

Equality versus priority: a useful distinction

Economics and Philosophy, 31 (2015), pp. 219–28

John Broome

Corpus Christi College, Oxford OX1 4JF, UK.

john.broome@philosophy.ox.ac.uk

<http://users.ox.ac.uk/~sfop0060/>

Abstract

Both egalitarianism and prioritarianism give value to equality. Prioritarianism has an additively separable value function whereas egalitarianism does not. I show that in some cases prioritarianism and egalitarianism necessarily have different implications: I describe two alternatives G and H such that egalitarianism necessarily implies G is better than H whereas prioritarianism necessarily implies G and H are equally good. I also raise a doubt about the intelligibility of prioritarianism.

Keywords

Egalitarianism, Prioritarianism, Utilitarianism, Fairness

Acknowledgements

This paper is a response to Fleurbaey (2015). It was commissioned by the World Health Organization, and I am grateful to the WHO for permission to publish it in this symposium. It was presented at a conference organized by the WHO in Thiruvananthapuram. I am grateful to the participants for their helpful comments. Apart from some minor copy-editing, this paper is the original version written in 2000.

My book *Weighing Goods* (Broome 1991) contains a discussion of egalitarianism and prioritarianism.¹ At the time the book was written, prioritarianism was well established amongst economists,² but was only just being discovered by philosophers. This note outlines some arguments from the book, very much abridged, in response to Marc Fleurbaey's 'Equality versus priority'.

1. *The distinction*

Take a fixed population of n people, and imagine various distributions of wellbeing across those people. Each distribution can be described by a vector of the form (w_1, w_2, \dots, w_n) , which lists the wellbeing of each person in turn.³ We wish to compare these distributions together, to determine which is better than which. In fact, we wish to put them into an order according to their goodness, with better distributions ranked higher in the order and worse ones lower. Different ethical views will order the distributions differently. I shall consider only views that are consistent with 'the principle of personal good', as I call it. This is the principle that, if one distribution gives some person more wellbeing than another distribution does, and if it gives no person less wellbeing than the other does, then it is better than the other.

One ethical view is *utilitarianism*. Utilitarianism says that one distribution is better than another if and only if it has a greater total of wellbeing, and two distributions are equally good if and only if they have the same total of wellbeing. This ordering of the distributions

¹Particularly in chapter 9. At the time, the term 'prioritarianism' had not been invented, and 'egalitarianism' had not acquired a meaning distinct from prioritarianism. I used the term 'additively-separable egalitarianism' for prioritarianism.

²For example see Atkinson and Stiglitz (1980: 340).

³Strictly, the vector (w_1, w_2, \dots, w_n) lists each person's wellbeing *apart from egalitarian considerations*. This is a technical correction that it is safe to ignore. Its point is mentioned in note 7.

can be *represented* by a particular *value function*. A value function assigns a value to each distribution. To say it represents the ordering means it assigns a higher value to one distribution than to another if and only if the former is better than the latter, and it assigns two distributions equal value if and only if they are equally good. A function that represents utilitarianism's ordering is the simple total of wellbeing:

$$w_1 + w_2 + \dots + w_n.$$

Utilitarianism gives no value to equality in the distribution of wellbeing. It cares only about the total of wellbeing, not about how wellbeing is spread amongst the people. In this paper, I shall concentrate on views that, unlike utilitarianism, do give some value to equality. More specifically, I shall concentrate on views that have this implication: that one distribution is better than another if it has the same total of wellbeing as the other, but has that wellbeing more equally distributed. I shall take this principle to be more formally expressed by the Pigou–Dalton condition (Fleurbaey 2015).

Any view that assigns value to equality in this sense might be called egalitarian. Daniel Hausman (2015) urges us to adopt this piece of terminology. But there is an important distinction to be drawn amongst views that give value to equality in this sense, and later in this note I shall argue that some of them cannot count as truly egalitarian. Those ones should be called prioritarian and not egalitarian.

How should we specify the distinction between egalitarianism and prioritarianism? Fleurbaey reminds us that the two views differ in the grounds they offer for valuing equality. Prioritarians believe that a distribution of wellbeing should be valued on the basis of each person's wellbeing taken separately, independently of its relationship to other people's. It is the absolute level of a person's wellbeing that matters to prioritarians, and not how her wellbeing compares with other people's. The lower a person's wellbeing, the greater the priority prioritarians assign to improving it. This leads them to value equality indirectly. On the other hand, egalitarians are directly concerned with how each person's wellbeing stands in comparison to other people's. They value equality directly.

Do these different grounds translate into a concrete difference between the two views' implications? Specifically, do they imply different orderings of distributions?

I think they do. I am going to suggest a distinction that I think correctly captures and makes precise the views of the two sides, and attributes to them concretely different conclusions. I take it for granted that both views include the Pigou–Dalton condition and the principle of personal good. I suggest that prioritarianism should be understood as the view that, also, the ordering of the distributions is *additively separable*, and egalitarianism as the view that, also, it is not.

What does this mean, exactly? A complication is that some orderings cannot be represented by a value function. But for brevity in this note, I shall concentrate on the ones that can be. Among those orderings, the additively separable ones are those that can be represented by an additively separable function. Given the Pigou-Dalton condition and the principle of personal good,⁴ an additively separable function is one that has the form:

$$f(w_1) + f(w_2) + \dots + f(w_n). \quad (*)$$

where f is some increasing strictly concave function (a function whose graph slopes upwards but bends downwards).

I think this additively separable formula accurately captures the prioritarian idea that each person's wellbeing should be evaluated independently of other people's wellbeing. Conversely, I think the requirement that the value function is not additively separable captures the egalitarian idea that comparisons between different people's wellbeing matter.

2. *Some differences between the theories*

Marc Fleurbaey adopts the same definition of prioritarianism, but (like Hausman) defines egalitarianism to include prioritarianism. I think he does so because he believes that egalitarianism, as I define it, is not a principled, distinctive view that can be opposed to prioritarianism. But I believe it is a distinctive view, and deserves a distinct name. This note

⁴These conditions ensure there is the same f function for each person, and this function is increasing and strictly concave.

explains why.

True, I define egalitarianism in a negative fashion, as the view that is not prioritarian. But this does not necessarily make it unprincipled or undistinctive. Here is a close parallel. Expected utility theory is the view that the ordering of uncertain prospects is additively separable amongst states of nature. Non-expected-utility theory is the view that it is not. Expected utility theory could formally be treated as a special, extreme case of non-expected-utility theory. But the people who believe in non-expected-utility theory think they have principled reasons for rejecting this special case. Similarly egalitarians think they have principled reasons for rejecting prioritarianism.

A prioritarian ordering of distributions is inevitably different from an egalitarian one, just because one is additively separable and the other is not. However, the difference cannot be effectively displayed in very simple examples that compare together only two distributions. For instance, the examples used in the ‘levelling-down objection’ are ineffective at distinguishing the views. Changing from A to B below is what is called a ‘levelling down’:

$$A = (100, 200)$$

$$B = (100, 100).$$

Prioritarians think A is better than B ; they are opposed to levelling down. So are egalitarians as I have defined them. The principle of personal good implies that A is better than B , and I included the principle of personal good as part of the definition of both egalitarianism and prioritarianism. Levelling down has no tendency to separate prioritarians from egalitarians, as I defined them.

True, some extreme egalitarians reject the principle of personal good. I excluded them from my definition because I see no merit in their view. An egalitarian who accepts the principle of personal good can still be a genuine egalitarian, so levelling down provides no objection to egalitarianism. I agree completely with Fleurbaey about this.

It takes more complicated examples to display the concrete difference between egalitarianism and prioritarianism. Here is one, which is roughly modelled on Maurice Allais’s (1979) famous counterexample to expected utility theory. Compare these four distributions:

$$C = (2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2)$$

$$D = (4, 1, 2, 2, 2, 2, 2, 2, 2, 2, 2)$$

$$E = (2, 2, 1, 1, 1, 1, 1, 1, 1, 1, 1)$$

$$F = (4, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1)$$

Prioritarianism implies that C is better than D if and only if E is better than F . This is easily checked from (*). The reason is that the only difference between C and D is in the wellbeings of the first two people, and there is exactly the same difference between E and F . Additive separability implies that the wellbeing of everyone else is irrelevant to the comparison.

However, C in this example has the merit of perfect equality, which E does not share. An egalitarian might find this a reason to think C better than D that is not also a reason to think E better than F . On the other hand, F is plausibly better than E because of its greater total of wellbeing. So an egalitarian may well think that C is better than D , and F better than E . This view is inconsistent with prioritarianism. It is a principled egalitarian view.

3. A crucial difference

Nevertheless, it is not incumbent on an egalitarian to think C better than D and F better than E . So this example does not definitively separate egalitarians and prioritarians. In *Weighing Goods*, I presented an example that does definitively separate them. Fleurbaey discusses this example in his paper. It is this:

G : Either (2, 2) or (1, 1) with equal probability.

H : Either (2, 1) or (1, 2) with equal probability.

The options G and H are uncertain prospects, so the example introduces a complication that we have not faced till now. I shall assume that expected utility theory tells us the right way to

make a valuation under uncertainty.⁵ According to expected utility theory, the way to value a prospect with uncertain results is first to assign something called a ‘utility’ to each of its possible outcomes, and then take the mathematical expectation of this utility. That expectation – ‘expected utility’ – represents the prospect’s value; one prospect is better than another if and only if it has a higher expected utility.

What does prioritarianism say about the relative goodness of G and H ? Each of these prospects has two possible outcomes, and each outcome is a distribution of wellbeing. Prioritarians value a distribution of wellbeing according to the additively separable formula (*) – in this case $f(w_1) + f(w_2)$. For example, the first outcome in option G has the value $f(2) + f(2)$, and the first outcome in H the value $f(2) + f(1)$. So far as expected utility theory is concerned, the utility of an outcome need not be exactly equal to its value found this way. But in fact prioritarians must identify utility with value, as I shall soon explain. For that reason, a prioritarian must value G and H by the expectation of their value, calculated according to the additively separable formula. For G this gives us

$$\frac{1}{2}\{f(2) + f(2)\} + \frac{1}{2}\{f(1) + f(1)\},$$

and for H

$$\frac{1}{2}\{f(2) + f(1)\} + \frac{1}{2}\{f(1) + f(2)\}.$$

These two expected values are the same. The conclusion is that prioritarians must take G and H to be equally good.

Fleurbaey points out that prioritarians might have a way to escape this conclusion. A prioritarian need not value G and H by taking the expectation of their value according to the additively separable formula. Instead, she might take the expectation of some transform of this value: $g(f(w_1) + f(w_2))$, where $g(\bullet)$ is some increasing function. Then she would calculate the value of G as

⁵This use of expected utility theory is defended in my *Weighing Goods*, chapter 6.

$$\frac{1}{2}g(f(2) + f(2)) + \frac{1}{2}g(f(1) + f(1)),$$

and the value of H as

$$\frac{1}{2}g(f(2) + f(1)) + \frac{1}{2}g(f(1) + f(2)).$$

A suitable transformation $g(\bullet)$ will make G come out better than H .

A prioritarian has potentially two ways to justify transforming the value function. First, the additively separable function (*) only represents the *ordering* of the distributions by their goodness; it does not pretend to measure their goodness more precisely than that. It only gets distributions in the right order: it assigns one distribution a higher value than another if and only if that distribution is better. Any increasing transform of the value function will represent the ordering equally well.

Secondly, even if the value function measures goodness more strictly than this, expected utility theory allows utility to be any increasing transform of value. So expected utility theory would itself permit the transformation. All this Fleurbaey points out.

However, this transformation of value is not really available to a prioritarian. Prioritarianism requires a distribution to be evaluated on the basis of each person's wellbeing, taken separately from other people's. Look at the two individual's separate wellbeings under the options G and H . Each option gives each individual either two units of wellbeing or one unit, with equal probability. So each person has exactly the same prospect of wellbeing under G as she has under H . A prioritarian must therefore take G and H to be equally good.

This argument extends prioritarianism a little. Prioritarianism was defined for distributions of wellbeing, and I am applying it to an uncertain prospect ('ex ante', in Fleurbaey's terms). But the case is special, because each individual has the very same prospect of wellbeing under G as under H . I do not think a prioritarian could justifiably deny the extension for this special case. So prioritarianism directly implies G and H are equally good. We do not need to consider the formulae to tell that.

Any manoeuvre on the part of a prioritarian to give H and G unequal value must therefore be mistaken. This conclusion blocks the transformation of the value function by $g(\bullet)$. Prioritarians must evaluate the prospects G and H by taking the expectation of the additively

separable value function (*), untransformed.⁶

On the other hand, a true egalitarian must think G better than H . When valuing a distribution, an egalitarian thinks we should compare the positions of the different people, to see how they stand relative to one another. In G they stand equal for sure, and in H they stand unequal for sure. So G must be better than H from the point of view of an egalitarian. No one deserves the name of ‘egalitarian’ unless she thinks G better than H .⁷

I conclude that this example strictly separates prioritarrians from egalitarians. Prioritarrians must think G and H equally good, whereas egalitarians must think G better than H .

4 The merits of the theories

⁶Someone who adopts a transformed function fits my definition of a prioritarian in this paper, since her ordering of distributions is additively separable. I count the transformed function as a mistaken form of prioritarianism. *Weighing Goods* gives a different definition, which excludes the transformed function.

⁷A technical problem lurks here. G and H are equally good for each person, taken separately. Yet according to egalitarians, G is better than H . But if we extend the principle of personal good to prospects, it implies that if two prospects are equally good for each person separately, they are equally good. So egalitarianism apparently conflicts with the principle of personal good applied to prospects. However, *Weighing Goods* argues in chapters 8 and 9 that the principle of personal good does indeed apply to prospects.

It reconciles the apparent conflict with egalitarianism by arguing in section 9.4 that egalitarianism should be *individualistic*. Egalitarians should recognize that the good of equality is an individual good – one component of an individual’s wellbeing. An egalitarian should therefore think that the figures shown in my distributions of wellbeing do not record the whole of wellbeing. As note 3 explains, these figures show wellbeing apart from egalitarian considerations. This allows the reconciliation.

On the other hand, a prioritarian does not recognize these egalitarian considerations, and takes the figures to show the whole of wellbeing.

Who are right, prioritarrians or egalitarians? *Weighing Goods* offers in chapter 10 an argument against prioritarianism, though not a conclusive one.⁸ Once more, I can only sketch it here. Its core is this. To give their theory meaning, prioritarrians need a measure of a person's wellbeing that is distinct from the value of her wellbeing. They may not be able to find one.

Take some given quantity of improvement in a person's wellbeing. According to prioritarrians, the *value* of this improvement depends on whom it comes to. By its 'value', I mean the amount it contributes to the overall goodness of a distribution. An improvement has more value if it comes to someone who is initially worse off than if it comes to someone who is initially better off. So prioritarrians make a distinction between the quantity of a change in wellbeing, and the value of the change. They distinguish the value of wellbeing from the quantity of wellbeing.

Compare these two distributions of wellbeing:

$$I = (3, 3)$$

$$J = (2, 4)$$

A prioritarian will think *I* better than *J*. Imagine a change from *J* to *I*. In this change, the first person gains one unit of wellbeing and the second person loses one. But the first person has priority, because she is worse off than the second. A unit change in her wellbeing is more valuable than a unit change in the second person's wellbeing. So in a change from *J* to *I*, the two people's wellbeings change by the same quantity – positive for one and negative for the other – but the changes differ in value.

Prioritarianism presupposes a quantitative scale for the quantity of wellbeing: a *cardinal* scale, to be exact. Philosophers often take it for granted that we have a rough quantitative notion of wellbeing. No doubt that is so, but we need to ask what its source is. To have a cardinal scale, we have to be able to make sense of statements that compare the quantity of changes in people's wellbeing. The example contains one change from 2 to 3 and another from 4 to 3, and these changes are supposed to be the same in quantity: one unit. What exactly does it mean to say they are the same in quantity?

⁸For a contrary argument, see Rabinowicz (2002).

One possible meaning is that they have the same value. But that cannot be a prioritarian's meaning, because a prioritarian thinks these changes differ in value. She has to make sense of the comparison some other way.

She may call on the valuation of uncertain prospects, because those also implicitly compare differences in wellbeing. Compare these two prospects, from the point of view of a single person:

$I' = b$ for sure,

$J' =$ Either a or c , with equal probability.

a , b and c are outcomes of some sort, for examples states of the person's health. Assume a gives our subject more wellbeing than b , and b more than c . Compared with I' , the risky prospect J' offers a possible gain in wellbeing from c to b , but also a possible loss from a to b . If we ask which prospect is better for the person, we are in effect comparing the possible gain with the possible loss. Suppose we conclude that I' and J' are equally good. We could take that to mean that the possible gain and the possible loss are the same in quantity: the difference in wellbeing between c and b is the same as the difference between b and a . We can assign meaning to quantities of wellbeing by generalizing this idea.

A prioritarian might think we give a quantitative meaning to wellbeing this way, through evaluations of uncertain prospects. Provided these evaluations are independent of the value of distributions of wellbeing, she is then free to distinguish the value of wellbeing from the quantity of wellbeing, as she needs to.

However, actually the evaluation of uncertain prospects is not independent of the value of distributions of wellbeing. The two correspond exactly, so that if the quantity of wellbeing gets its meaning through the value of uncertain prospects, the quantity of wellbeing will turn out to be exactly the same as the value of wellbeing. This conclusion can be drawn from a theorem that originates with John Harsanyi (1955).⁹ I am sorry to say I cannot explain it here. It means the prioritarian cannot get her quantities of wellbeing this way.

⁹The theorem must be reinterpreted for the purposes of this argument; see *Weighing Goods*, chapter 10. Rabinowicz (2002) objects to one of the theorem's assumptions.

A prioritarian still needs to separate the value of wellbeing from the quantity of wellbeing. Perhaps she can find a way of doing so. But until she does, her theory is shaky.

Finally, on the other side, a good case can be made for the egalitarian view. It accords well with our natural, intuitive understand of fairness. Fairness comes into play when some good is to be distributed amongst people. It requires people to receive a share of the good that is in proportion to their claim to it. It is not concerned with how much of the good they receive absolutely. Even if each person receives very little, or indeed none at all, each is fairly treated so long as each receives a share in proportion to her claim.

I have been assuming implicitly that people have an equal claim to wellbeing. If they do, then fairness requires that they receive an equal share of wellbeing. In the example of G and H , G for sure leads to a fair result, whereas H for sure leads to an unfair result. So G is undoubtedly better from the point of view of fairness. This is the egalitarian conclusion.

References

- Allais, M. 1979. The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American School. In *Expected Utility Hypothesis and the Allais Paradox*, ed. M. Allais and O. Hagen, 27–145. Dordrecht: Reidel.
- Atkinson, A. B., and J. E. Stiglitz. 1980. *Lectures on Public Economics*. New York and London: McGraw-Hill.
- Broome, J. 1991. *Weighing Goods: Equality, Uncertainty and Time*. Oxford and Cambridge, MA: Blackwell.
- Fleurbaey, M. 2015. Equality versus priority: how relevant is the distinction? *Economics and Philosophy*, this issue.
- Harsanyi, J. C. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy* 63: 309–21.
- Hausman, D. M. 2015. Equality versus priority: a badly misleading distinction. *Economics and Philosophy*, this issue.
- Rabinowicz, W. 2002. Prioritarianism for prospects, *Utilitas* 14: 2–21.