Economics and Philosophy, **19** (2003) 111–134 Copyright © Cambridge University Press

GIVING UP LEVELLING DOWN

Campbell Brown
Australian National University

1. THE ORTHODOX PICTURE

There is a great deal of common ground between egalitarians and prioritarians: theorists of those two persuasions will commonly find themselves recommending one and the same ranking of outcomes. However, as is often stressed, one should take care not to allow this superficial agreement to obscure the very significant disagreement that may occur deeper below the surface. For, although an egalitarian and a prioritarian may both say that, as it might be, outcome P is better than outcome Q, they will have very different reasons for saying that this is so; they may agree, that is to say, that P is better than Q, *all things considered*, yet disagree as to the relevant *respects* in which P is better, or indeed worse, than Q.

This feature of the debate between egalitarians and prioritarians is widely thought to be exemplified in cases of "levelling down" – that is, cases in which, roughly, some are made worse off, while none is made better off. Both egalitarians and prioritarians may say that levelling down makes things worse, all things considered. However, while egalitarianism implies that levelling down may nonetheless make things better in at least one respect (i.e., in respect of equality), prioritarians deny that it may make things better in *any* respect. This deeper disagreement is usually thought to tell in favour of prioritarianism, since levelling down benefits no one (by definition), and many people, so it is claimed, have the intuition that things cannot be better, in any respect, unless there is someone *for whom* they are

I owe a debt of gratitude to the following people for their many helpful comments, criticisms and suggestions: John Broome, Daniel Cohen, Marc Fleurbaey, Natalie Gold, Christian List, Dan Marshall, Wlodek Rabinowicz, Mike Ridge, Michael Smith and Larry Temkin. I should also like to thank the organizers and participants of the Workshop on Equality at the Fourth European Congress of Analytic Philosophy, at which event this paper was presented.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

better – that is, someone who is benefited by things being that way.[1] Hence, prioritarians are apt to invoke the so-called "Levelling Down Objection" in defending their position against egalitarianism. The objection is commonly thought to constitute the best reason one may have to adopt prioritarianism ahead of egalitarianism; the status of prioritarianism as a serious rival to egalitarianism, so the thought goes, lives or dies to a very great extent with the case of levelling down.

Such is the "Orthodox Picture", as I shall refer to it. It would appear that Bertil Tungodden (2003) does not find the levelling down objection very compelling:

> Admittedly, I am not sure how to rebut the intuition of people endorsing the levelling down objection, but in my view it does not really challenge egalitarianism as a viable normative position or the nature of the betterness relation you eventually have to defend. (p. 9)

Still, I think he does buy into the orthodox picture to a significant extent; for he seems to believe that, while there are some versions of egalitarianism that are vulnerable to the levelling down objection, prioritarianism is plainly and simply invulnerable to it.

As I shall argue, Tungodden is too quick in his (albeit qualified) acceptance of the orthodox picture, as are many others. For, although that picture hinges crucially on a certain notion – the notion of "being better in this or that respect" – that notion has hitherto received very little in the way of thoroughgoing analysis; it seems that authors writing in this area have been content to rest their arguments on an intuitive grasp of the notion.

The present paper grew out of an attempt to fill that gap – an attempt to analyse "in-a-respect betterness" in such a way as to fit snugly with the orthodox picture. However, the attempt failed. So, what I offer here instead is an analysis that crumples the orthodox picture into a ball and tosses it in the bin. On my account, thinking carefully about the nature of respects of value reveals two things: firstly, there is nothing absurd or obviously mistaken in the belief that levelling down may make things better in some respect or other; and secondly, the prioritarian is no less vulnerable to the levelling down objection than is the egalitarian. If the latter were true, then the levelling down objection could not serve to divide egalitarians and prioritarians, as it is commonly thought to do; hence, it would be entirely beside the point, at least so far as the debate between those positions goes.

Thus, the analysis suggests both a stronger and a weaker conclusion. The stronger conclusion is that the orthodox picture should be seriously revised and the levelling down objection entirely abandoned. I must

[1] Larry Temkin (2000) discusses this intuition at length.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

confess an inclination to accept this conclusion; but I cannot pretend that what I have to say here will constitute a knock-down argument in its favour. What I shall present here is the best analysis I can think of, and, so far as I can tell, it is presently the only game in town. But it remains to be seen whether rival analyses are forthcoming, and if so, whether any of these might be more congenial to the orthodox picture.

Still, there is a weaker conclusion to be drawn. And it is this: if philosophers and economists are going to continue framing this debate in terms of the orthodox picture – and, in particular, if prioritarians are going to continue appealing to the levelling down objection – then they *must* tell us exactly what it means, in their view, to say that one thing is better than another in a respect. I propose here an interpretation according to which what is said makes a nonsense of the orthodox picture. And, while some may complain of being misinterpreted, such a complaint should be taken seriously only if accompanied by an explanation of what a better interpretation would be. The levelling down objection may not be refuted; but its proponents sure have some explaining to do.

2. THE LEVELLING DOWN OBJECTION

Consider *Inegalitaria*, an imaginary mid-sized village in which, though everyone is at least reasonably well off, some are clearly more well off than others are. To make our description a little more precise, let us say that the population of Inegalitaria is divided roughly into two equal-sized classes such that (a) everyone within the same class is equally well off, and (b) everyone in one class is better off than everyone in the other. In short, Inegalitaria has significant inequality.

Now suppose that one day Inegalitaria is struck by a bomb. Fortunately, because of the wide availability of bomb-shelters and advance warning, there are no casualties. Nonetheless, the results of the bombing are devastating; the infrastructure of the village is entirely destroyed. Consequently, everyone is reduced to roughly the same low level of welfare; everyone is made rather badly off, but more or less equally so. Call the state of affairs that obtains prior to the bombing A, and the one that obtains post-bombing B. For the sake of convenience, I shall attach some numbers to those two states, representing the levels of welfare of the two groups in each: $A = (10, 2)$, $B = (1, 1)$.

When it soon comes to light that the perpetrators of the bombing are the dwellers of the nearby village *Egalitaria* – or, in other words, the Egalitarians – they are called upon to account for their actions. Seeing that the Inegalitarians are outraged by what they regard as the needless infliction of great harm, the Egalitarians reply, "Look on the bright side: at least things are better now in one important respect – there's

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

more equality." The Inegalitarians, however, are unmoved: "Are you crazy?" they retort, "things are not better now in any respect!"

Many people's intuitions are with the Inegalitarians, or at least so we are led to believe – for example, by Derek Parfit (2000, p. 98). Egalitarianism says that equality is valuable in itself; therefore, it implies that the bombing was, in at least one respect, an improvement – that state *B* is, in at least one respect, better than state *A*. Many people have the intuition that this implication is false; they judge, intuitively, that *B* is in no way an improvement relative to *A*. Thus, we may instantiate the levelling down objection as follows:

(1) egalitarianism implies that *B* is better than *A* in at least one respect;
(2) *B* is not better than *A* in any respect;
therefore,
(3) egalitarianism is false.

The target of this paper is premise (2). I shall try to show that once we get a clearer understanding of the egalitarian's claim – that levelled-down equality may be good in one respect – we shall see that, far from being obviously false, as the levelling down objection maintains, the claim looks quite reasonable – the kind of thing that most people would assent to. The crucial part in understanding the egalitarian's claim, I believe, is seeing what is meant by "being better in at least one respect". Thus, I turn now to a brief, general analysis of in-a-respect betterness claims.

3. AN ANALYSIS OF BEING BETTER IN A RESPECT

Suppose I say that, although Shaun Pollock is a better all-rounder than Chris Cairns, *in respect of bowling prowess*, Chris Cairns is the better all-rounder of the two, *all things considered*. Most people, I take it, (provided they know something about cricket[2]) would regard my claim as perfectly intelligible – as making good intuitive sense, at least. But what precisely do I mean? The basic idea seems to be something like the following. Other things being equal, the difference between Pollock and Cairns in respect of bowling prowess would suffice to make Pollock the better all-rounder. However, other things are not equal; there is some other difference (or some combination of other differences) that overshadows the contribution that the difference in bowling prowess makes to the comparison. Were it

[2] For the uninitiated, a brief cricket lesson may be in order. Put very simply, cricket is a sport in which two competing teams take turns at batting and bowling; while one side bats, the other bowls. The winning team is the one that performs best at the combination of these two tasks. Most cricketers specialize in just one or other task – either they are good at batting or bowling, but not both. However, there are a select few cricketers who excel at both facets of the game: these are the all-rounders.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

not for the countervailing difference (or differences) Pollock would be the better all-rounder, all things considered.

This suggests that a promising strategy for analysing my in-a-respect betterness claim might be to transform it into a counter-factual all-things-considered betterness claim. Such a transformation should help, it seems, because we already have a fairly clear idea of what it is for one thing to be better than another all things considered. As a first stab, we might imagine two counter-factual cricketers, Pollock* and Cairns*, such that (a) they differ in respect of bowling ability in precisely the way that Pollock and Cairns do (where Pollock and Pollock* are equal in this respect, as are Cairns and Cairns*), and (b) they do not differ in any other respect. And then we might say that the claim that Pollock is better than Cairns in respect of bowling prowess amounts to the claim that Pollock* is better than Cairns*, all things considered.

However, there is a slight wrinkle: the proposed analysis requires us to imagine the impossible. The conditions (a) and (b) are inconsistent; if Pollock* and Cairns* differ in respect of bowling prowess, then they cannot but differ in some other ways too. But the wrinkle is easily ironed out. For even if there must be some other differences between them, these other differences might not be relevant to the comparison. It might well be possible, therefore, to imagine Pollock* and Cairns* in such a way that they differ in respect of bowling prowess, yet do not differ in any other *relevant* respect. In this context, a relevant respect will be one such that the two players' differing in that respect will count towards one player's being a better all-rounder than the other; it will be, in other words, a better-all-rounder-making respect. So, a suitable revision of condition (b) seems to be this: Pollock* and Cairns* do not differ in any better-all-rounder-making respects besides bowling prowess.

Let us say that there are only two such respects: bowling prowess and batting prowess (actually, this cannot be far from the truth[3]). My proposal is that in order to see what the claim that Pollock is better than Cairns in respect of bowling prowess comes to, we need to entertain these three propositions: (a) Pollock* and Pollock do not differ in respect of bowling prowess; (b) Cairns* and Cairns do not differ in respect of bowling prowess; and (c) Pollock* and Cairns* do not differ in respect of batting prowess. Given (a), (b), and (c), Pollock is better than Cairns, in respect of bowling prowess, just in case Pollock* is better than Cairns*, all things considered. In this way, we can analyse an in-a-respect betterness claim by reducing it to a (counter-factual) all-things-considered betterness claim. As we shall see, one happy result of this approach is that it enables us to assess the former claim by thinking about the latter.

[3] See note 2.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

However, the analysis, as it stands, is incomplete. In our cricket example it was fairly settled ahead of time what were the "other relevant respects" – that is, the respects that needed to be held constant across the counter-factual objects (Cairns* and Pollock*) under comparison. For it is more or less a part of the concept of an all-rounder in cricket, that a player's status as such depends on both his batting and his bowling, and on nothing else besides. But one should not expect this to be so for every instance of the form "x is better than y in respect of r". Consider another example: two film critics – call them David and Margaret – agree that, as it might be, *The Phantom Menace* is a better film than *Star Wars* in one respect – say, in respect of having realistic special effects – even though there is no consensus at all between them regarding what other respects, besides this one upon which they agree, are relevant in evaluating movies. Suppose, for simplicity's sake, that David thinks the only two relevant respects here are "having realistic special effects" and "having an interesting story", whereas Margaret thinks they are "having realistic special effects" and "having convincing acting". It seems that, despite this difference, David and Margaret may each speak with the same meaning when they say some such thing as "*The Phantom Menace* is better than *Star Wars* in respect of special effects". So we do not want an analysis of in-a-respect betterness claims that ties the meaning of such a sentence to any *particular* set of other relevant respects.

The following definition, I believe, captures the spirit of the analysis given thus far, while also satisfying the desideratum just mentioned – namely, that of remaining neutral between different views regarding the relevant better-making respects.

(D1) For any two objects of evaluation x and y, and any respect r, x is better than y, in respect of r, if and only if there is some finite set of respects $r_1, r_2, \ldots, r_n$, such that, for any two objects of evaluation x^* and y^*, if
(1) x and x^* do not differ in respect of r;
(2) y and y^* do not differ in respect of r;
(3) x^* and y^* do not differ in respect of any $r_1, r_2, \ldots, r_n$;
(4) r is independent of $r_1, r_2, \ldots, r_n$;[4]
then x^* is better than y^*, all things considered.

The point of the analysis is to reduce sentences of the form "x is better than y, in respect of r" to sentences of the form "x^* is better than y^*, all things considered", thereby eliminating all talk of respects of betterness.

[4] The relevant notion of independence here is the of *differing* independently; it must be possible for two states to differ in respect of r without differing in respect of any $r_1, r_2, \ldots, r_n$. More precisely the requirement is this: there is some x^* and y^* such that conditions 1–3 are satisfied. The point of this independence condition is to prevent the right-hand-side of the biconditional (D1) from being made vacuously true.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

Conditions 1, 2 and 3 facilitate the reduction: 1 and 2 require that x^* and y^* differ in respect of r in just the same way as x and y; and 3 requires, in effect, that x^* and y^* do not differ in any other relevant respect.

Consider again the movie example. The analysis says that both David and Margaret must believe – or, at any rate, be disposed to believe – that there is some set of relevant respects such that, if the two movies in question did not differ in any of those respects, though they did still differ in respect of having realistic special effects (in just the way that they actually do differ), then *The Phantom Menace* would be better than *Star Wars*, all things considered. Putting the point a little more formally, on the proposed analysis, David and Margaret must both believe that there is some set of respects, $r_1, r_2, \ldots, r_n$, such that the following is a valid inference.

(1) *The Phantom Menace* and *The Phantom Menace** do not differ in respect of having realistic special effects;
(2) *Star Wars* and *Star Wars** do not differ in respect of having realistic special effects;
(3) *The Phantom Menace** and *Star Wars** do not differ in respect of any $r_1, r_2, \ldots, r_n$;
therefore,
(4) *The Phantom Menace** is better than *Star Wars**, all things considered.

Thus, the analysis shows how David and Margaret may both be committed to the claim that *The Phantom Menace* is a better film than *Star Wars* in respect of having realistic special effects, even though they disagree as to what the relevant set of respects is: for David it is just "having an interesting story", whereas for Margaret it is just "having convincing acting". In order to see David's reason for accepting the inference, we may substitute (3) with

(3_D) *The Phantom Menace** and *Star Wars** do not differ in respect of having an interesting story.

Whereas, for Margaret, we may substitute it with

(3_M) *The Phantom Menace** and *Star Wars** do not differ in respect of having convincing acting.

Given the substitutions (3_D) and (3_M), David and Margaret, respectively, must be disposed to assent to the validity of the above inference; they must be disposed, that is, to assent to the conclusion (4) under the supposition of (1–3).

4. THE ANALYSIS DEFENDED

So now we have an analysis of in-a-respect betterness. In the sequel, I shall apply the analysis to the case of levelling down. But first we need to address two obvious objections.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

4.1 First Objection

An objector might say the following:

> "It might be that, for a given respect r, there is one set of respects $r_1, r_2, \ldots, r_n$ such that, for any x^* and y^*, x^* is better than y^* all things considered, and yet *another* set of respects $r'_1, r'_2, \ldots, r'_n$, such that, for any x'^* and y'^*, y'^* is better than x'^* all things considered. In such a case, your analysis would entail both that x is better than y in respect of r, *and* that y is better than x in the *very same* respect. Surely that is a contradiction."

I agree: that would indeed be a contradiction. But it does not follow that the analysis should be rejected, since one may equally well avoid the contradiction by rejecting the objector's first claim. Thus, the analysis places a certain rational constraint on one's all-things-considered judgements: if one's all-things-considered judgements imply that there is some $r_1, r_2, \ldots, r_n$, such that, for any x^* and y^*, x^* is better than y^*, all things considered, then they must not, on pain of contradiction, also imply that there is some $r'_1, r'_2, \ldots, r'_n$, such that, for any x'^* and y'^*, y'^* is better than x'^*, all things considered.

It is, of course, perfectly plausible that one's all-things-considered betterness judgements are subject to rational constraints: for example, surely few would deny that if x is judged better than y, all things considered, then y cannot be judged better than x, all things considered. More controversially, many maintain that one's all-things-considered betterness judgements must be transitive. Is the additional constraint proposed here a plausible one? In the absence of a clear example in which the constraint is violated (I have not been able to think of one) it is hard to adjudicate this issue. (One might expect that an example of the kind used to illustrate the next objection – i.e., an example involving some kind of non-seperability or complementarity – would do the trick; but, as we shall see, that is not so.) If it turns out not to be an acceptable constraint, then I believe there are ways in which the analysis might be amended; but, for the time being at least, I think it best to stick with what we have.

4.2 Second Objection

An objector might say the following:

> "Your analysis is inconsistent with the context-dependent nature of some respects of value; and so, to that extent, it is deficient."

In discussing this objection, I think it will be useful to have an example on hand. To that end, consider the case of Daniel, who is contemplating a change in his facial hair. Although presently Daniel is fully bearded, he is entertaining three alternatives to his current appearance: firstly, he might shave his beard off entirely; secondly, he might shave off only the left-hand

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

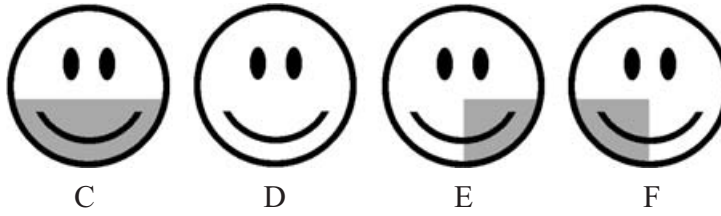


FIGURE 1. Daniel's facial hair.

side of his beard, leaving the rest of it intact; finally, he might shave off only the right-hand side of his beard. (Needless to say, Daniel is peculiarly partial to non-standard facial-hair configurations.) Thus, we may consider four states of affairs: in state *C*, Daniel has a full beard (this is the status quo); in state *D*, Daniel is perfectly clean shaven; in state *E*, Daniel has only half a beard, which is wholly located on the left side of his face; and in state *F*, Daniel again has only half a beard, but in this case it is on the right side of his face (the shaved region in *E* occupies the same facial location as the bearded region in *F*). These four states are represented in Figure 1, where the shaded regions indicate, in the obvious way, which regions of Daniel's face are bearded in the relevant states of affairs.

Now, plausibly, one may wish to make the following two all-things-considered judgements:

(C1) *C* is better than *E*, all things considered.
(C2) *D* is better than *F*, all things considered.

In other words, it is better, all things considered, that Daniel shaves not at all than that he shaves only the right side of his face; and it is also better, all things considered, that he shaves off his entire beard than that he shaves only the left side of his face. All of this is, I think, quite uncontroversial.

Now, suppose we coin a new expression – "right-side beardedness" – to denote a putative respect of value, where the degree of right-side beardedness manifested by a given state of affairs is directly proportional to the extent to which, in that state, the right side of Daniel's face is bearded. For example, in *C*, the right side of Daniel's face is fully bearded, but in *D*, it is not bearded at all; so *C* has more right-side beardedness than *D*. In addition to the above all-things-considered claims, then, some may also wish to make the following two in-a-respect judgements:

(C3) *C* is better than *E*, in respect of right-side beardedness.
(C4) *D* is better than *F*, in respect of right-side beardedness.

Notice that there is a kind of context-dependence involved in the latter two claims. In the context of comparing *C* and *E* – that is, in the context of an otherwise *bearded* face – it is regarded as a *good* thing for the right

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

side of Daniel's face to be bearded; yet in the context of comparing *D* and *F* – that is, in the context of an otherwise *clean shaven* face – it is regarded as a *bad* thing for the right side of Daniel's face to be bearded. Thus, the right half of Daniel's beard appears to change its value from one context to the next; in the former context it has value, but in the latter it has *dis*value.[5]

However, on my analysis, the four claims (C1–C4) may well be inconsistent. Assume, for simplicity, that the only possible "other relevant respect" here is "left-side beardedness" (where this is defined in a way parallel to right-side beardedness). By applying our analysis, we may translate (C3) as follows:

(C3*) For all states of affairs *C** and *E**, if
(1) *C* and *C** do not differ in respect of right-side beardedness;
(2) *E* and *E** do not differ in respect of right-side beardedness;
(3) *C** and *E** do not differ in respect of left-side beardedness;
then *C** is better than *E**, all things considered.

But (C3*) is inconsistent with (C2). To see this, suppose that $C^* = F$ and $E^* = D$. Under this supposition the three conditions in the antecedent of (C3*) are satisfied. Making the relevant substitutions, we get the following:

(1) *C* and *F* do not differ in respect of right-side beardedness (because in both *C* and *F*, the right side of Daniel's face is bearded);
(2) *E* and *D* do not differ in respect of right-side beardedness (because in both *E* and *D*, the right side of Daniel's face is clean shaven);
(3) *F* and *D* do not differ in respect of left-side beardedness (because in both *D* and *F*, the left side Daniel's face is clean shaven).

However, under the same supposition, the consequent of (C3*) is that *F* is better than *D*, all things considered, and this clearly contradicts (C2). By a parallel line of reasoning (which, for want of space, I will not work through here) it may be seen that, on my analysis, (C4) implies that *E* is better than *C*, all things considered; and the latter contradicts (C1).

So it appears that, if one adopts my analysis, one will have to deny the possibility of context-dependent respects of value. Is this a deficiency of the analysis? It is not clear to me that it is. For there is a very common and widely-accepted rival account of the phenomenon illustrated by our beard example, and it is perfectly consistent with my analysis. This alternative account is provided by G. E. Moore's doctrine of "organic unities", which he states as follows:

> [t]he value of a whole must not be assumed to be the same as the sum of the values of its parts. (1903, p. 28)

[5] Jonathan Dancy (2000, pp. 137–41) argues for a context-dependent view of this kind, which he calls "value holism".

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

Although the doctrine is stated in the language of "parts" and "wholes", rather than that of "in-a-respect" and "all-things-considered", I think it is easily enough translated: in the terms we have employed here, it is simply the thesis that there is no such thing as in-a-respect betterness – that is, that there are no relevant true instances of the form "*x* is better than *y*, in respect of *r*". One who holds to this doctrine may judge that the objects of evaluation, within a given domain, are better or worse than each other, all things considered; but he will insist that these all-things-considered judgements resist decomposition into respects of betterness; he will say, that is, that although, as it might be, *P* is better than *Q*, *all things considered*, there are no *respects* in which *P* is better (or for that matter worse) than *Q*.

As we saw, given (C1) and (C2) and the simplifying assumption (i.e., that the only possible respects of value are right-side beardedness and left-side beardedness), my analysis implies that there are no relevant respects of value in the case of Daniel's facial hair. One who adopts my analysis may explain this result, it seems to me, by saying that each of the four states of affairs under comparison constitutes a Moorean organic unity; and, more generally, he may say that any such case – that is, any case in which the only putative respects of value are context-dependent ones – must be a case involving organic unities (though, of course, this more general claim needs further defence – a task for another occasion). On this Moorean approach, it is just a mistake to go looking for respects in which, say, *C* is better or worse than *E*; for one can apprehend which of these two states is better, all things considered, only by comparing them *as wholes*. Breaking the comparison down into smaller parts – that is, breaking the betterness into respects – just will not work; in this case, at least, there is all-things-considered value, and that is *all* there is.

These last remarks are, of course, not so much decisive as suggestive. I have merely gestured at an important connection between my analysis of in-a-respect betterness and Moore's doctrine of organic unities; and there is a great deal more to be said before such a connection is properly established.[6] Fortunately, however, that task may be put off for another occasion. For the issue of context-dependence may safely be set aside here. I might simply say, for example, that mine is an analysis only of *context-independent* respects of value; that is, I might revise (D1) so as to say " . . . *x* is *context-independently* better than *y*, in respect of *r* . . . ".

[6] Indeed, some may take issue with my claiming Moore on my side; they will insist that Moore is better interpreted as holding a context-dependent view. It is worth noting, however, that Jonathan Dancy – one of the chief proponents of the context-dependent style of view – takes his "value holism" to be opposed to Moore's organic unity theory (Dancy 2000, p. 139). Also, Thomas Hurka (1998) distinguishes two interpretations of Moore's doctrine – the "holistic interpretation" and the "conditionality interpretation" (where the latter corresponds to the view I have called context-dependent) – and argues that it is the former, context-independent interpretation that should be attributed to Moore.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

And that would be quite sufficient for present purposes, since my claim is just that given a plausible analysis of being better in a respect, the prioritarian is committed to saying that levelling down may make things context-independently better, in some respect.

5. THE ANALYSIS APPLIED

Returning now to our central example, recall that the egalitarian claims that B, the state of affairs obtaining in Inegalitaria after the bombing, is better, in respect of equality, than A, the state of affairs obtaining in that place prior to the bombing. Applying our analysis, we see that, by this claim, she means that if A and B differed in respect of equality in precisely the way that they actually do, yet did not differ in any other relevant respect, B would be better than A, all things considered. More formally, she means that there is some set of respects $r_1, r_2, \ldots, r_n$ such that the following is a valid inference:

(1) B^* and B do not differ in respect of equality;
(2) A^* and A do not differ in respect of equality;
(3) B^* and A^* do not differ in respect of any $r_1, r_2, \ldots, r_n$;
therefore,
(4) B^* is better than A^*, all things considered.

But what are the other relevant respects? Here egalitarians divide. To keep things simple, though, let us say that ours is an egalitarian who believes that the goodness of states of affairs is determined by some function that weighs equality in the distribution of welfare against total welfare. On his view, there are two relevant respects: equality of welfare (henceforth abbreviated to "equality") and total welfare.[7] Thus, we may substitute (3) with

(3_E) B^* and A^* do not differ in respect of total welfare.

In Table 1 I have assigned to the relevant states of affairs values that, given the characterization of of A and B in Section 2, satisfy (1–3_E).[8]

[7] We might just as well say "average welfare" as "total welfare". For the sake of simplicity, I assume throughout that we are dealing only with "same-people" cases – that is, that all comparisons are between states of affairs sharing a common population. Given this assumption, average welfare and total welfare come to the same thing.

[8] Some may wish to quibble about (1), claiming that B and B^* do, in fact, differ in respect of equality. Although it strikes me as almost obviously true that the values assigned to states in Table 1 are consistent with (1), my argument does not require that this is so. To be sure, the argument assumes that there is some possible assignment of values that satisfies (1–3_E); but that assumption seems a perfectly reasonable one.

It may also be objected that, though this particular case is an "easy" one (since it involves perfect equality), there are other possible cases in which it will be much harder to

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

		States of Affairs			
		A	*B*	*A**	*B**
Classes of People	*x*	10	1	10	6
	y	2	1	2	6

TABLE 1. The bombing of Egalitaria.

Now, our egalitarian does indeed believe that the distribution (6,6) is better than (10,2), all things considered; surely any egalitarian must believe this. Therefore, given the present assignment of values to A^* and B^*, the egalitarian must say that the above inference goes from true premises to a true conclusion. And, while there may be other assignments of values that satisfy (1–3_E), I see no reason at all to think that any of these would require the egalitarian to withdraw her assent to the conclusion (4). The choice of numbers here seems quite arbitrary; if the inference goes through according to one assignment of values satisfying (1–3_E), it should go through for all such assignments. Therefore, according to our analysis, the egalitarian is indeed committed to the claim that *B* is better than *A*, in respect of equality.

Others, however, want to deny the egalitarian's conclusion; these objectors think it absurd to suggest that *B* is better than *A* in *any* respect. We might well ask, then, what grounds there could be for objecting. On our analysis, the objector must show that, given the above assignment of values, B^* is not better than A^*, all things considered; for otherwise, he cannot deny that there is a set of respects, $r_1, r_2, \ldots, r_n$, such that, for any two states of affairs A^* and B^*, (1–3) entail that B^* is better than A^*, all things considered. If the distribution (6, 6) is indeed better than (10, 2), all things considered, then, as we have just seen, there *is* such a set of respects: it is the set that contains only "total welfare".

Suppose, then, that the objector denies that B^* is better that A^*. Such a position is, of course, not unheard of; utilitarians, for example, would say that we should be indifferent between those two states. On this supposition, the objector can consistently deny that *B* is better than *A* in any respect. But there are two important points to note here. Firstly, even if it is the case that B^* is no better than A^*, surely this is not obviously so; it is not plainly *absurd* to suggest that B^* is better than A^*. Yet, according to

determine whether or not two states differ in respect of equality. Many rival (in)equality indices have been proposed; and, of course, these do not always agree. In my view, each of these indices provides a distinct putative respect of value; hence, it may be that, while one state is better than another in respect of, say, Equality_1, it is not better in respect of Equality_2. In general, I assume that in order for any *r* to count as a well-defined respect of value, it must be determinate whether or not, for any two states, those states differ in respect of *r*.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

standard accounts of the levelling down objection, there is supposed to be something obviously mistaken about the suggestion that levelling down might improve things in some respect – for example, *B* might be better than *A* in some respect. So our analysis reveals, at the very least, that the levelling down objection is far less powerful than is commonly supposed; even if it did show that egalitarianism was false, it certainly would not show that view to be absurd. Secondly, and more importantly, if there is disagreement over *A** and *B**, then we need not bother with *A* and *B*; all the levelling down business becomes quite redundant. For, if it were false that *B** was better than *A**, all things considered, then that in itself would be perfectly sufficient grounds for rejecting egalitarianism; hence, there would be no point in considering what the right thing to say about *A* and *B* was.

Thus the advocate of the levelling down objection faces a dilemma. Either *B** is better than *A**, all things considered, or it is not. In the event of the former, *B* is better than *A* in at least one respect and, hence, premise (2) of the levelling down objection is false. In the event of the latter, there is no need for the levelling down objection, since the following, much simpler argument, will do at least as well: it is not the case that, all things considered, *B** is better than *A**; egalitarianism says that, all things considered, *B** *is* better than *A**; therefore, egalitarianism is false. In either event, the levelling down objection is not of much use; for either it is unsound, or it is redundant.

6. PRIORITARIANISM

I now want to draw out a perhaps surprising implication of the argument of the foregoing section: prioritarians, I want to suggest, are no less vulnerable to the levelling down objection than egalitarians.

The central idea of prioritarianism is that we should give priority to benefiting those who are worse off, since benefiting them will do more good. As I like to think of it, the view offers a certain kind of justification for assigning such priority to the worse off, which is nicely brought out in contrast with another, more familiar, theory in the same neighbourhood. Utilitarianism says famously that the value of a state of affairs is the sum of the utilities – or, as I shall say, the levels of welfare – of its occupants. More formally, it says the value of any state of affairs x is given by the function $U(x)$ defined as follows:

$$U(x) = w_1 + w_2 + \ldots + w_n,$$

where n is the number of individuals in x, and w_i is the welfare level of the ith individual. Prioritarianism says something very similar, except it attaches "individual weights" that are positive and inversely related to the individual's level of welfare (Fleurbaey, 2001, p. 4). Rather than

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

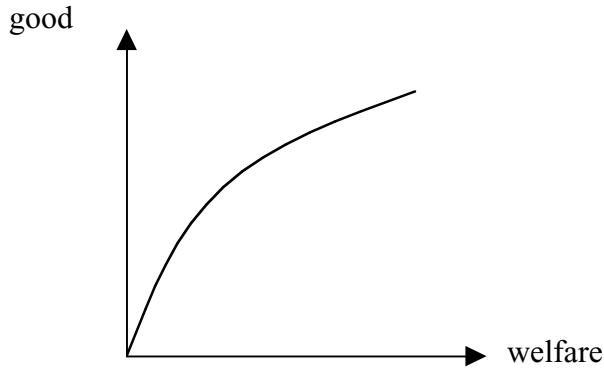


FIGURE 2. Prioritarianism.

summing individual welfare levels, prioritarianism sums something a little different: what one might call the contributions made to overall value by the individuals' having the particular levels of welfare that they do. For the utilitarian, these contributions simply come to the same thing as the welfare levels themselves; but this is not so for the prioritarian. The latter says that an individual's contribution is a diminishing function of his welfare. More formally, prioritarianism says the value of any state of affairs x is given by the function $P(x)$ defined as follows:

$$P(x) = g(w_1) + g(w_2) + \ldots + g(w_n),$$

where g is an increasing *strictly concave* function.[9]

This is well captured graphically. Consider the graph in Figure 2. Along the x-axis is plotted a person's welfare; along the y-axis is plotted the contribution that this person's having the level of welfare that he does makes to the goodness of the relevant states of affairs. Notice that the function plotted is an increasing, strictly concave one. Greater quantities of welfare always map onto greater quantities of goodness, so that as the person's welfare rises, so too does the contribution to goodness. But the rate at which the latter magnitude rises, relative to the former, gradually declines; if we were continuously to increase the person's welfare in equal-sized increments, each successive increment would cause a smaller increase in goodness than the one before. Prioritarianism says that the overall goodness of a state of affairs is given by aggregating individual contributions determined in this way; one might say, therefore, that it asserts the *diminishing marginal value* of welfare.[10]

[9] John Broome (1991, p. 179) gives this definition.
[10] This observation is made by Parfit (2000, p. 105) and Broome (1991, p. 179).

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

Now, suppose that we must choose between benefiting a well-off person and benefiting a badly-off person. Utilitarianism says that the right choice here depends solely on the *size* of the possible benefits: we should choose to benefit whoever stands to receive the greater benefit. The mere fact that one person is worse off than the other is, for the utilitarian, perfectly irrelevant. The prioritarian, however, denies this; on his view, since the value of a person's welfare is marginally diminishing, one does more good by benefiting a badly-off person than one does by benefiting a well-off person (at least, if the possible benefits are equal in size). Hence, one might say helpfully that prioritarianism is utilitarianism modified to give priority to the worse off.

According to the official story, prioritarianism implies that levelling down does no good at all – for example, *B* is not better than *A* in any respect. Parfit (2000) puts it this way:

> Egalitarians face [the levelling down objection] because they believe that inequality is bad in itself. If we accept [prioritarianism], we avoid this objection. We are more concerned for people the worse off these people are. But . . . it makes no difference to our concern whether there are other people who are better off. On this view, when inequality is not bad for people, it simply does not matter. If the better off suffer some misfortune, so that they become as badly off as anyone else, we do not think this in any way a change for the better. (p. 105)

Tungodden (2003) appears to endorse this story:

> Certainly, on the basis of a concern for the absolute circumstances of people, there is nothing to be gained by reducing the well-being of the better off. (p. 25)

However, this is all too quick. To be sure, the prioritarian can explain his judgement that, for instance, *A* is better than *B* without explicitly mentioning respects; he need not say that, while *A* is better in one respect, *B* is better in another, and *A*'s betterness in the first respect outweighs *B*'s betterness in the second. He can avoid altogether such talk of weighing competing respects against one another. But this is not to say that his view does not *imply* such a respect-laden justification. It may be no part of the official prioritarian story that *B* is better than *A* in a respect; but prioritarians might be committed to saying that this is so, nonetheless. (I note, as an aside, that the egalitarian also seems able to justify his judgement that *A* is better than *B* without resorting to respects-talk. Suppose he says that what is ultimately of value is a disjunctive property called "equality-or-total-welfare". If he specified this property in such a way that *A* had more equality-or-total-welfare than *B*, he could offer this as his reason for preferring *A* to *B*, thereby doing away with all talk of respects.)

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

Our analysis of the egalitarian's claim – that *B* is better than *A* in one respect – reveals that the prioritarian, for all his talk to the contrary, is in no position to say any different; for he too must hold that, given the assignment of values in Table 1, *B** is better than *A**, all things considered. The prioritarian is committed to this principle: given a fixed total of welfare, an equal distribution of welfare is always better than an unequal one; in other words, the best distribution of any fixed amount of welfare is an equal distribution. (It is important not to conflate this principle with another: given a fixed total of welfare, the more equally welfare is distributed the better. The prioritarian might not be committed to this latter principle.) Now, since (a) total welfare in *A** equals total welfare in *B**, and (b) welfare is equally distributed in *B**, yet unequally distributed in *A**, the prioritarian must say that *B** is better than *A**, all things considered. But now the argument of the foregoing section kicks in: if the prioritarian believes that *B** is better than *A**, all things considered, then he is committed to the claim that *A* is better than *B*, in respect of equality.

The prioritarian may wish to protest that he simply does not *care* about equality (at any rate, not in a non-derivative way). But that is really beside the point. For mine is not intended to be an analysis of the typical prioritarian's psychology or behaviour: it is not intended merely to capture those features of the world that stir the moral intuitions of self-proclaimed prioritarians; nor is it intended to describe either the conscious mental processes by which they arrive at judgements of moral value, or the kinds of considerations they are apt to invoke in defending those judgements. In short, I am not offering an account of what prioritarians care about. In my opinion, none of the phenomena just listed holds much promise as a basis for analysing the concept of in-a-respect betterness.

I prefer instead to allow prioritarianism – that is, the view itself – to, as it were, speak for itself, where the "speaking" is done via the all-things-considered betterness ordering that the view entails. On my analysis, talk of respects of value is nothing more than a convenient shorthand way in which to describe certain patterns in the all-things-considered betterness ordering. In this sense, my definition of "better in a respect" is reductive: respects-talk is reduced to all-things-considered talk. In principle, I believe we could do away with in-a-respect betterness claims altogether, and still get by; there would be no consequent impoverishment of the information we were capable of conveying. However, referring directly to patterns in the all-things-considered betterness ordering – that is, without the intermediary of respects-talk – would likely be terribly inconvenient.

The reason I like the reductive approach is that it provides a neat account of the explanatory nature of respects-of-value-talk. It is uncontroversial, I take it, that respects of value subserve an explanatory function; such is their "point". When called upon to explain one's judgement that, as it might be, *P* is better than *Q*, all things considered,

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

one will commonly proceed as follows: first, one cites those respects in which one judges *P* to be better than *Q*; second, one cites those respects in which one judges *Q* to be better than *P*; and third, one says that *P*'s betterness in the former set of respects outweighs *Q*'s betterness in the latter (in some simple cases the last two steps may not be required). But how is it that one's appealing to respects of value in this way might help to *explain* one's all-things-considered judgement? My answer is that, at each step in the process just described, one is saying something about the all-things-considered betterness ordering; one is saying that, rather than being completely random, there are certain patterns that may be discerned in it. Thus, by pointing to respects of value, one explains a particular all-things-considered claim – one renders it intelligible – by showing how it fits into a larger pattern of such claims.

By way of illustration, one might say that, on a certain view of causation, causal-talk performs an explanatory role in much the same way. One may explain a particular event – say, my falling off my bike – by pointing out that it was caused by another event – say, my taking the corner too quickly. On the so-called "Humean" theory of causation, this causal explanation works by exposing a certain regularity (or pattern) in events more generally – namely, that falling-off-bike-type events tend to follow taking-corner-too-fast-type events. Thus, my bike accident is explained by showing that, rather than being some random event, it is part of a larger pattern of such events.[11] (Of course, I do not mean to suggest that there is any obvious logical connection between Hume's analysis of causation and my analysis of in-a-respect betterness; as far as I can see, one can perfectly consistently accept either one without accepting the other. I mention Hume's analysis only to illustrate the contribution that patterns or regularities may make to explanation.) I do not take this to be a decisive case for adopting such a reductive approach. Nonetheless, I think it does enough to show that the analysis on offer here warrants further attention; it is not simply a "non-starter".

On a more rhetorical note, one might also mention the fact that distribution-sensitivity is often taken to be a virtue of prioritarianism.[12] It is often for this reason – that is, for the reason that prioritarianism tends to favour equal distributions – that the view is seen as superior to utilitarianism, which is widely taken to be, as one might say, "equality blind". I have always found something dubious in the thought that one can have such distribution-sensitivity while remaining invulnerable to the

[11] Bas C. Van Fraassen (1985, p. 639) attributes to Wesley Salmon the view that " . . . science provides us with understanding of a phenomenon to the extent that it provides an explanation, and provides *explanation* to the extent that it fits the phenomenon in a discernible *pattern* of causal relations" (emphasis added).

[12] See, for example, Nils Holtug (1999, pp. 17–21).

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

levelling down objection. For once you allow that the way in which welfare is distributed in a state of affairs may make some contribution to that state's overall value, it seems you must admit the possibility of levellings down in which, though the total welfare is drastically reduced, the distribution of welfare is improved. To put the point differently, it is hard to see how one could consistently endorse distribution-sensitivity and yet deny that there are any distributional respects in which some states of affairs might be regarded as better than others. In this sense, prioritarianism, as normally understood, just looks too good to be true – like an attempt at having one's cake and eating it too. I take it as a virtue of my analysis, therefore, that it exposes this duplicity in prioritarianism, as it is rendered in the orthodox picture (though I concede that others may not see it in this way).

7. THE ANALYSIS REFINED

In this section I want to slightly refine the analysis proposed earlier. I think the analysis, as it stands, is fine as far as it goes, and is certainly on the right track; but it has what seems to me to be a moderately counter-intuitive implication, which one would do well to avoid.

As we have seen, the way the analysis works is to specify, for any given in-a-respect betterness claim, what one might term a "test class": that is, a class of pairwise comparisons by reference to which one may test the truth of the claim in question. In our central example, we were concerned to determine the meaning of the egalitarian's claim "*B* is better than *A*, in respect of equality". For this purpose, the analysis specifies a test class consisting in all those comparisons between pairs of states of affairs that are such that (a) neither member of the pair contains greater total welfare than the other, and (b) the two members of the pair differ in respect of equality *in precisely the way that A and B do*. The egalitarian's claim is then said to be true if and only if a certain pattern of all-things-considered betterness is present in the test class – that is, if and only if, in every comparison, the state with more equality is better, all things considered, than the state with less.

My worry is that the test class is too narrowly specified, and, in particular, that condition (b) seems too exclusive. Consider, for example, states $G = (3, 8)$ and $H = (2, 9)$: these two differ in respect of equality, but not in the way that *A* and *B* do; notice that, in contrast to *A* and *B*, neither *G* nor *H* contains *perfect* equality, but that, none the less, the distribution of welfare in *G* is more equal than that in *H* (I assume here that equality is not simply a binary, on-off notion – i.e., that it comes in degrees). Condition (b) excludes the comparison between *G* and *H* from the test class. As a consequence, the analysis implies that one may say consistently that *B* is better than *A* in respect of equality, *and* that *G* is *not* better than *H* in the same respect. So the analysis violates the following principle:

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

(D2) For any respect r, if there is some pair of objects of evaluation x and y such that (a) x is more r than y, and (b) x is better than y in respect of r, then, for any two objects of evaluation x^* and y^*, if x^* is more r than y^*, then x^* is better than y^*, in respect of r.

Given (D2), if B is better than A in respect of equality, then G must be better than H in that respect also.

Some may deny (D2). If they are right, then we can rest content with the arguments of the previous sections; those arguments will suffice to prove my central thesis. However, I find (D2) intuitively compelling; hence, I now want to propose a way of broadening the specification of the relevant test class, so as to admit the comparison between G and H, thereby reconciling the analysis with (D2).

To begin, it will be helpful to have a more general definition:

(D3) For any two objects of evaluation x and y, and any respect r, x is better than y in respect of r if and only if x is more (or less) r than y, and r is a respect of value (or disvalue).

The general definition splits the in-a-respect claim "x is better than y in respect of r" into a descriptive component, "x is more (or less) r than y", and an evaluative component, "r is a respect of value (or disvalue)". I take the descriptive claim to be fairly uncontroversial, and so the substance of the analysis consists in fleshing out the evaluative claim – namely, in elaborating what it is to be a respect of either value or disvalue. Here is my definition:

(D4) For any respect r, r is a respect of value (or disvalue) if and only if there exists some finite set of respects R such that:
(1) r is a member of R;
(2) each member of R is independent of each other member;[13]
(3) for any respect r^* that is a member of R, and any two objects of value x^* and y^*, if x^* is more (or less) r^* than y^*, and x^* and y^* do not differ in respect of any member of R besides r^*, then x^* is better than y^*, all things considered.

Together, these two definitions, (D3) and (D4), constitute my refined proposal.

Now, by way of illustration, I shall apply the analysis to a simple case. Suppose that in order to ascertain a certain individual's beliefs regarding the value of bundles of apples and oranges, we have subjected her to the following test. The test subject – call her Mia – has been shown a number of

[13] More precisely, condition 2 is this: $\forall r \in R\ (\exists x, y(x \not\approx^r y \wedge \forall s \in R\ (r \neq s \rightarrow x \approx^s y)))$, where $x \approx^r y$= "x and y do not differ in respect of r". The basic idea is that for each member of R it must be possible to vary that member while holding the other members of R fixed.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

x		y		
Apples	Oranges	Apples	Oranges	Mia's judgement
3	3	2	3	$x > y$
3	3	1	3	$x > y$
2	3	1	3	$x > y$
3	2	2	2	$x > y$
3	2	1	2	$x > y$
2	2	1	2	$x > y$
3	1	2	1	$x > y$
3	1	1	1	$x > y$
2	1	1	1	$x > y$

TABLE 2. Varying apples, holding oranges fixed.

x		y		
Apples	Oranges	Apples	Oranges	Mia's judgement
3	3	3	2	$x > y$
3	3	3	1	$x > y$
3	2	3	1	$x > y$
2	3	2	2	$x > y$
2	3	2	1	$x > y$
2	2	2	1	$x > y$
1	3	1	2	$x > y$
1	3	1	1	$x > y$
1	2	1	1	$x > y$

TABLE 3. Varying oranges, holding apples fixed.

bundles of fruit each of which consists of a number of apples and oranges (and nothing else). She has been asked to perform a series of pairwise comparisons on the bundles of fruit and, in each case, she has reported which of the two bundles under comparison, if either, she believes to be the better of the two, all things considered. The data thereby collected is represented in Tables 2 and 3 (where "$x > y$" means "x is better than y").

This data reveals evidence of such patterns in Mia's dispositions to judge as would commit her to there being certain respects of value relevant in the evaluation of bundles of fruit. Consider first Table 2. There is a clear pattern here. Notice that two features are present in each and every comparison: first, the number of oranges does not vary (e.g., bundles of fruit containing three oranges are compared only with other bundles containing three oranges); and second, the bundle with more apples is judged by Mia better than the bundle with less. Thus, there is evidence of a clear tendency in Mia's judgements to conform to the following claim:

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

(C5) For any two bundles of fruit x and y, if x and y do not differ in respect of the number of oranges, and x has a greater number of apples than y, then x is better than y, all things considered.

Now, consider Table 3. Here we see the same pattern only with the roles reversed: the number of apples is held fixed, and variation in the number of oranges is positively correlated with variation in the overall value of the bundles (at least by Mia's lights). Thus, Mia's reported judgements also conform to this claim:

(C6) For any two bundles of fruit x and y, if x and y do not differ in respect of the number of apples, and x has a greater number of oranges than y, then x is better than y, all things considered.

According to the analysis set out above, therefore, there is evidence that Mia is committed to two respects of value: the number of apples (call this respect "a"), and the number of oranges (call this respect "o"). For the conjunction of (C5) and (C6) entails that there *is* a finite set of respects R such that, if $r = a$, then conditions 1–3 of (D4) are all satisfied; that set of respects is $\mathrm{R} = \{a, o\}$. More precisely, the conjunction of (C5) and (C6) entails that:

(1) a is a member of $\{a, o\}$;
(2) every member of $\{a, o\}$ is independent of every other member;
(3) for any respect r^* that is a member of $\{a, o\}$, and any two objects of value x and y, if x is more r^* than y, and x and y do not differ in respect of any member of $\{a, o\}$ besides r^*, then x is better than y, all things considered.

The third condition may be a little less obvious than the first two. Basically, when r and R are instantiated by a and $\{a, o\}$, respectively, condition 3 says two things: first, when o is held fixed, a is positively correlated with the all-things-considered value of bundles of apples and oranges; and second, when a is held fixed, o is positively correlated with the all-things-considered value of bundles of apples and oranges. Notice that the former claim is just another way of stating (C5), and the latter another way of stating (C6).

If there is a possible R, such that, given $r = a$, conditions 1–3 are all satisfied, then (D4) says that a is a respect of value. Insofar as Mia is committed to (C5) and (C6), therefore, she is committed to a's being a respect of value. As we have seen, those of Mia's judgements that are represented in Tables 2 and 3 conform to (C5) and (C6); hence, there is at least some reason to think that she is so committed. (And, of course, since the same argument goes through if we stipulate that $r = o$, there is just as

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

much reason to think that Mia is committed to *o*'s being a respect of value also.)

Now we may return to the case of prioritarianism. As Marc Fleurbaey points out (2001, p. 7), prioritarianism may be represented by a function that weighs average (or total) welfare against an inequality index that conforms to the Pigou–Dalton principle of transfer. He specifies the form of the relevant function as follows:

$$W = B.(1 - IN)$$

where W represents the social ranking, B is the average level of benefit (or the sum total), and IN the inequality index. It is clear that, given the analysis developed in the foregoing section, this function entails two respects of value (or, more accurately, one respect of value and one of disvalue); for it entails that (a) if we hold fixed IN, then the greater B is, the better the outcome is, and that (b) if we hold fixed B, then the greater IN is, the worse the state is. The set $R = \{B, IN\}$ satisfies the three conditions of (D4);[14] hence, (D4) says that these two respects – B and IN – are both respects of (dis)value, by the prioritarian's lights.

If welfare is perfectly equally distributed – that is, if each person has precisely the same welfare level – then $IN = 0$; otherwise $0 < IN < 1$. In our central example, therefore, IN in state A must be greater than IN in state B; and so B is better than A in this respect, according to the prioritarian. It may be a matter of dispute whether the respect of disvalue captured by IN deserves the label "inequality". But that is not important here. What is important is that, according to a plausible analysis of in-a-respect betterness, the prioritarian is committed to saying that levelling down may make things better in at least one respect.

8. CONCLUSION

As I stated at the outset, the orthodox view is that the levelling down objection shows egalitarianism to be less plausible than prioritarianism. I have argued that this is a mistake; if I am right, the levelling down objection in no way serves to divide egalitarians and prioritarians. Therefore, the holder of the orthodox view, if he is convinced by my arguments, faces a choice: either he comes to see egalitarianism as more plausible than he previously thought, or he comes to see prioritarianism as less plausible than he previously thought. I want to suggest that the former option is best; I think our analysis shows that there is nothing particularly counter-intuitive, or implausible about the claim that levelling down may make things better in one respect.

[14] I take conditions 1 and 2 to be obviously satisfied.

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press

But that leaves us with a puzzle: how might we explain the apparently widespread intuition that levelling down can do no good at all? Recall the story of the bombing of Inegalitaria. What is so objectionable about the Egalitarians' defending their actions by saying "Look on the bright side: at least things are better now in one important respect – there's more equality"? My suggestion is that the Egalitarians' defence is objectionable, not because it is false, but simply because it is inappropriate. True enough, we may grant, things have been made better now in one respect; but that is no cause for celebration.

Where does this leave the debate between prioritarians and egalitarians? It seems to me that if the orthodox picture with which we began has not now been entirely eclipsed, then it has at least had a good deal of doubt cast upon it. If the prioritarian wishes to persist with the levelling down objection, then he must reject my analysis. But that is not all. For, even if the prioritarian can show that he is not, after all, committed to the claim that levelling down may improve things in some respect, thereby getting himself off the hook, he will gain nothing if, in the process, he gets the egalitarian off the hook too. He needs to show not only that he is *not* committed to this claim, but also that the egalitarian *is* committed to it. Therefore, if he rejects my analysis, he must replace it with another that serves this dual purpose. My hunch is that no such analysis will be forthcoming; my challenge to others is to prove me wrong.

REFERENCES

Broome, J. 1991. *Weighing Goods.* Blackwell

Dancy, J. 2000. The particularist's progress. In *Moral Particularism*, pp. 130–56. B. Hooker and M. Little (eds.). Oxford University Press

Fleurbaey, M. 2001. Equality versus priority: how relevant is the distinction? Typescript

Holtug, N. 1999. Utility, priority and possible people. *Utilitas,* 11:16–36

Hurka, T. 1998. Two kinds of organic unity. *The Journal of Ethics,* 2:299–320

Moore, G. E. 1903. *Principia Ethica.* Cambridge University Press

Parfit, D. 2000. Equality or priority? In *The Ideal of Equality*, pp. 81–125. M. Clayton and A. Williams (eds.). St. Martin's Press Inc. and Macmillan Press Ltd.

Temkin, L. 2000. Equality, priority, and the levelling down objection. In *The Ideal of Equality*, pp. 126–61. M. Clayton and A. Williams (eds.). St. Martin's Press Inc. and Macmillan Press Ltd.

Tungodden, B. 2003. The value of equality. *Economics and Philosophy*, 19:1–44

Van Fraassen, B. 1985. Salmon on explanation. *The Journal of Philosophy,* 82:639–51

https://doi.org/10.1017/S0266267103001044 Published online by Cambridge University Press