

PRIORITY AND DESERT

Post-print of a paper published in the BSET conference issue of *Ethical Theory and Moral Practice* 16/5 (November 2013): 939-51. The final version is available at Springer via
<http://dx.doi.org/10.1007/s10677-013-9420-9>

Matthew Rendall
Lecturer
School of Politics and International Relations
University of Nottingham
University Park NG7 2RD
United Kingdom
44-(0)115-846-6231 (tel.)
44-(0)115-951-4859 (fax)
Matthew.Rendall@nottingham.ac.uk

Abstract: Michael Otsuka, Alex Voorhoeve and Marc Fleurbaey have challenged the priority view in favour of a theory based on competing claims. The present paper shows how their argument can be used to recast the priority view. All desert claims in distributive justice are comparative. The stronger a party's claims to a given benefit, the greater is the value of her receiving it. *Ceteris paribus*, the worse-off have stronger claims on welfare, and benefits to them matter more. This can account for intuitions that at first appear egalitarian, as the analysis of an example of Larry Temkin's shows. The priority view, properly understood, is desert-adjusted utilitarianism under the assumption that no other claims pertain.

Keywords: Priority view; egalitarianism; desert-adjusted utilitarianism; risk

1 Introduction: A New Challenge to the Priority View

Recently the priority view has given egalitarians a run for their money. In Derek Parfit's (2002) seminal statement, benefits matter more, the worse off their recipients. Equality may be instrumentally valuable, but is of no intrinsic importance. Prioritarians give priority to the worse-off because their *absolute* position is worse. It is as important to help them whether or not there are other people who are better off. Prioritarians favour a more equal distribution of resources, but this is because those with fewer resources are usually worse-off. Not only does this reasoning seem to capture many people's true concerns, but it avoids a key objection to teleological egalitarianism: that the latter implies there would be *something* good about dragging the fortunate down to the level of the unfortunate. Prioritarians favour redistribution only when it *does* help the worse-off (McKerlie 1996: 288).

Prioritarianism has caught on rapidly, to the point that Brad Hooker (2011) describes it as 'the dominant form of consequentialism'. This lends great importance to Michael Otsuka's, Alex Voorhoeve's and Marc Fleurbaey's recent critique of the priority view. In a series of papers, they have presented arguments and cases which I restate here in simplified and schematized form, modifying some of the terminology. None of my changes, I believe, affects anything of substantive importance. Suppose we have two small boys, Ernie and Bert, both with severely restricted vision that is getting worse. We have recently discovered an operation that has an equal chance of completely restoring their vision—raising their overall welfare to 1.0—or failing and allowing a continued decline to blindness at 0.65. Alternatively, we have a tried-and-true operation which will, with certainty, stabilize their vision, leaving both at 0.8. Choosing the new-fangled operation means balancing the chance of a gain against the risk of a loss for each boy. Given that the gain is greater, many would consider it

appropriate to gamble. Let us call this the *Two-Person Intrapersonal Risk Case*.¹

Suppose next that the new-fangled operation can only be performed on Ernie. It is sure to succeed, but will leave no money to perform the tried-and-true operation on Bert, who will go blind. In this *Interpersonal Trade-off Case*, with the new-fangled operation Ernie's welfare will rise to 1 and Bert's decline to 0.65.² Surely we are more likely in this case to perform the tried-and-true operation (Otsuka and Voorhoeve 2009: 171-75; Voorhoeve and Fleurbaey 2012: 384-88).

Operation	Patient	State of the world (equiprobable)	
		S ₁	S ₂
Two-Person Intrapersonal Risk Case			
New-fangled operation	Ernie	1	0.65
	Bert	1	0.65
Tried-and-true operation	Ernie	0.8	0.8
	Bert	0.8	0.8
Interpersonal Trade-off Case			
New-fangled operation	Ernie	1	1
	Bert	0.65	0.65
Tried-and-true operation	Ernie	0.8	0.8
	Bert	0.8	0.8

Table 1: Utility distributions in the Two-Person Intrapersonal Risk and Interpersonal Trade-off Cases (adapted from Voorhoeve and Fleurbaey 2012).

This creates a puzzle for prioritarrians. The expected utility of the new-fangled operation is greater than that of the tried-and-true operation, and the difference is equally great in both cases. Utilitarians will choose the new-fangled operation in each. Since prioritarrians accord greater weight to benefits that come at lower levels, they might prefer the tried-and-true operation unless the gains from the new-fangled operation are significantly greater. What the priority view has trouble explaining is

¹ Otsuka and Voorhoeve's version of this case involved decision-making on behalf of a single person. They call it the 'one-person case', as do Voorhoeve and Fleurbaey (2012). I have preserved the structure of the case while expanding the effects to both people. This is consistent with Otsuka's and Voorhoeve's observation that 'whatever claims we make ... about what one ought to do in cases involving single persons apply, *mutatis mutandis*, to groups of identically fated people created by such replication' (2009: 175-76 n. 8).

² The Interpersonal Trade-off Case corresponds to what Otsuka and Voorhoeve call their 'multi-person case with certainty' (2009: 179 n. 16).

why we give *greater* priority to preventing blindness in the Interpersonal Trade-off Case. If only absolute levels of wellbeing matter, then one chance in two of helping two badly off people has no greater expected value than the certainty of helping one.

Does this sink the priority view? Otsuka and Voorhoeve evidently think it does, holding that '[e]galitarian or otherwise comparative views...offer a better account of why those who are worse off than others have greater moral claims' (2009: 198). In response, Parfit has suggested that prioritarrians should care about the *prospects* of the worst-off, as well as how they fare in outcomes. In the Intrapersonal Risk Case, Ernie's and Bert's prospects will remain equal whatever we choose. In contrast, if we choose the new-fangled operation in the Interpersonal Trade-off Case, Bert's prospects will be worse than with the tried-and-true one. This concern with improving the prospects of the worse-off justifies our reluctance to gamble in cases where some stand to gain and others to lose (2012: 429-36). Yet Parfit offers no explanation of why people's prospects should matter in and of themselves. The claim is hardly self-evident. As Otsuka observes, 'it would seem, on the face of it, that how well or badly a person's life actually ends up going is what matters, rather than how well or badly it is expected to go' (2012: 379).

In this paper I argue that competing claims can serve as the basis for a *relational* prioritarianism that explains our reluctance to risk losses to some for the sake of gains to others. The priority view is best understood as part of desert-adjusted utilitarianism, in which the strength of claims affects the value of benefits.

2 Why does it matter that some are worse off than others?

Otsuka and Voorhoeve attribute the change in our judgment in cases such as the foregoing ones to the separateness of persons. In the Two-Person Intrapersonal Risk Case, for both parties 'the prospect of a greater gain is the desirable flip side of

exposure to the risk of a lesser loss'. In contrast, in the Interpersonal Trade-off Case, Ernie stands to gain and *Bert* to lose. That the latter has no chance of coming out ahead explains, they argue, why we are less willing to accept such trade-offs (2009: 179-80, 185). Many believe that a key reason—perhaps *the* reason—that the separateness of persons matters is that more compensation is possible within lives than between them (Jeske 1993). But as Voorhoeve and Fleurbaey point out, this is not true when people have only a *chance* of receiving a benefit: If the new-fangled operation fails, neither Ernie nor Bert will receive any compensating benefit (2012: 381-82). In Dylan's words, failure's no success at all.

Alternatively, one might think that the outcome in the Interpersonal Trade-off Case is worse because the outcome is unequal. But that explanation runs into the levelling down objection (Otsuka and Voorhoeve 2009: 182-83). Moreover, as Voorhoeve and Fleurbaey (2012: 390) show, we can formulate versions of the trade-offs with identically unequal outcomes, in which the new-fangled operation still seems more defensible in the intrapersonal than the interpersonal case. Suppose that with the new-fangled operation, Ernie, as before, has a fifty percent chance of regaining full vision, and a fifty percent chance of its getting worse. With the tried-and-true operation we can stabilize his sight at its present level. Bert's vision has stabilized but is inoperable; he will remain at 0.8 whatever we do. Call this the *Single-Person Intrapersonal Risk Case*. If Ernie's potential gains from the new-fangled operation sufficiently outweigh his potential loss, it would seem just as appropriate to gamble as in the Two-Person Intrapersonal Risk Case.

Suppose now, in contrast, that we can only perform either the new-fangled or the tried-and-true operation on both boys. The new-fangled operation, if it succeeds, will restore Ernie's eyesight and stabilize Bert's; if it fails, it will merely stabilize

Ernie's eyesight, allowing Bert to go blind. While this example is more fantastic than the others, I see no reason to expect our reactions to be systematically biased. Call this the *Interpersonal Risk Trade-off Case*.

Operation	Patient	State of the world (equiprobable)	
		S ₁	S ₂
Single-Person Intrapersonal Risk Case			
New-fangled operation	Ernie	1	0.65
	Bert	0.8	0.8
Tried-and-true operation	Ernie	0.8	0.8
	Bert	0.8	0.8
Interpersonal Risk Trade-off Case			
New-fangled operation	Ernie	1	0.8
	Bert	0.8	0.65
Tried-and-true operation	Ernie	0.8	0.8
	Bert	0.8	0.8

Table 2: Utility distributions under the Single-Person Intrapersonal Risk Case and Interpersonal Risk Trade-off Case (adapted from Voorhoeve and Fleurbaey 2012: 387).³

In both cases, the new-fangled operation produces an unequal outcome. In the Single-Person Intrapersonal Risk Case, Ernie will be better *or* worse off than Bert. In the Interpersonal Risk Trade-off Case, Bert is sure to end up worse off than Ernie. In each case, the distribution of outcome utilities is just as unequal. Yet in the latter, we're more likely to think it best to perform the tried-and-true operation. Why?

One reason might be that while the distribution of *outcome* utility is as unequal in both cases, the distribution of *expected* utility is not. In the Single Person Intrapersonal Risk Case Ernie stands to gain or lose, while Bert's prospects remain unchanged. In the Interpersonal Risk Trade-off Case, Ernie has only a chance of gain, while Bert has only a chance of loss. Egalitarians, Voorhoeve and Fleurbaey observe, might consider the new-fangled operation in the Interpersonal Case more objectionable 'because it involves a more unequal distribution of the chances of being

³ Voorhoeve and Fleurbaey call their equivalent cases the 'Two-Person Intrapersonal Trade-off Case' and the 'Two-Person Interpersonal Trade-off Case'.

relatively advantaged' (2012: 392-393). But this explanation falls foul of the levelling down objection. If unequal chances are bad *because they are unequal*, then it is also objectionable in the Single-Person Intrapersonal Risk Case that Ernie's expected utility is higher than Bert's. Because performing the tried-and-true operation on Ernie would equalize their prospects, there would be something good about choosing it, even though it offers Ernie less expected utility. If the expected utility of the two operations were close enough, Bert's *mere existence* could determine which operation is best to perform on Ernie, even if Bert lived in another galaxy. That is absurd.

Fortunately, the authors have another explanation. Whenever we distribute benefits and burdens among more than one person, the parties have competing claims on our solicitude. 'Those who are relatively worse off have stronger claims to a given increment of improvement', as Otsuka and Voorhoeve note, 'simply by virtue of the fact that it is, other things equal, harder to justify improving the situation of someone who is better off rather than someone who is worse off...[T]here is not even a pro tanto objection on this approach to a benefit going to the better off person when it cannot instead go to the worse off person' (2009: 183-84). In the Single-Person Intrapersonal Risk Case, our decision affects only Ernie. Only he has any claims. That his expected utility will be higher than Bert's if we refuse to intervene is irrelevant. In contrast, both parties have equal claims on us in the Interpersonal Risk Trade-off Case. It is harder to justify giving all the chances of benefit to Ernie than dividing them between him and Bert (Voorhoeve and Fleurbaey 2012: 397-98).

This Competing Claims View is a much better argument. Parfit responds by maintaining that the priority view can also explain our reluctance to gamble in interpersonal cases by appealing to the value of chances. We can believe, he writes, that '[w]e have a stronger reason to give people benefits, or expectable benefits, the

worse off these people are, or expectably are' (Parfit 2012: 436). If we choose the new-fangled operation in the Single-Person Intrapersonal Risk Case, Ernie's expected utility will be higher, and Bert's will be the same. But if we do the same in the Interpersonal Risk Trade-off Case, Ernie's expected utility will be 0.9, whereas Bert's will be only 0.725. Unless we choose the tried-and-true operation, Bert will be expectably worse-off. Parfit's argument gives the intuitively correct results. But he does not explain why we should *care* about prospects independently of outcomes.

Voorhoeve and Fleurbaey concede that even if individuals obtain any welfare from the chance of obtaining a good, it 'evaporates once it is clear that this chance is unrealized'. They suggest, however, that people might have 'an interest in being treated fairly', and that *this* would be better served by a fair chance whether or not they actually win a gamble (Fleurbaey and Voorhoeve 2011a: 16-17; Voorhoeve and Fleurbaey 2012: 396). The same argument is not available to Parfit. Fairness is a relational concept. It makes no sense to say that Bert is being treated unfairly if he is not compared with Ernie at all. In contrast, the Competing Claims View can appeal to Ernie's and Bert's equal claims on a given amount of expected utility.

Nevertheless, it seems unlikely that it can bear much weight. It implies that for the sake of affording people equal chances, it could be better to choose an option that we know will have *more* unequal outcomes (Fleurbaey and Voorhoeve 2011a: 19-20). As Fleurbaey and Voorhoeve suggested in an earlier draft of their paper, this seems plausible only if the difference in the outcomes is marginal (2011b: 15)—if that.

In the next section I argue that the Competing Claims View can explain why the *outcomes* of choosing the new-fangled operation are better in the Intrapersonal cases. We will see this better through examining desert-adjusted utilitarianism.

3 Desert-Adjusted Utilitarianism

Desert is often considered a matter of merit. On this view, to deserve something you must ‘come to earn it, or to be able to claim credit (or discredit) for it’. This is what distinguishes ‘genuine desert’ from other sources of claims. Nevertheless, we do speak of deserving things we do not merit (Miller 2003: 27; Olsaretti 2003: 5). Bert may deserve a tricycle not because he cleans his room, but simply because Ernie has already received one.⁴ Fred Feldman cites various factors that can affect what we deserve, including not only moral worth, but also how much we have received in the past (1997a: 162). Desert in this sense amounts to a *claim*, whether earned or not.

Some claims are absolute. You deserve a fair trial, whatever they do with your fellow citizens. But in *distributive* justice, so I will argue, all claims are comparative. Distribution involves dividing benefits and burdens. If we assign one person a lighter burden, then another person must bear a heavier one. If we give some people more of a limited good, others must receive less. In a rich society, people can lay claim to more resources than in a poor one, because there is a larger pie to divide. Even if we believe that people deserve whatever the market allows them to earn, how much this amounts to depends on the overall size of the economy (Hurka 2003; Miller 2003: 31-32). Some suggest that people deserve a certain amount of welfare simply ‘in virtue of *being a person*’ (Feldman, 1997b: 203, emphasis in original). But is this true? Why just so much and no more or less (Olsaretti 2003b: 397)? Alternatively, we might think that everyone deserves some basic minimum (Olsaretti, 2003a: 20). But do people deserve just as much utility in an absolute sense during a famine or a natural disaster—when giving the full allotment to some would mean giving far less to others? Even if this were true, it would surely be trumped by comparative desert when

⁴ By this I mean that Bert has a stronger *claim* on a tricycle. This assertion can be challenged, as Patrick Tomlin (2012) shows. I cannot go into the matter here.

we made our decision. Claims to scarce goods are inherently comparative. If we do *not* compare them, then we are bound to give some people less than they deserve.

Distributive justice involves goods that can be distributed or redistributed. Only people who can receive a benefit have claims to it. Suppose that Peter has \$10, Paul has \$6, and Marvin has \$1, each through no merit or fault of his own. We have a free lunch to distribute. *Ceteris paribus*, it would be best for Marvin to receive it. However, Marvin is on Mars. By the time he returns to Earth, the lunch will have spoiled. This is hard luck for Marvin. He has the greatest claim on our sympathy. But this in no way affects Peter's and Paul's claims to the lunch. Paul's claim is just as strong as if Marvin did not exist. Similarly, when food and medical treatment are apportioned according to people's claims, the distribution is as fair during a famine as in a rich society. This does not mean that the outcome is as *good*. The size of the pie matters, as well as how we cut it (McLeod, 2003: 143). But fairness matters too. A society in which everyone gets what she deserves is better than an unjust society that contains the same amount of welfare. This is not because the just society's welfare is more valuable to its members. Rather, it has 'more value for the world' (Feldman 1997b).

One way to capture the total value of an outcome might be to count both welfare's personal value to each recipient and also the extent to which people receive their just deserts. A poor world in which goods are fairly shared would rank high on 'desert value', but low on personal value. Conversely, a rich world full of injustice would rank low on desert value, but high on personal value. We would then sum the values to calculate the world's intrinsic value (cf. Arrhenius 2006). However, this approach raises problems. On the one hand, we might hold that the desert value of welfare is always positive, though greater when it goes to the right person. This would

imply, however, that even a horrendously unjust world could be better than no world at all, not only all things considered, but also *in terms of desert*. Surely that cannot be right. Alternatively, perhaps it can be *bad* if people receive less than they deserve. If that is true, if a world contained enough injustice, it would be better for it not to exist, even if all of its inhabitants had lives that were worth living (Vallentyne 1995: 212). Feldman welcomes this as a way of avoiding the repugnant conclusion (1997b). But this seems perverse. How could it be bad, all things considered, for a world to exist in which everyone's life was worth living?

A better way to understand desert, as Peter Vallentyne has proposed, is as a coefficient modifying the value of welfare. 'For any given increment in welfare...the desert-adjusted welfare increment is the *product* of [the recipient's] marginal desert factor and the welfare increment' (1995: 213; emphasis in original). Additional welfare is always good, or at least of neutral value. But the stronger a person's claims, the greater value of a given increment of welfare she receives has for the world. 'Suppose we start with equality', Shelly Kagan observes (1999: 303).

We have—or at least many of us have—the intuition that if A is worse off than B, and we can help one or not both by some fixed amount, then it is better to help A. Helping A makes a greater improvement in the overall goodness of the outcome than does helping B; we can say, in this light, that A has a greater claim than B.

If someone has less than her share of benefits, then it follows that others must have more. Benefits that go to her will improve the world more than benefits that go to them.

Moreover, the moral importance of each increment of welfare increases the further the person falls short of what she is due. Conversely, though benefits that

exceed people's desert still improve the world, they do so less and less. If A and B have equally strong claims, then it is better to give a benefit of a given size to A. This might be, Kagan suggests, because A is equally deserving but comparatively worse off than B, or because A is further from what she deserves in an absolute sense, or both (1999: 303). I have argued that *all* desert is comparative. A has as strong claims on a given level of welfare as B, but she is worse off. She thus deserves more welfare and B deserves less. That renders B's surplus welfare less valuable. Levelling down B would still destroy value, and make the world worse. But the welfare would have more value for the world if it were transferred to A.

Each additional increment of welfare affects the strength of people's claims. If we can give one unit of welfare to A or to B, and both have equally strong claims, then it does not matter whom we pick. But if we may either give 30 to A and 70 to B, or 50 to both, then the second choice is better. Suppose we nevertheless choose the first option. The first unit of utility that A and B receive counts the most. The second counts a little less, and the third less still, with the marginal moral value steadily declining. Let us say, borrowing terms from Gustaf Arrhenius (2003: 227), that each unit a person receives up to her fair share is *over-deserved*, and each unit a person receives in excess of her fair share is *under-deserved*. A's and B's initial increments of welfare are *equally* over-deserved, and the first thirty received by B count just as much as A's (cf. Holtug 1999: 27; Porter 2011: 201-2; Vallentyne 1995: 208). Above 30, however, the marginal value of additional increments for B continues to decline, and from the fifty-first unit onward they are under-deserved. True, we are imagining giving B all 70 units at once. We cannot say which unit came first and which came last. But we can still say, as Thomas Porter notes, that the units had different weight, just one can 'say that what enables me to pay off a £20 debt is just one of two £20

notes that fall into my lap in an unexpected windfall' (2011: 205). Conversely, each unit below 50 foregone by A is increasingly over-deserved, and the opportunity cost of giving it to B grows.

4 Desert and the Priority View

We now have a natural justification for the priority view. The marginal value of each unit of welfare increases or decreases in proportion to the strength of the recipient's claims (Vallentyne 1995: 208-9). The worse off someone is in comparison to others, the stronger her claim to receive a given benefit, and the greater the value of her receiving it. Conversely, as she becomes better off, the weaker her claims, and the less each additional unit matters. Priority, then, is inherently comparative, as its etymology suggests, and we should be *relational*, or *comparative*, *prioritarians*.⁵

Ingmar Persson has pioneered this view, but argued that unjust *inequality* reduces the value of benefits to the better-off (2001: 35-36; 2008: 297-300). Desert-adjusted utilitarianism offers a simpler explanation: *Ceteris paribus*, it is always better if a benefit goes to the worse-off, because this increases its value. Nevertheless, it is not always best, all things considered, for everyone to get just what she deserves. Sometimes it is better to distribute to people who do not have the strongest claims, if they will benefit much more, or there are many more of them (Feldman 1997a: 161).⁶ Like the standard version of the priority view, desert-adjusted utilitarianism gives only weighted priority to the worse-off.

Moreover, there is one kind of claim that does not affect the goodness of a person's receiving a benefit or avoiding a burden. This is when the justification for the claims is simply that they promote an independently better outcome. Sometimes we

⁵ The first term is Ingmar Persson's, and the second is Paula Casal's (2007: 309).

⁶ Compare John Broome's argument that fairness is 'a relative matter' that must sometimes be weighed against maximizing the 'general good' (1984: 43-45). We perform this balancing, I am arguing here, by treating claims as a coefficient of utility.

create entitlement claims for instrumental reasons.⁷ In these cases, the recipients are not more deserving. They would have no grounds for complaint if we awarded someone else the entitlement instead. Children, for example, do not *deserve* to inherit money. But it might be socially valuable to create and enforce certain inheritance rights. If so, the reason would be not that the benefits matter more, but rather that recognizing such claims results in greater utility, or in an overall utility distribution that better reflects desert. When legacies benefit the better-off, as they usually do, those benefits, other things being equal, count for *less* than their face value.

Desert-adjusted utilitarianism explains why we might risk the new-fangled operation in the intrapersonal cases, yet prefer the tried-and-true operation in the interpersonal ones. We have benefits to distribute, and Ernie and Bert, with equal levels of welfare, have equal claims on our assistance. In both the Single-Person and Two-Person Intrapersonal Risk Cases, there is no benefit *to be divided*, so the question of desert never arises. Benefits to Ernie and Bert are not weighted, and carry their utilitarian face value. We prefer the new-fangled operation, because its expected utility is higher. In the Two-Person Intrapersonal Risk Case, both boys will end up with the same level of welfare; in the Single-Person Intrapersonal Risk Case, one will end up better off. Nevertheless, since any gains that Ernie makes will not come at Bert's expense, we still assess benefits to Ernie at their face value, and maximize not only raw, but also desert-adjusted utility (which in this case amount to the same thing) by opting for the new-fangled operation.

In the Interpersonal Cases, there *is* a benefit to be divided. Ernie and Bert have equally strong claims on benefits, and thus deserve *equal* benefits. If we opt for the new-fangled operation, we give Ernie more than he deserves, and Bert less. In the

⁷ I am grateful to Erin Taylor for raising this point.

Interpersonal Trade-off Case, it is certain that Bert will end up with less welfare than he deserves, and Ernie with more. The portion of Ernie's welfare that exceeds his fair share will be under-deserved and count for less than its face value. Moreover, there is the *opportunity cost* of the welfare that Bert would have received had we opted for the tried-and-true operation. This foregone welfare—the increments from 0.65 up to Bert's fair share of 0.8—is over-deserved, and would have counted for more.⁸ With the tried-and-true operation, total 'raw' utility is lower, but since competing claims are satisfied, *desert-adjusted* utility may be higher. Similarly, in the Interpersonal Risk Trade-off Case, if we opt for the new-fangled operation, whether it succeeds or fails, Bert will be avoidably worse off than Ernie through no fault of his own. This reduces the value of the *outcome* compared with that of the tried-and-true operation. Even though the raw expected utility of the new-fangled operation is higher, its desert-adjusted expected utility may be lower.

5 Parfit's objections

Responding to Otsuka and Voorhoeve's critique, Parfit denies the distinction between the outcome value of intrapersonal and interpersonal gambles. He argues along the following lines: In the Single-Person Intrapersonal Risk Case, if we choose the new-fangled operation, we are judging that restoring Ernie's vision from 0.8 to 1.0 would do more to improve the outcome than preventing its decline from 0.8 to 0.65. Suppose now that in the Interpersonal Trade-off Case, it is Bert who stands to gain, rather than Ernie. If we still choose the tried-and-true operation—as it seems we should—we are judging that stabilizing Ernie's vision at 0.8 would do more to improve the outcome than restoring Bert's vision to 1.0. We now seem committed to the following judgement:

⁸ One might think that Bert's 0.65 units of welfare would also be over-deserved and worth more than their face value, thus *increasing* the value of the outcome. Both boys, however, have equally strong claims to the first 0.65 units of welfare.

Restoring Ernie's vision > Stabilizing Ernie's vision > Restoring Bert's vision

Yet we have no reason to favour Ernie. How can this be (Parfit 2012: 427-8)?

Parfit's argument shows that if we wish to retain the foregoing judgements, we must accept that how an outcome arises can affect its value. In the intrapersonal cases, both the new-fangled and tried-and-true operations bear their raw utilitarian face value. But our choices in the interpersonal cases affect the strength of the parties' claims, and thus the desert-adjusted outcome value. This need not result in intransitive value judgements (cf. Temkin 1987). Instead, we can maintain transitivity by individuating the outcomes (Broome 1991, ch. 5):

Restoring Ernie's vision *at no one's expense* > Stabilizing Ernie's vision *at no one's expense* > Restoring Bert's vision *at Ernie's expense*

Some utilitarians will deny that the value of a given benefit can vary according to circumstances. But that is just what desert-adjusted utilitarianism entails.

We seem to have good reason to distinguish between the outcome values of the gambles. But Parfit raises a second objection. Suppose that in the Interpersonal Trade-off Case, rather than simply awarding Ernie the new-fangled operation, we give Ernie and Bert equal chances of receiving it. In this *Interpersonal Lottery Case*, each boy has an equal chance of a loss or a greater gain—indeed, exactly the same odds that he faced in the *Two-Person Intrapersonal Risk Case*. In each case, we could defend our choice to both boys as being in their *ex ante* interest. How, one might wonder, could one justify choosing the new-fangled operation in one case, and the tried-and-true operation in the other (Parfit 2012: 413-16)?

Operation	Patient	State of the world (equiprobable)	
		S ₁	S ₂
Interpersonal Lottery Case			
New-fangled operation	Ernie	1	0.65
	Bert	0.65	1
Tried-and-true	Ernie	0.8	0.8

operation	Bert	0.8	0.8
Two-Person Intrapersonal Risk Case			
New-fangled operation	Ernie	1	0.65
	Bert	1	0.65
Tried-and-true operation	Ernie	0.8	0.8
	Bert	0.8	0.8

Table 3: Utility distributions under the Interpersonal Lottery Case and Two-Person Intrapersonal Risk Case.

The difference is not that in the Lottery Case the new-fangled operation results in inequality. That explanation would face the levelling down objection. Rather, only in this case do we have a benefit *to divide*. If we choose the new-fangled operation, one boy will receive more than his share by sheer luck. Yet this outcome was avoidable, and the loser—whoever that turns out to be—will have as strong a claim to the benefit as the winner (Fleurbaey and Voorhoeve 2011a; Otsuka 2012: 374). Part of the winner’s utility will thus be under-deserved and contribute less value to the outcome than if it were evenly distributed.

6 Temkin’s asteroid

We see similar considerations at work in an example of Larry Temkin’s. In one scenario we can divert an asteroid full of valuable minerals to a planet. This planet, unlike any others, already has a rich endowment of resources, with many other asteroids having struck it in the past. In the second scenario, the planet’s inhabitants are equally well-off in absolute terms, but ‘they have been terribly *unlucky*’: while every other planet has benefited from numerous asteroids, theirs has never received a single one. As a result, it is poor in resources, ‘and they have worked incredibly hard to achieve their well-being’. If only absolute levels of welfare matter, Temkin argues, prioritarrians will see no more reason to divert the asteroid in the second scenario. Yet here ‘the people are the victims of natural unfairness’, whereas ‘[i]n the first, they are the beneficiaries of it. To my mind, however much I should sacrifice for those below

in the first scenario, I should sacrifice more, if necessary, in the second scenario, where the situation exerts a greater claim on me' (2003: 69-70, emphasis in original).

Temkin thinks that his example supports egalitarianism over prioritarianism. In fact, desert claims are doing the work. First, we should note that he has loaded his example with an extraneous consideration—the people in the second case have *worked hard* to scrape out what they have. When we compare how deserving people are, we often consider their contribution, how hard they have worked, or both (Hurka 2003: 57-58). In the second case the people have contributed a great deal to their wealth, instead of having mineral manna dropped into their laps, and they have certainly put in more effort. Even if we don't believe that contribution or effort themselves deserve reward, we may suspect that because the people have had to struggle to overcome their natural poverty, they have enjoyed less happiness. But in fact Temkin has stipulated that they are just as well off as the inhabitants of the other planets. Any suffering and frustration from the poor endowment must have been offset by something else. At this point, we may lose the intuition that it is more important to help them. Alternatively, we may insist that contribution and effort themselves increase desert.

Now, if this were a question of choosing which planet would receive the asteroid, it would be a straightforward case of distributive justice. If the people on the second planet were more deserving, the impersonal value of their receiving the minerals would be higher. Temkin, however, has not stated the question in this way. He is asking in which case we should feel more *motivated*. Since there is only one planet in each scenario, it may seem hard to see how comparative claims come into play. Yet we can still make hypothetical comparisons, if only implicitly. Suppose somebody reproached us in the first case for not making more of an effort to divert the

asteroid. We might retort, ‘We’d be trying a lot harder if they hadn’t already had so many asteroids land there’. Moreover, note Temkin’s remark that he should be prepared to *sacrifice* more in the second case because it makes a stronger claim on *him*. The claims of the second planet’s people are weightier compared to our own; welfare that goes to them counts more in comparison with welfare that goes to us.

What matters is the strength of the claims of any parties who can receive a good, not how they stack up against third parties. In the second case, Temkin tells us, ‘every other populated planet is *loaded* with natural resources, and has benefited from *countless* mineral-rich asteroids’ (2003: 69, emphasis in original). But so what? The relevant comparison is not between the claims of the planet that can receive the asteroid and those of other planets that can’t, but between that planet’s claim on our assistance and *our* claim to be spared the effort of assisting its inhabitants. That intervening would make the planet better off compared to other planets is beside the point. Otherwise, how hard we should work to help this planet would depend on the endowments of planets of which we are entirely unaware. If we discovered inhabited planets in the Andromeda galaxy that were mineral-poor, we would have some reason to do less for this one (cf. Parfit 2002: 88). This would be as silly as allowing Bert’s mere existence at the other end of the universe to determine which operation to perform on Ernie—and for the same reason.

7 Priority and population

Desert-adjusted utilitarianism solves another puzzle that has bedeviled the priority view. Prioritarians hold that benefits of a given size matter more when they go to people who are worse off. If bringing a person into existence benefits this person, then it may seem that it must be particularly valuable to bring badly-off people into existence. We will contribute more value if we create a vast throng whose lives are

barely worth living than a smaller number, all of whom are very well off. This resembles Parfit's (1987) famous repugnant conclusion, but here, we are told it would be better to create the impoverished multitude not only if aggregate utility is no greater, but even if it is somewhat *less* (Holtug 1999: 32-33; Persson 2001: 33-34). That cannot possibly be right.

Desert-adjusted utilitarianism avoids this objection. Benefits matter more when they go to the worse-off because it is harder to justify giving them to a well-off person when one could use them to improve the position of a badly-off person. But we can't *improve* a non-existent person's position at all. This is not to deny that merely possible people have claims on us. Notably they have a strong claim not to be brought into existence if their lives would be miserable (cf. Sterba 1980: 434-39). But they do not have claims *in virtue of being worse off*.

It is true that adding any lives worth living increases total utility—both raw and desert-adjusted. If we add a population of people with impoverished lives to an existing population that is better off—as occurs in the mere addition paradox—we will increase total value. Moreover, in so doing we will create a situation in which the former have stronger claims to welfare. We will increase value further by improving their lives at the better-off group's expense. Desert-adjusted utilitarianism does not solve Parfit's paradox; other arguments are needed for that.⁹ But neither does it make the paradox worse.

8 Conclusion

Otsuka, Voorhoeve and Fleurbaey have undermined the orthodox priority view. With their Competing Claims View, they have also shown a way to recast it. Distributive justice depends on the extent to which competing claims are satisfied—

⁹ See my criticism of Feldman above. I have proposed a solution to the mere addition paradox in Rendall (2012).

the extent to which people receive their just deserts. Claims modify the moral value of welfare. People don't deserve to be worse-off through no fault of their own, and if we have a benefit to distribute, those who are worse-off have, *ceteris paribus*, stronger claims. This makes it better to give them benefits of a given size. Bernard Gans has recently argued that if prioritarianism is to correspond more closely to our everyday moral intuitions than utilitarianism, it must capture desert-related considerations (2011: 124). Gans is right. The priority view, correctly understood, is simply desert-adjusted utilitarianism applied to a case where no other claims pertain.¹⁰ If imprudence and lack of self-discipline reduce people's claims, then this reduces the value of compensating them for losses. Desert-adjusted utilitarianism will capture luck egalitarians' distinction between brute and option luck. But if we are not responsible for our characters or our efforts, then it may be that the only competing claims arise from differences in well-being. In either case, the goodness of outcomes depends on total welfare, adjusted for desert.

Let me close with a caveat. Desert-adjusted utilitarianism—including the priority view—tells us which outcomes are better than others. Naturally, this helps explain our intuitions about cases such as the ones discussed here. Yet desert-adjusted utilitarianism may not tell us directly what to do. It will do so only if we should always seek to maximize value directly. It may be, however, that we will maximize value by following some other principle. I believe that under conditions of moderate scarcity, desert-adjusted utilitarians should adopt a priority-weighted version of the *sufficiency view* as their 'prima facie principle' in distributing some resources (cf. Hare 1981). But that is the subject for another paper.

¹⁰ Cf. Shelly Kagan's similar argument about egalitarianism (1999).

Acknowledgement I thank Gerald Lang and audience members at the 2012 annual conference of the British Society for Ethical Theory for comments.

References

- Arrhenius G (2003) Feldman's desert-adjusted utilitarianism and population ethics. *Utilitas* 15: 225-36
- Arrhenius G (2006) Desert as fit: an axiomatic analysis. In McDaniel K, Raibley JR, Feldman R and Zimmerman MJ (eds.) *The good, the right, life and death: essays in honour of Fred Feldman*. Ashgate, Aldershot, pp 3-17
- Broome J (1984) Selecting people randomly. *Ethics* 95: 38-55
- Broome J (1991) *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Basil Blackwell
- Casal P (2007) Why sufficiency is not enough. *Ethics* 117: 296-326
- Feldman F (1997a) Adjusting utility for justice: a consequentialist reply to the objection from justice. In Feldman F *Utilitarianism, hedonism and desert: essays in moral philosophy*. Cambridge University Press, Cambridge, pp 154-74
- Feldman F (1997b) Justice, desert and the Repugnant Conclusion. In Feldman F *Utilitarianism, hedonism and desert: essays in moral philosophy*. Cambridge University Press, Cambridge, pp 195-214
- Fleurbaey and Voorhoeve (2011a) Decide as you would with full information! An argument against *ex ante* Pareto. Unpublished manuscript, 30 November
- Fleurbaey M and Voorhoeve A (2011b) On the evaluation of expectedly beneficial treatments that will disadvantage the worst off. Unpublished manuscript, 29 April

- Gesang, B (2011) *Klimaethik*. Suhrkamp, Berlin
- Hare, RM (1981) *Moral thinking: its levels, method and point*. Clarendon Press, Oxford
- Holtug N (1999) Utility, priority and possible people. *Utilitas* 11: 16-36
- Hooker B (2008) Rule consequentialism. In Zalta, EN (ed) *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition).
<http://plato.stanford.edu/archives/spr2011/entries/consequentialism-rule/>
 (accessed 14 November 2011)
- Hurka T (2003) Desert: individualistic and holistic. In: Olsaretti S (ed) *Desert and Justice*. Clarendon Press, Oxford, pp 45-68
- Jeske D (1993) Persons, compensation, and utilitarianism. *Philosophical Review* 102: 541-75
- Kagan S (1999) Equality and desert. In: Pojman, LP and McLeod O (eds) *What do we deserve? A reader on justice and desert*. Oxford University Press, New York, pp 298-314
- McKerlie D (1996) Equality. *Ethics* 106: 274-96
- McLeod O (2003). On the comparative element of justice. In: Olsaretti S (ed) *Desert and justice*. Clarendon Press, Oxford, pp 123-44
- Miller D (2003) Comparative and non-comparative desert. In: Olsaretti S (ed) *Desert and justice*. Clarendon Press, Oxford, pp 25-44
- Olsaretti S (2003a) Introduction: debating desert and justice. In: Olsaretti S (ed) *Desert and justice*. Clarendon Press, Oxford, pp 1-24
- Olsaretti S (2003b) Unmasking equality? Kagan on equality and desert. *Utilitas* 14: 387-400
- Otsuka M (2012) Prioritarianism and the separateness of persons. *Utilitas* 24: 365-80

- Otsuka M and Voorhoeve A (2009) Why it matters that some are worse off than others: an argument against the priority view. *Philos Public Aff* 37: 171-99
- Parfit D (1987) *Reasons and Persons*. Clarendon Press, Oxford
- Parfit D (2002) Equality or priority? In: Clayton M and Williams A (eds) *The ideal of equality*. Palgrave Macmillan, Basingstoke, pp 81-125
- Parfit D (2012) Another defence of the priority view. *Utilitas* 24: 399-440
- Persson I (2001) Equality, priority, and person-affecting value. *Ethical Theory Moral Pract* 4: 23-39
- Persson I (2008) Why levelling down could be worse for prioritarianism than egalitarianism. *Ethical Theory Moral Pract* 11: 295-303
- Porter T (2011) Prioritarianism and the levelling down objection. *Ethical Theory Moral Pract* 14: 197-206
- Rendall M (2012) Mere addition and the separateness of persons: a candidate for Theory X. Paper presented at the 12th International Conference of the International Society for Utilitarian Studies, 9 August
- Sterba, JP (1980) Abortion, distant peoples, and future generations. *Journal of Philosophy* 77: 424-40.
- Temkin, LS (1987) Intransitivity and the mere addition paradox. *Philos Public Aff* 16: 138-87
- Temkin LS (2003) Equality, priority or what? *Econ Philos* 19: 61-87
- Tomlin P (2012) On fairness and claims. *Utilitas* 24: 200-13
- Vallentyne P (1995) Taking justice too seriously. *Utilitas* 7: 207-16
- Voorhoeve A and Fleurbaey M (2012) Egalitarianism and the separateness of persons. *Utilitas* 24: 381-98