, albeit less starkly than in the case described above. It is also the kind of choice we face when deciding whether it should be legal to pay live organ and tissue donors, or whether compensating these donors should be forbidden. More generally, it is the kind of choice we face when deciding to what extent society should rely on rewards or coercion – rather than virtue – to induce its members to behave in desirable ways. For all such choices, I argue, altruism's effects on desert serve to reduce its moral desirability relative to the alternatives. And when all other

²To my knowledge, the only existing treatments of a related phenomenon are Shelly Kagan, 'Indeterminate Desert', *The Good, the Right, Life and Death: Essays in Honor of Fred Feldman*, ed. Kris McDaniel, Jason R. Raibley, Richard Feldman and Michael J. Zimmerman (New York, 2017), pp. 45–70, and Bradford Skow, 'A Solution to the Problem of Indeterminate Desert', *Mind* 121 (2012), pp. 37–65. I explain, in section II, the relationship of their claims to mine.

things are close enough to being equal, these effects render the promotion of altruism morally wrong.

This conclusion is likely to strike many of us as surprising. A major reason for this, I suspect, is that the promotion of altruistic sacrifice has an additional effect that may seem morally desirable: it also causes some agents to become more *virtuous*. The idea that the promotion of altruism is made more desirable by its effects on virtue has some history in the literature. Richard Titmuss, for instance, argues that those who provide blood for transfusions should not be paid for doing so; among other considerations, he cites the fact that non-compensated blood donation provides the opportunity for virtuous altruism in a way that the sale of blood does not.³

If virtue is in fact morally desirable, then I will need to show that the badness of altruism's effects on desert can sometimes outweigh the goodness of its effects on virtue. To this end, I spend the next two sections examining Nell's choice between Ad Campaign and Global Tax in greater detail. In the second section I build an intuitive case for the claim that Ad Campaign's effects on desert make it, on balance, morally worse than the alternative; in the third section I bolster this intuitive case by appeal to a theoretical framework offered by Thomas Hurka.⁴ Key to the intuitive case is the observation that Ad Campaign has two distinct undesirable effects on desert: it not only prevents individual agents from receiving what they deserve, but it also systematically disadvantages virtuous agents relative to non-virtuous ones. Hurka's framework, I believe, supports my contention that these two effects should be treated as distinct sources of negative value; the combined badness of these two effects should, I argue, be sufficient to outweigh the goodness of Ad Campaign's desirable effects on virtue. In the final section, I turn from this test case to the real world. Realistic cases are more complex, but we can identify some in which considerations of desert could make a difference to which policies we morally ought to pursue; I suggest the debate over human kidney markets as one promising example.

II. The intuitive case

I assume at the outset that some pluralistic form of maximizing act-consequentialism is correct, and that both virtue and desert are morally important. The presence of virtue adds value to a state of affairs, and the failure of agents to receive what they deserve adds disvalue. For clarity, I refer to that value contributed to a state of affairs by virtue as *virtue-goodness*, and that disvalue contributed by the failure of agents to receive what they deserve as *desert-badness*.⁵ My aim is to show that the desert-badness produced by the promotion of altruism can sometimes outweigh the virtue-goodness which it also produces.

We need not take a position, at least in the present section, on whether desert is non-comparative or comparative – that is, on whether an agent deserves a particular level of

³Richard Titmuss, *The Gift Relationship: From Human Blood to Social Policy* (New York, 1997).

⁴Thomas Hurka, *Virtue, Vice, and Value* (New York, 2001); Thomas Hurka, 'The Common Structure of Virtue and Desert', *Ethics* 112 (2001), pp. 6–31.

⁵In this I am adapting Hurka's ('Common Structure') practice of referring to the value contributed by an agent's receiving his desert as 'desert-good', and the disvalue contributed by his failing to receive it as 'desert-evil'. Throughout the article I refer to 'bad' and 'badness' where Hurka refers to 'evil'. This is for the sake of clarity; I wish to avoid the potentially misleading connotations of malice or ill will associated with 'evil'.

well-being which is fixed by the objective quality of his character, or whether he instead deserves a level of well-being that is *better than* that of less-virtuous agents and *worse than* that of more-virtuous agents.⁶ Nor need we decide between virtuous character and virtuous actions as the basis of desert. That is, we might hold that agents deserve certain things on the basis of their character alone; we might alternatively hold that agents deserve certain things on the basis of whether they *express* good or bad character by *acting* virtuously or viciously. Either view would allow the following discussion to move forward.

One restriction, however, is required. We cannot allow that what an agent deserves is based on the *deontic* status of his actions – that is, on whether he acts in ways that are morally *right* as opposed to morally *wrong*. Shelly Kagan points out that such an account of the basis of desert generates unacceptable conclusions when combined with the claim that desert is morally important.⁷ Kagan's basic observation is this: if desert is morally important, then what an agent deserves sometimes determines which actions towards that agent are morally right. And if the deontic status of an agent's actions in turn determines what that agent deserves, then there can be cases in which what an agent deserves is dependent, in an objectionably circular way, on *itself*. Suppose that I deserve to be treated badly. If I treat myself badly, I act rightly. In acting rightly, I may make it the case that I deserve to be treated well. If so, then in treating myself badly, I act wrongly. In acting wrongly, I make it the case that I deserve to be treated badly... And so on. There is no stable assignation of desert in this scenario, and Kagan shows that the phenomenon can be replicated even in cases which are considerably more complex and in which considerations of an agent's *own* desert play no role.⁸

I mention Kagan both to rule out one possible account of the basis of desert and to distinguish the phenomenon which he describes from the one which is of interest in this article. If we assume that desert is based either on virtuous character or on the expression thereof, then we are not vulnerable to the kinds of problems that Kagan envisages. That is because the *deontic status* of an action is distinct from the quality of the character that it expresses (and the character of the agent who performs it). If desert is based either on character or the expression thereof, then what an agent deserves will not change even if the deontic status of her actions changes. If Nell is moved to give to Oxfam because she cares deeply about the well-being of others, then she is virtuous and expresses her virtue by acting. Her character and her actions would *remain* virtuous even if the act of donating were made wrong by its side-effects on Nell's desert; thus there is no potential for the objectionable indeterminacy and instability which trouble Kagan.

⁶I do wish to assume that each agent deserves some *level* of well-being, as opposed to deserving specific rewards or punishments as the result of particular actions. An anonymous reviewer suggests a case in which an agent finds a bag of money on the street; not having earned the money, there is some sense in which she does not deserve it. But this is not the sense of 'desert' that is of interest here. On the kind of view which I am assuming, the discovery of a bag of unearned money may or may not have a desirable effect on desert – this depends on whether the money brings the agent closer to, or further away from, the level of well-being she deserves.

⁷Kagan, 'Indeterminate Desert'.

⁸For simplicity I have chosen a case in which there is no stable assignation of desert, although Kagan is more troubled by cases in which there are multiple stable assignations and in which an agent's desert is therefore indeterminate. Not everyone is convinced by Kagan's argument; see Skow's 'Solution' for a dissenting view.

So there is not, on my view, anything paradoxical or problematic about desert's implications for the promotion of altruism, nor are these implications meant to be a *reductio* of anything; I hold that altruism's effects on desert really *do* make it less morally desirable. My aim in this section and the next is to show that these effects can be so significant as to make a difference to what we morally ought to do, using Nell's choice between Ad Campaign and Global Tax as a test case.

It is no cause for concern that this case is unrealistically simple, since our present goal is merely to demonstrate that it is *possible* for the desert-badness caused by the promotion of altruism to outweigh its virtue-goodness. To this end, we are entitled to make further simplifying assumptions as required; the idea is to finagle the case so that the *only* morally significant differences between the two plans relate to their effects on virtue and desert in the developed world. We may assume, for instance, that the involuntary nature of the poverty-relief tax does not affect its desirability – we may set aside, for the time being, any worries about the moral significance of *coercion*.

We may further assume that the total reduction in well-being in the developed world is identical on both plans. This represents a simplification because, in reality, it is possible to remove the same amount of *money* from the developed world without causing an identical reduction in *well-being*. The marginal utility of additional wealth diminishes; thus a relatively poor agent who parts with a certain sum of money would suffer a greater reduction in well-being than a wealthier agent who parts with the same sum. But we need not account for this complication in considering our test case. We may assume that the total reduction in well-being caused by the two plans is the same; the only difference is that on Global Tax, the reduction is distributed evenly across agents in the developed world, whereas on Ad Campaign, it is concentrated upon the virtuous.

Finally, we may assume that the two plans have *completely* identical effects on the beneficiaries of the poverty-relief. We have already assumed that they raise a comparable amount of money, but we may additionally assume that it makes no difference to the beneficiaries whether this money comes from voluntary charity or an involuntary tax. As noted, this assumption is important because it prevents the effects on the beneficiaries' well-being – which might otherwise be the dominant factor in determining which plan is morally best – from overwhelming the moral significance of virtue and desert. It also prevents us from having to consider the effects of poverty relief on the virtue of its beneficiaries, or on the extent to which they receive their deserts. Its effects on the latter are presumably desirable – many, if not all, of those suffering from severe poverty deserve to have lives that are better than they currently are, and thus the alleviation of their poverty brings these agents closer to the level of well-being that they deserve. Perhaps – if a certain level of material well-being facilitates the cultivation of virtue – its effects on the former are also desirable. Taken together, these effects may be of great intrinsic moral importance. But they do not need to factor into our evaluation of Ad Campaign and Global Tax because they are the *same* for both plans and therefore cannot ground any moral difference between the two.

From this point forward, therefore, our discussion of these two plans will have a restricted scope: when I write that a plan will result in virtue-goodness or desert-badness, I mean that it will produce such goodness or badness *in the developed world*.⁹ Global Tax is meant to affect neither virtue nor desert in the developed

⁹Alternatively: it is acceptable to read such claims as 'Ad Campaign produces virtue-goodness' as 'Ad Campaign produces virtue-goodness *relative to* Global Tax', or 'Ad Campaign is virtue-better than Global Tax', meaning that its overall effects on virtue are more desirable than those of the alternative.

world.¹⁰ Ad Campaign is meant to affect both; its effects on virtue are desirable, and its effects on desert undesirable. That it has a morally desirable effect on virtue should be clear. In causing some agents to care about global poverty, we have supposed, Ad Campaign makes those agents more virtuous.

That it has a morally undesirable effect on desert may be less obvious. One might point out, first, that all agents whom Ad Campaign induces to make sacrifices do so voluntarily; the plan cannot negatively affect any agent's level of well-being without that agent's consent. Perhaps this entails that Ad Campaign's effects cannot be morally undesirable, given that the affected agents have all chosen to waive their entitlements to what they deserve. But this suggestion is in conflict with the consequentialist assumptions made at the outset of the present section. If desert is morally important *from the perspective of the universe*, as I take a properly consequentialist view to entail, then the agent whose desert it is cannot nullify its moral significance at will. This is most easily seen if we consider cases of agents who deserve to have their lives go *poorly*. A murderer has a moral reason to submit to the authorities so as to receive the punishment he deserves, and it is not as though he can neutralize this reason by waiving his entitlement to his desert. The fact that a given set of agents have consented to not receiving what they deserve does not abolish the badness of their not receiving it, nor preclude us from having to factor that badness into our moral calculus.

A different worry: one might point out that, on the view which we are assuming, virtue is *good*. Does it not follow that the agents affected by Ad Campaign are made *better off* rather than *worse off* when they become more virtuous? It does not, in fact, follow. To reiterate, we have assumed that virtue is good from the perspective of the universe – that its presence makes a state of affairs more valuable than it would otherwise have been. This does *not* imply that virtue is good *for the virtuous agent*.¹¹ Even if virtue *is* good for the virtuous agent, it does not follow that virtue is *so* good as to outweigh the badness of any potential reduction in material well-being. Some reflection, I think, will reveal that most of us do not think that virtue is *that* good: most of us will find it quite appropriate to describe certain agents as both extraordinarily virtuous and extraordinarily unfortunate in virtue of their material circumstances.¹²

I assume here that those agents who donate enough of their wealth to poverty-relief efforts fall into that category: their virtue, even if it is good for them, is not sufficient compensation for their material sacrifice, and so they *do* make themselves worse off by donating. Do they make themselves worse off than they *deserve*? This seems almost certain to be the case, so long as the sacrifice they make is reasonably significant. On a

¹⁰Since Global Tax makes all agents in the developed world somewhat worse off, it may make some of them worse off than they deserve and thus have a minor undesirable effect on desert. Because this effect is not concentrated on the virtuous, it is dwarfed by Ad Campaign's undesirable effects on desert, and can be ignored for the purposes of comparing the two plans.

¹¹Oddly, Hurka – whose view will be the subject of the next section – seems to conflate the claim that virtue is good *simpliciter* with the claim that it is good *for* the virtuous agent (see e.g. Hurka, *Virtue*, pp. 55, 153). It seems to me that he *ought* to hold that virtue is good in the former rather than the latter sense, especially in light of the structural analogies he suggests between virtue and desert. It is good for a vicious agent to be punished and therefore to receive his desert, but surely it is not good *for* the punished agent – rather, it is good *simpliciter*.

¹²An anonymous reviewer suggests that agents are likely to derive some psychological satisfaction from acting altruistically, and offset the reduction in their well-being thereby. But I take it that this satisfaction is limited in magnitude and that it cannot be expected to outweigh or even significantly mitigate the reduction in well-being incurred by a sufficiently large material sacrifice.

non-comparative account, each agent deserves some particular level of well-being which corresponds to his level of virtue or virtuous action, and we may safely assume that many of the most virtuous agents deserve lives that are as good as, if not better than, the ones they currently enjoy. (To deny this would require us to hold that even the most virtuous agents in the developed world currently enjoy lives that are much *better* than what they deserve; this seems implausible.) On a comparative account, an agent does not deserve a particular level of well-being that corresponds to the objective quality of his character, but instead deserves a life that is worse than that of more-virtuous agents and better than that of less-virtuous agents. Since Ad Campaign induces virtuous agents to give away their wealth without affecting non-virtuous agents, it reduces the well-being of the virtuous relative to that of the non-virtuous; thus, it also prevents them from receiving what they deserve on a comparative account.

The question is whether Ad Campaign's undesirable effects on desert outweigh its desirable effects on virtue. One might think that the answer will be different depending on how we fill in the details of the case – I have not specified precisely how many agents are moved to become more virtuous, nor how much, exactly, those agents are sacrificing when they donate. Perhaps different sets of details would result in different amounts of virtue-goodness and desert-badness, in turn determining which effect is more significant. If indeed this is our initial reaction to the case, it tells us something important – it suggests that we are already intuitively committed to the claim that neither virtue-goodness nor desert-badness is lexically superior to the other. If we think that either virtue-goodness or desert-badness could outweigh the other, depending on the exact quantity produced, then we must think that the moral significance of these two effects is on roughly the same order of magnitude.

For my part, I think that small adjustments to the number of agents affected by Ad Campaign or to the exact magnitude of their sacrifices are unlikely to make a difference – so long as the number of agents remains reasonably large and the sacrifices these agents make remain reasonably significant. This is not because either virtue-goodness or desert-badness is lexically superior, but because, in cases like Ad Campaign, the two scale in proportion to one another. As we increase the number of agents who are made virtuous, we also increase the number of agents who are caused not to receive their deserts. And as we increase the sacrifice made by each agent, we increase not only the extent to which he fails to receive his desert but also the degree of virtue he displays. Thus, in evaluating Ad Campaign, we are not measuring two independently generated quantities of virtue-goodness and desert-badness against one another. We are asking whether the virtue-goodness of making a certain group of agents more virtuous can be outweighed by the desert-badness of preventing that *same* group of agents from receiving their deserts.

It is not clear how much precision our intuitions can provide in answering a question like this. We can ask: how *good* is it when many agents are made significantly more virtuous? I take it that a natural response would be to say that this is *very good*. We can also ask: how *bad* is it to prevent that same group of virtuous agents from receiving what they deserve? I take it that a natural response here would be to say that this is *very bad*. But it seems unlikely that our intuitions can tell us which of these effects is stronger than the other; intuitive judgements simply do not have a fine enough grain. They may be able to tell us that an effect is *very* (as opposed to *somewhat* or *slightly*) important. But they do not generally seem to be able to tell us which of two very important effects is *more* important.¹³

¹³I take it that the saving of lives and the alleviation of intense suffering are likely to strike us as *extremely* important, which explains why they can outweigh any effects on virtue and desert.

Even so, I think that there is a strong intuitive case for the claim that desert-badness outweighs virtue-goodness in Ad Campaign. That is because the promotion of altruism in this case has an additional undesirable effect on desert. We have already considered the fact that Ad Campaign will cause many individual agents to fail to receive what they deserve. But there is something intuitively significant about the fact that this effect is systematic and iterated across the entire developed world. *Each* agent who is made more virtuous is also caused not to receive what he deserves, while those agents who are not virtuous (and who remain non-virtuous) are not affected. The end result is a society in which the virtuous are, as a whole, worse off than the less-virtuous, and I suspect that many of us will be inclined to describe this effect – call it the *social* effect – as a *very bad* one. Since, I contend, it is a *distinct* effect from Ad Campaign's effect on individual desert, it means that Ad Campaign has a total of *one* very good effect and *two* very bad effects. That gives us a reason to think that the bad effects outweigh the good.

It is only a *prima facie* reason, of course. More will need to be said about *how* good and bad these effects are; two very bad effects can still be outweighed by a good effect if it is good *enough*. But given the limited precision of our intuitive judgments, it is about as strong a reason as an appeal to intuitions can provide. We can, I think, say more about the relative strengths of these moral considerations, but that will require us to move into the realm of theory. That is the task of the next section.

First, though: why think that Ad Campaign's social effects represent an additional source of badness? I have remained neutral here with respect to comparative and non-comparative accounts of individual desert. On either kind of account, however, we might find reason to object. First, suppose that we are non-comparativists. What an agent deserves, we think, is some particular level of well-being that is determined by the objective quality of his character or his actions. Comparisons with other agents are totally irrelevant to what an agent deserves, and it should by extension be irrelevant that less-virtuous agents do not have their well-being affected by Ad Campaign. The worry for the non-comparativist, therefore, is that Ad Campaign's social effects have nothing to do with desert.

Suppose instead that we are comparativists. What an agent deserves, we think, is a level of well-being that is proportional to his relative level of virtue: he deserves to be better off than those agents who are less virtuous than he is, and worse off than those who are more virtuous. Our worry, as comparativists, is not that the social effects of Ad Campaign are irrelevant to desert; rather, it is that they simply represent the aggregation of its individual effects on desert. To make a virtuous agent worse off than his less-virtuous peers *is* to prevent him, individually, from receiving his desert. And to make *all* virtuous agents worse off than their less-virtuous peers *is* to cause all of them, individually, to fail to receive their deserts. How, then, could the society-wide effects on the distribution of goods represent an additional source of desert-badness – one over and above the summed badness of the effects on each individual agent?

As we shall see, Hurka's framework gives us reason to think that Ad Campaign's social effects *do* represent an additional source of *desert*-badness. But we need not accept that claim for the sake of the present intuitive case. All we need to accept is that these social effects represent *some* kind of additional badness which is significant enough to make a moral difference; nothing turns here on the linguistic question of whether this badness is properly described as relating to 'desert'. We might instead say that a society in which the virtuous are systematically made worse off than the non-

virtuous is 'unjust', or use any other term we like.¹⁴ The end result is that Ad Campaign still has two very bad effects – one on desert and one on justice (or whatever we prefer to call it) – as against a single very good effect on virtue. This, in short, is the response to the worries of both the comparativist and the non-comparativist.

That there *is* something additionally undesirable about Ad Campaign's social effects may require some further argumentation, but in the end I think that this claim is also likely to be supported by common intuitions. Many of us, I take it, have the intuitive sense that a society can be good or bad *qua* society, such that a society's goodness or badness is not wholly exhausted by the goodness or badness of what happens to its members. And the fact that virtue is *punished* in a given society should seem to make it worse, *qua* society. Ad Campaign results in a society in which virtue *is*, in effect, punished¹⁵ – the absence of other mechanisms to alleviate global poverty compels the virtuous to make massive sacrifices, while the non-virtuous are unaffected. In so far as this does make the society in Ad Campaign worse – and in so far as it matters morally whether our societies are good or bad – this represents a moral reason to choose Global Tax instead.

III. Hurka's framework

I have offered an intuitive case for the claim that Ad Campaign's combined individual and social effects on desert (or, its effects on desert plus its social effects, if we prefer to think of the latter as unrelated to desert) make it all-things-considered worse than Global Tax. So far as intuitive cases go, I think it is a strong one. But it would be better if we could appeal to a theoretical framework, one which would allow us to assess with greater clarity the relationship between Ad Campaign's social effects and its effects on the desert of individuals. The most suitable framework, in my view, is that offered by Thomas Hurka.¹⁶

Hurka's framework requires us to assume, first, that there are features of a state of affairs which are unrelated either to virtue or to desert but which can, by virtue of their presence or absence, contribute moral value and disvalue. These are the *base-level goods* and *base-level bads*. Precisely which features these are makes no difference to the structural claims which are of interest here; for illustrative purposes, I follow Hurka in supposing that pleasure contributes value to a state of affairs, while pain contributes disvalue.¹⁷

Once we have assumed that certain features of a state of affairs are good and bad in this way, Hurka suggests, it seems reasonable to think that certain attitudes *towards* these features are also good and bad. In loving the pleasure of others, for instance,

¹⁴Perhaps this seems like cheating, given my earlier admonition that Ad Campaign and Global Tax should be simplified in such a way that they differ *only* in their respective effects on virtue and desert. But I do not think that it is cheating. Unlike other potential differences, the social effects of Ad Campaign are an unavoidable consequence of mass altruism and cannot be simplified away.

¹⁵If this characterization seems doubtful, it is helpful to consider an analogous case. Suppose that I leave my desperately sick dog on your doorstep, knowing that you, a virtuous person, will feel obligated to take it to the veterinarian and incur the significant costs of treatment. Although I have not literally *compelled* you to do anything, it still seems that I have *wronged* you, and that I have, at least in a loose sense, punished you – by targeting you – for your virtue. My suggestion here is that a society which does not systematically work to alleviate poverty does essentially the same thing, though perhaps unintentionally, on a massive scale – it leaves the problem of poverty to be taken care of, and the attendant costs to be incurred by, the virtuous.

¹⁶Hurka, *Virtue*; Hurka, 'Common Structure'.

¹⁷Hurka, *Virtue*, pp. 11–12.

an agent loves that which is good: is not a love of the good itself good, *qua* attitude?¹⁸ Hurka holds that it is. In fact, he identifies virtue with *appropriate* attitudes towards goods and bads – positive attitudes towards that which is good are virtuous, as are negative attitudes towards that which is bad. And he identifies vice (partly, as we shall see) with *inappropriate* attitudes towards goods and bads – negative attitudes towards that which is good are vicious, as are positive attitudes towards that which is bad. Virtues and vices, then, are *higher-level goods and bads*. They are themselves good and bad, but they derive their goodness and badness from their appropriateness or inappropriateness towards good or bad objects.¹⁹

To many of us it also seems that different *outcomes* are appropriate for different *agents*, depending on whether these agents are themselves good or bad – that is, that agents *deserve* to have their lives go well or poorly based on their character. Hurka agrees, holding that it is good when a virtuous agent prospers or a vicious agent withers, but bad when a virtuous agent withers or a vicious agent prospers. And therein lies the crucial structural similarity which Hurka proposes between virtue and desert: both derive their moral value or disvalue from their *appropriateness* or *inappropriateness*. In the case of virtue, value derives from a match between a good or a bad object and a corresponding positive or negative attitude; in the case of desert, disvalue derives from a mismatch between a good agent and a bad life, or between a bad agent and a good life.²⁰

Hurka's framework seems especially apt for our purposes, because it harmonizes three claims to which we are committed in this article: that some form of maximizing act-consequentialism is correct, that virtue is good, and that it is bad when agents fail to receive their deserts. Hurka notes that it can be difficult to explain how virtue could be good, on a consequentialist normative theory, without resorting to an account of virtue which is fundamentally deflationary – one which defines virtues simply as those traits with desirable consequences, for instance.²¹ But Hurka's account offers an explanation: virtue derives its goodness from the correspondence of the constitutive attitude to a good or bad object. Since the goodness and badness associated with desert can be similarly explained in terms of correspondence, we have a framework with significant explanatory power.

In exchange for this explanatory power, however, Hurka's framework imposes some constraints on which accounts of desert are viable. Specifically, the structural analogy with virtue implies that desert must be assessed at least partly on a *non-comparative* basis. This is because the status of a particular attitude as virtuous or vicious must itself be assessed non-comparatively. An attitude is virtuous only if its valence and intensity are appropriate to its object. In order to be virtuous, for instance, my attitude towards your intense suffering must be negatively valenced – it must be an attitude like *sorrow* – and it must itself be intense. Comparisons with other attitudes seem to be irrelevant here. Mild annoyance at your intense suffering would not be virtuous even if it were more intense than my milder annoyance at a less-bad object, such as the stubbing of my own toe. If desert is analogous, then it seems that comparisons are also irrelevant when evaluating whether a particular agent receives what she deserves. On the basis of the quality of her character, an agent deserves a life with a corresponding level of

¹⁸Hurka, *Virtue*, p. 30.

¹⁹Hurka, *Virtue*, pp. 11–21; Hurka, 'Common Structure', pp. 8–10.

²⁰Hurka, 'Common Structure', pp. 10–12.

²¹Hurka, *Virtue*, pp. 3–4.

quality. Thus a virtuous agent deserves to be well off to some particular degree; it is not sufficient, for her to receive what she deserves, that she be *less badly off* than less-virtuous agents.

But crucially, for our purposes, Hurka's framework *also* accommodates a kind of comparative desert-badness. This is another feature which follows from the structural analogy with virtue. While individual attitudes can be intrinsically bad (and therefore vicious) only if they have the wrong valence or the wrong strength, Hurka holds that a set of attitudes taken as a *whole* can be intrinsically bad (and therefore vicious) if the strengths of the (individually good) attitudes therein are incorrectly *proportioned*. Indeed, Hurka thinks, we require our account of virtue to have this feature if we are to explain the viciousness of character defects like selfishness. Since one's own happiness is good, it is virtuous to love one's own happiness, yet some agents seem vicious because they love their own happiness *too much*. The solution is to recognize that the attitudes of such agents are incorrectly proportioned relative to one another; an agent who cares about his own happiness ten times more strongly than he cares about the happiness of others has a set of attitudes which are vicious and therefore bad *when taken as a whole*, even if each individual attitude is virtuous.²² Such failures of proportionality represent an *additional* source of badness when some of the individual attitudes are themselves bad: an agent who *hates* the happiness of others is vicious not only in so far as this attitude is itself vicious, but *also* in so far as it negatively affects the quality of his attitudes taken as a whole.²³

Desert, on Hurka's view, works similarly. There can be states of affairs in which each agent receives a significant portion of what he deserves – and in which each agent therefore seems to contribute some desert-goodness, individually – but in which the overall situation seems to contribute desert-badness because of the way in which the goods are distributed. Two equally virtuous agents might prosper to different degrees, for instance. If the difference in prosperity is dramatic enough, Hurka holds, it can make the distribution of well-being across these two agents bad as a whole.²⁴ The same phenomenon can exacerbate the badness of a scenario in which some agents do *not* receive what they deserve. If a virtuous agent withers instead of prospering, that is bad in itself; if a virtuous agent withers while less-virtuous agents prosper, then that distribution of desert is *additionally* bad as a whole.

So Hurka's framework allows for both comparative and non-comparative desert badness. Whether or not an individual receives her desert is evaluated non-comparatively, and the failure of particular agents to enjoy particular levels of well-being contributes one kind of desert-badness – call this *individual desert-badness*. But the overall distribution of well-being is evaluated comparatively, and, if it is a bad distribution, it contributes additional desert-badness – call this *proportional desert-badness*.²⁵ If Hurka is right, then Ad Campaign does produce two distinct bad effects on desert – it generates individual desert-badness by virtue of preventing individuals from receiving their

²²Hurka, *Virtue*, pp. 83–7; Hurka, 'Common Structure', pp. 24–5.

²³Hurka, *Virtue*, pp. 88–9.

²⁴Hurka, 'Common Structure', pp. 26–7.

²⁵My usage of 'individual' and 'proportional' follows Hurka's usage (*Virtue*, pp. 193–7) of the same to describe two facets of desert, although he does not use them to distinguish different species of desert-badness.

deserts, and proportional desert-badness by virtue of its effects on the distribution of goods across the population of the developed world.²⁶

A crucial question remains: *how* bad are Ad Campaign's combined individual and proportional effects on desert, relative to its desirable effect on virtue? Hurka proposes the following principle: the moral goodness of a virtuous attitude cannot exceed the goodness or badness of its object.²⁷ Thus, though my sorrow that you are in pain is good, its goodness cannot outweigh the badness of your pain itself. Hurka offers several arguments for this principle; most are not relevant for our purposes because they seem unlikely to apply analogously to desert.²⁸ One, however, is worth stopping to consider: virtue's status as a 'dependent' good means that its moral significance cannot exceed that of the original good upon which it depends. The goodness of my sorrow at your pain, for instance, depends upon the badness of your pain. It is the badness of your pain – in conjunction with the fact that sorrow is an appropriate attitude towards a bad object – which makes my sorrow good. Thus – to put the point more evocatively – the goodness of my sorrow is in some sense merely a *reflection* of the badness of your pain. It would be odd, Hurka suggests, if this reflected moral significance were able to outshine its original source.²⁹

Suppose that this reasoning is sound. If so, it could imply that there are analogous limits on the moral weight of individual desert-badness relative to that of virtue-goodness. Consider a virtuous agent who withers. The badness of his withering is a reflection of the goodness of his virtue; it is only because his virtue is good that his withering represents a mismatch between the quality of his character and the quality of his life. And it might seem similarly odd for the badness of his withering to outshine the goodness of the virtue it reflects. Perhaps, then, it is impossible for the individual desert-badness of a given agent's failure to receive his desert to outweigh the goodness of that agent's virtue.³⁰

²⁶Might Ad Campaign also produce two *desirable* effects on virtue? In causing an agent to care about global poverty, we cause him to acquire a new attitude which is individually virtuous. It seems likely that we also cause his attitudes to become better proportioned relative to one another – this would be the case if he had previously cared about global poverty less than he should have relative to other objects. However, this proportional effect on virtue does not represent an additional source of goodness which needs to be factored independently into our evaluation of Ad Campaign. In so far as this proportional effect increases the virtue of the agents in Ad Campaign, it makes them deserve even better lives and increases the extent to which they fail to receive their deserts; thus, the effect is offset by a corresponding increase in desert-badness.

²⁷Hurka, *Virtue*, p. 129.

²⁸For instance, Hurka thinks that this principle is required to prevent base-level goods and bads from being overwhelmed by the goodness of higher-order attitudes *towards* virtue. Since virtue is good, a correct attitude towards a virtuous attitude is also good – so it is not only good that I feel happiness at your pleasure, but also good that I feel happiness at my happiness at your pleasure, that I feel happiness at my happiness at my happiness at your pleasure, and so on. To prevent these 'infinite hierarchies of virtuous attitudes' (*Virtue*, p. 136) from outweighing the base-level goods at the bottom, Hurka thinks that the goodness of each level must be less than that of its object at the level below it. This argument does not seem to apply to desert, however, because desert admits of no analogues to these infinite chains of higher-level attitudes.

²⁹Hurka, *Virtue*, p. 141.

³⁰For somewhat technical reasons, Hurka himself does *not* think that an analogous principle limits the weight of desert-badness ('Common Structure', pp. 17–20). There is insufficient space to discuss these reasons here, and they need not concern us. In this section I have assumed that we accept the broad contours

Such a principle would threaten my overall argument only if it similarly limited the moral weight of *proportional* desert-badness. And there is good reason to think that it does not. Upon what, we should ask, does the moral importance of proportional desert-badness depend? This kind of badness does not result from a simple mismatch between the quality of an agent's character and the quality of his life; it results from a state of affairs in which a virtuous agent does *less* well than others even though he is *more* virtuous than they are. It is bad, in part, because the virtuous agent's virtue is good, but that is not the whole story; to explain its badness, we must also appeal to the fact that the character of less-virtuous agents is *bad*, or at any rate *worse*. Thus proportional desert-badness reflects not only the goodness of one agent's virtue but also the *badness* of every less-virtuous (and more prosperous) agent's *lack* of virtue.

Such limits as there may be on the weight of Ad Campaign's proportional desert-badness therefore do not seem to be very restrictive. This badness cannot be greater in magnitude than the combined badness of all the vices present in the developed world; but that level of badness, though hard to quantify, is presumably enormous. At any rate it is very likely to be greater than the virtue-goodness caused by Ad Campaign. We may assume that the group of agents who are made more virtuous by the advertisements represent only a fraction of the developed world's population; were we to add together the sloth, selfishness, and other vices of the majority, their badness would, I take it, vastly outweigh the goodness of this smaller group's virtue.

To be clear: the preceding discussion is merely meant to establish the upper bounds on Ad Campaign's proportional desert-badness. But in showing that these upper bounds are high, I have also given us some reason to think that Ad Campaign's *actual* proportional desert-badness is significant enough to outweigh its virtue-goodness. That this desert-badness reflects a vast quantity of good and bad suggests not only that its magnitude *can* be great, but also that it *is*. Possibly this proportional desert-badness is enough, on its own, to make Ad Campaign morally worse than Global Tax; when combined with Ad Campaign's individual desert-badness, it seems even more likely to do so.

IV. Altruism in the real world

So: Hurka's framework does not imply that the desert-badness of Ad Campaign must outweigh its virtue-goodness. But it is compatible with – and, I think, strongly suggestive of – this possibility. If we take Hurka's framework seriously, we have some theoretical justification for our intuitive reactions; if not, the intuitions will have to stand on their own. In either case I take myself to have shown that we have good reason to think that the promotion of altruism *can*, in principle, be made morally wrong by its effects on desert. I now turn to the question of whether this occurs in any real-world cases.

One problem is that it can be difficult to identify realistic cases in which the same goal can be accomplished equally well by either altruistic or non-altruistic means. Titmuss's discussion of blood probably does not provide a good example. His appeal to the value of altruism is one among many arguments; among his others is the claim that donated blood is less likely to be contaminated than commercially sourced blood. If this is correct, then the better health outcomes associated with donated blood most likely outweigh any desert-badness. Another problem is that blood donors generally do not make themselves *significantly* worse off when they donate. On the

of Hurka's framework, but we are not compelled to agree with all of its details. And I describe here a reason to worry that we should accept such a limiting principle, even if Hurka does not.

assumption that blood donors experience at worst some inconvenience, it seems unlikely that donating can generate significant desert-badness.

The live donation of human kidneys would seem to provide a more promising example. As it stands, agents are generally not (legally) permitted to buy or sell kidneys as they would other commodities, but the shortage of donor kidneys has led to discussions of whether this policy ought to be changed. Typically, the two options seriously considered are the status quo – on which kidney sales are banned and on which there is a chronic shortage of kidneys – and the legalization of kidney markets. Suppose that we consider the status quo unacceptable, and that *some* means of addressing the shortage must be found. One alternative way of doing this might be by encouraging the altruistic, non-compensated donation of kidneys. This currently happens on far too small a scale to meet the demand for kidneys, but social mores evolve over time, and perhaps, with the right sort of encouragement on the part of public health authorities, more and more people might come to regard the donation of a kidney to a stranger as a live option – a charitable act more radical than, but fundamentally similar to, the donation of blood. Suppose, then, that we have the choice of pursuing one of two possible futures: one in which the kidney shortage is addressed by legal markets in these organs, and one in which it is addressed by an increase in non-compensated live organ donations. In both cases the number and quality of kidneys available for transplantation will be the same; the difference lies in how they are obtained.

I make four (realistic, I think) assumptions here. First, I assume that an agent who donates a kidney without being compensated makes herself meaningfully worse off by doing so. *How much* worse off she makes herself is currently the subject of empirical investigation,³¹ but it seems reasonable to assume that voluntarily undergoing any non-trivial surgical procedure represents a significant sacrifice, irrespective of any health effects that it may or may not have. Second, I assume that only the most virtuous agents will donate a kidney to a stranger without compensation. Though on one of these possible futures such donations have become relatively routine, we should imagine that they still represent a much more significant exercise of virtue than the donation of blood. Third, I assume that there is nothing intrinsically morally objectionable about the purchase of kidneys, such as would render it wrong automatically or regardless of the circumstances. If it turns out that such purchases are, for instance, exploitative in a way that makes them wrong *no matter what*, then the following discussion will not be very interesting; the possible future on which we obtain more kidneys through the promotion of altruism will win by default. But it has been plausibly argued that kidney markets are not intrinsically exploitative in this way, and that the grotesque abuses which sometimes attend kidney sales in the real world result from the fact that these sales must currently occur on the black market.³² My fourth assumption is that such abuses would not attend kidney sales if they were legalized.

Given these assumptions, there is a case to be made that the desert-badness caused by non-compensated kidney donations would outweigh the associated virtue-goodness, and that this effect would make a market regime morally preferable. If, as I have assumed, only the most virtuous agents will be moved to donate their kidneys without compensation, then such donations will disproportionately harm the most virtuous.

³¹See e.g. Bertram L. Kasiske, 'Outcomes after Living Kidney Donation: What We Still Need to Know and Why', *American Journal of Kidney Disease* 64 (2014), pp. 335–7.

³²See e.g. Janet Radcliffe Richards, *Careless Thought Costs Lives: The Ethics of Transplants* (New York, 2012).

That can be expected to have an undesirable effect not only on the individual desert of these particular virtuous agents, but also on the distribution of well-being as a whole, as less-virtuous agents will refrain from making the significant sacrifice entailed by non-compensated donation.

There is a complication, however. Whereas in Nell's case the alternative to altruism was to distribute the burden of poverty relief across all agents in the developed world, the burden of kidney donation cannot be comparably distributed. There is no way to donate a fraction of a kidney; an agent either parts with a kidney or she does not. And since the alternative to altruism involves *paying* agents for their organs, we can expect the burden to fall disproportionately on those most in need of money – that is, upon the poorest agents.

That the poorest agents would bear this burden might cause us to wonder whether a market regime would also result in significant desert-badness. The worry incorporates implicitly the assumption that kidney donation represents a genuine burden *even if* the donors are paid – that is, that whatever compensation is offered to the donors does not prevent them from being worse off, on balance, as the result of selling their kidneys. We have already assumed that the abuses which attend real-world kidney sales are not a factor in the possible future in which such sales are legalized; this entails that the removal of kidneys will occur under hygienic conditions, that sellers will receive the compensation which they were promised, and so on. Even given these assumptions, it is possible that the sale of a kidney entails a net reduction in well-being; suppose, for the sake of argument, that it does.³³ It is likely that many of the poorest agents already fall below the level of well-being that they deserve; on the assumption that selling their kidneys would entail an additional reduction in well-being, it would exacerbate this problem and therefore generate additional desert-badness.

But it is far from clear that this generates *more* desert-badness than the non-market alternative. For on the non-market system, it is the most virtuous agents who bear the burden of providing kidneys. And while it is reasonable to assume that the poor generally do not deserve to have their lives go worse, the most virtuous agents are *guaranteed* not to deserve to have their lives go worse. It therefore seems to me that a regime which systematically harms the most virtuous agents is likely to have *worse* effects on desert than a regime which systematically harms the poorest ones.

So I tentatively conclude that it is morally better for us to work towards a future in which kidney markets are legal, rather than one in which the shortage in organs is addressed through a change in social attitudes towards donation. Irrespective of whether I am right about this, I hope at least to have illustrated that altruism's effects on desert *can* make a decisive moral difference in realistic cases under empirically plausible assumptions. Whether these effects make such a difference in any particular case depends in part on the severity of the moral costs associated with the alternatives. If we are to achieve a particular result through non-altruistic means, we must have some

³³ As Richards (*Careless*, pp. 50–1) points out, it is reasonable to presume that agents who sell their kidneys make themselves *better off* by doing so; they have, after all, chosen to act in the way that they judge to be in their best interests. This presumption is rebuttable; the agents in question could be mistaken about where their best interests lie. Some research suggests that kidney sales are generally harmful to the sellers, though this seems to reflect conditions under the current black market regime rather than one in which kidney sales are legal; see e.g. Madhav Goyal, Ravindra L. Mehta, Lawrence J. Schneiderman and Ashwini R. Sehgal, 'Economic and Health Consequences of Selling a Kidney in India', *Journal of the American Medical Association* 288 (2002), pp. 1589–93.

other method of inducing agents to act in the desired way. One method, as we have seen, is to compensate agents for the desired behaviour, and this can have the morally undesirable effect of disproportionately harming those agents who are worst off.

Another method, broadly speaking, is the application of coercion. Coercion, in the form of taxation, was at work in Global Tax, and it is also likely to strike many of us as carrying a moral cost – either because it is intrinsically undesirable to coerce others, or because coercion can be expected to have undesirable side-effects, such as on happiness.

There may be other methods. An intriguing possibility is suggested by the recent literature on ‘nudging’ – a general term for a collection of non-coercive techniques that exploit various cognitive and behavioural biases to elicit desirable behaviour.³⁴ One kind of nudging, for instance, exploits the human bias towards the status quo. Perhaps a global fund for poverty relief could be supported by an ‘opt-out’ tax which agents can decline to pay by taking positive steps to be exempted. Under certain conditions, a significant portion of agents might decline to opt out and would therefore pay the tax; their choosing to pay would not prevent them from receiving their deserts, since it would be an expression of bias towards the status quo rather than of virtue. Nudging is likely to be limited in what it can accomplish, however. If the required sacrifice were large enough, the most selfish agents would be likely to overcome their status quo bias and decide to opt-out; the burden would then be borne disproportionately by virtuous agents, and the undesirable effects on desert would be similar to those in Ad Campaign.

Each of these alternatives, then, has something to be said against it. In the case of compensation for desirable behaviour, there is the possibility that the poor will be disproportionately harmed. In the case of coercion, there is the possibility of additional moral badness resulting from the coercion itself. And in the case of nudging, the ability to induce desirable behaviour is likely to be limited by the strength of the human bias towards the status quo. The point which I have tried to make is that each of these alternatives *also* has something to be said in its favour – namely, that it does not promote altruism and does not have the attendant morally undesirable effects on desert. I think that in some cases this will make the relevant alternative morally best. In *all* cases, however, the effect which I have described here should serve to *increase* the moral desirability of these alternatives relative to the promotion of altruism.³⁵

³⁴For the seminal article, see Cass R. Sunstein and Richard H. Thaler, ‘Libertarian Paternalism is Not an Oxymoron’, *The University of Chicago Law Review* 70 (2003), pp. 1159–202; for an overview, see Richard Noggle, ‘The Ethics of Manipulation’, *Stanford Encyclopedia of Philosophy*, <<https://plato.stanford.edu/entries/ethics-manipulation/>> (2018).

³⁵Many former colleagues at Syracuse University offered helpful feedback on earlier versions of this article; I wish to thank, in particular, Ben Bradley, Hille Paakkunainen and Travis Timmerman. Thanks also to two anonymous referees for *Utilitas* for their comments. This research was supported by a grant from China’s Fundamental Research Funds for the Central Universities programme.