

Prioritarianism

Nils Holtug

1. Introduction

Prioritarianism is the view that, roughly, a benefit has greater moral value the worse the situation of the individual to whom it accrues. This view has received a great deal of attention in recent years.¹ In part, this is due to the fact that prioritarianism has only recently been distinguished from egalitarianism, at least by political philosophers. Thus, it has been suggested that perhaps some political theorists who have thought of themselves as egalitarians are in fact better described as prioritarians (Parfit 1991: 19–22). A second, more important reason is that it has recently been argued that, unlike egalitarianism, prioritarianism is not vulnerable to the so-called Levelling Down Objection. According to this objection, it cannot, in any respect, be better to increase equality when this means lowering the welfare of some and increasing the welfare of none (Parfit 1991: 23). Since egalitarians are committed to the claim that it is in at least one respect better to increase equality, even if by levelling down, they are subject to the Levelling Down Objection. Prioritarians, on the other hand, do not

Thanks to Roger Crisp, Karsten Klint Jensen, Kasper Lippert-Rasmussen, Dennis McKerlie, Larry Temkin, Peter Vallentyne, participants at the Seventh Conference of the International Society for Utilitarian Studies in Lisbon 2003, and participants at the conference on egalitarianism at the University of Copenhagen, 2004, for comments on this chapter and my work on prioritarianism that preceded it.

¹ Arneson (2000; Ch. 11 in this volume); Broome (forthcoming); Brown (2003); Crisp (2003); Fleurbaey (forthcoming); Hausman (forthcoming); Holtug (1999, forthcoming); Hooker (2000: 55–65); Hurley (Ch. 13 in this volume); Jensen (2003); McKerlie (1994, 1997, 2003; Ch. 6 in this volume); Nagel (1991); Parfit (1991); Persson (2001); Raz (1986: 217–44); Scanlon (1982: 123); Scheffler (1982: 31); Temkin (1993a, b, 2000, 2003a, b); Tungodden (2003); Weirich (1983).

value equality but are simply concerned with the fate of the worse off and so see no value in levelling down. Several political philosophers have therefore considered prioritarianism superior to egalitarianism (Arneson 2000; Crisp 2003; Holtug 1999, forthcoming; Parfit 1991).

These claims about prioritarianism and its relative merits have been challenged in various ways. Thus, some have argued that there really is no distinction to be made between egalitarianism and prioritarianism (Fleurbaey forthcoming; Hausman forthcoming). Others have argued that even if there is a distinction to be made, it does not serve to distinguish them with respect to the Levelling Down Objection (Broome forthcoming; Fleurbaey forthcoming). Furthermore, some have raised doubts with respect to the person-affecting basis of prioritarianism (Temkin 1993a: 245–82; 2000: 151–3). Finally, it has been argued that prioritarianism fails to capture some common and important intuitions about relational justice and is in this respect inferior to egalitarianism (Temkin 2003b: 12–19), and fails to appropriately capture the virtue on which it is based, namely that of compassion (Crisp 2003: 755–63).

In the following, I shall first make a suggestion as to how the distinction between egalitarianism and prioritarianism should be drawn. I shall then argue that, according to this distinction, the Levelling Down Objection does serve to distinguish between these two distributive views. Furthermore, I shall argue that while prioritarianism includes an impersonal element, it is quite compatible with the person-affecting basis of the Levelling Down Objection. Finally, I shall argue that prioritarianism may well accommodate our intuitions about relational justice and the virtue of compassion. Thus, my argument will amount to a defence of prioritarianism against a number of objections.

Before I begin my defence, I need to say something about the framework I shall employ. I am concerned with axiology. More precisely, my concern is with principles that order outcomes in terms of welfare. In fact, I shall assume a welfarist axiological doctrine according to which outcome value is a function only of the welfare of individuals.² This means that the versions of prioritarianism and egalitarianism I shall consider are welfarist axiological versions.³

Obviously, the welfarist doctrine I shall assume is highly controversial, but it allows me to focus on some particular aspects of distributive justice and to leave out others. And even if one is not a welfarist, my discussion may be of interest because even non-welfarists may want to attach some importance to the distribution of welfare.

² For similar characterisations of welfarism, see Blackorby *et al.* (1984: 328–32); Ng (1990: 171); Sen (1987: 39). In Holtug (2003b), I point to some problems with this characterisation, but argue that different characterizations may be appropriate in different contexts. And in the present context, this characterization will serve us just fine.

³ Sometimes axiological versions of such principles are called 'teleological'; see Parfit (1991: 3–4).

I need to be a bit more precise about my claim that egalitarianism is vulnerable to the Levelling Down Objection. Strictly speaking, I shall only be concerned with axiological welfare egalitarianism and so I shall not consider how this objection applies to other egalitarian views.⁴ And since axiological welfare egalitarianism is the only kind of egalitarianism I shall be concerned with, I shall simply refer to it as 'egalitarianism'.

2. Egalitarianism

In this section and the next, I consider the issue of how to distinguish prioritarianism from egalitarianism. Consider first egalitarianism. Egalitarians value equality, and equality is a *relation*.⁵ But what does it mean that egalitarians *value* this relation? First, they value it *intrinsically*. That is, egalitarians value equality because they take it to be *good in itself* and not (merely) good because it tends to further some other goal, say, fraternity, political stability, or the general welfare. Thus, a utilitarian may have reasons to prefer some degree of equality, but surely that does not make her an egalitarian. To be an egalitarian, a person must value equality *for its own sake*.

Second, to qualify as an egalitarian, a person must value equality in the sense that she considers more equal outcomes *in one respect better* than less equal outcomes.⁶ However, she need not consider more equal outcomes better all things considered. After all, clearly she may have other concerns, say, for liberty, autonomy, or the general welfare. Combining these two claims, what we get is

Outcome Welfare Egalitarianism. An outcome is in one respect intrinsically better, the more equal a distribution of individual welfare it includes.

Of course, different measures of equality will rank outcomes differently and so there may be genuine disagreements about which outcomes are better than others with respect to this value. Nevertheless, I want to suggest that in order for a measure to be a measure of *equality*, it must imply

⁴ But for the suggestion that the Levelling Down Objection applies to many other forms of egalitarianism as well, see Lippert-Rasmussen (Ch. 4 in this volume).

⁵ Thus, it is common to stress the *relational* nature of egalitarianism. For instance, Parfit writes: 'Egalitarians are concerned with *relativities*: with how each person's level compares with the level of other people' (1991: 23). And Temkin states: 'The egalitarian has no intrinsic concern with how much people have, her concern is with how much people have *relative to others*' (1993a: 200). Others, however, use the term 'egalitarianism' in a wider sense, where it includes at least one view that does not attach intrinsic value to relations between people, namely prioritarianism (see e.g. McKerlie 1996: 277).

⁶ For a similar principle and the suggestion that it is central to egalitarianism, see Tungodden (2003: 6).

The Perfect Equality Claim. An outcome in which everyone has the same share of welfare is more equal than an outcome in which individuals have different shares.

Like the claim that, according to egalitarians, more equal outcomes are in one respect intrinsically better than less equal outcomes, the Perfect Equality Claim seems to me to be a part of our ordinary conception of what (welfare) equality is (see also Vallentyne 2000: 4).

While Outcome Welfare Egalitarianism and the Perfect Equality Claim each seem innocent enough, together they imply

The Egalitarian Relational Claim. An outcome in which everyone has the same share of welfare is in one respect intrinsically better than an outcome in which individuals have different shares.

The Egalitarian Relational Claim, it seems to me, nicely captures an important relational aspect of egalitarianism. A principle that satisfies it will imply that an increase in an individual's welfare from n to $n+1$ intrinsically improves an outcome in one respect if everyone else is at $n+1$, but makes it intrinsically worse in one respect if everyone else is at n . Thus, the value of such an increase depends on the recipient's welfare level relative to that of others. However, as we shall see in Section 4, the Egalitarian Relational Claim renders egalitarianism vulnerable to the Levelling Down Objection.

In the following, then, I shall assume that egalitarians are committed to Outcome Welfare Egalitarianism, the Perfect Equality Claim, and (hence) the Egalitarian Relational Claim. Of course, these three principles amount only to a very limited characterization of egalitarianism. Thus, they only tell us how to rank equal outcomes against unequal ones and only in one respect. Let me address these two limitations separately.

The proposed egalitarian principles do not tell us anything about how to rank various patterns of inequality. However, they are of course compatible with a number of further restrictions. For simplicity, let us momentarily assume that we hold an all things considered version of Outcome Welfare Egalitarianism (such a version is obtained by deleting 'in one respect' from the original formulation). This version may satisfy, for instance, the Pigou–Dalton Principle of transfer. According to this principle, if the sum of welfare remains constant, equality is increased by a transfer of welfare from a better-off person to a worse-off person, as long as their relative positions are not reversed. What this principle captures is the intuition that *equalizing* transfers between two individuals that do not affect the total sum of welfare improve an outcome. However, I do not intend to specify the egalitarian ordering of different patterns of inequality here because, as I shall argue in the next section, the characterization I have already provided suffices to distinguish egalitarianism from prioritarianism.

My characterization is also limited in the sense that it does not tell us anything about how to order outcomes *all things considered*. Consider again Outcome Welfare Egalitarianism. Egalitarians will not want to claim that equality is all that matters and so will not simply hold this principle; they will want to combine the concern for equality with certain other distributive concerns. After all, Outcome Welfare Egalitarianism does not even imply that welfare equality at high welfare levels is better than welfare equality at low welfare levels. Furthermore, if we combine the concern for equality with certain other distributive concerns, there is no guarantee that the resulting principle will imply that an equal distribution is better than an unequal one, all things considered.

Some such principles that combine egalitarian and non-egalitarian concerns will satisfy the Pareto Principle, according to which two outcomes are equally good if they are equally good for everyone, and one is better than the other if it is better for some and worse for none.⁷ Let us call such versions Pareto Outcome Welfare Egalitarianism.⁸ Because it satisfies the Pareto Principle, Pareto Outcome Welfare Egalitarianism implies that an increase in an individual's welfare (holding the welfare of others constant) improves an outcome even if it changes it from a situation of equality to one of inequality. Furthermore, since Pareto Outcome Welfare Egalitarianism is supposed to capture not just the concern for equality but also other relevant moral concerns, I shall take it to be a principle that orders outcomes all things considered.

Note that while Pareto Outcome Welfare Egalitarianism implies that some unequal outcomes are better than some equal outcomes, it nevertheless satisfies the Egalitarian Relational Claim. This is because I have defined Pareto Outcome Welfare Egalitarianism as a principle that combines egalitarian and non-egalitarian concerns such as to satisfy the Pareto Principle, and I have characterised egalitarian concerns in terms of this relational claim. Thus, Pareto Outcome Welfare Egalitarianism implies that an outcome in which everyone has the same share of welfare is *in one respect* intrinsically better than an outcome in which individuals have different shares.

⁷ This version differs from standard versions of the Pareto Principle, since standard versions are formulated in terms of preferences rather than welfare; see e.g. Broome (1991: 152). It corresponds to Broome's principle of personal good; see (1991: 165).

⁸ While many egalitarians seem to accept the Pareto Principle, not all do. Brian Barry reconstructs (and endorses) an argument of Rawls's, according to which equality is unjust when it is Pareto—inferior and, in particular, when there is a Pareto—superior outcome in which the worst off are better off; see Barry (1989: 213–54). Furthermore, Broome defends the Pareto Principle and combines it with egalitarian concerns in Broome (1991: 165–201). Also, Tungodden endorses the Pareto Principle from within an egalitarian framework; Tungodden (2003: 10). Incidentally, both Broome and Tungodden seem to consider the Pareto Principle a plausible restriction on egalitarianism in its own right, rather than an implication of a plausible weighing of egalitarian and non-egalitarian (e.g. utilitarian) concerns. On the other side of this egalitarian divide, McKerlie suggests that there is no particular reason why egalitarians should weight equality and welfare such that the Pareto Principle is always satisfied; McKerlie (1996: 287). Temkin holds a similar view (2003b: 80).

To put it differently, we might say that egalitarians are committed to a particular *reason* for holding the particular all things considered ordering that they hold. Of course, egalitarians do not share all their reasons (then they would hold the same orderings all things considered), but they have at least *one* reason in common. And this reason is itself an axiological ordering, albeit a partial ordering of outcomes in one respect only.⁹

However, Marc Fleurbaey rejects the idea that egalitarianism should be characterized in terms of reasons. He considers what he calls the minimal egalitarian statement that 'unequal distributions have something bad that equal distributions do not have', but claims that this statement has almost no implications for the social ranking (Fleurbaey forthcoming). Obviously, what Fleurbaey calls the minimal egalitarian statement is very similar to the Egalitarian Relational Claim. So what, exactly, is the problem with this statement supposed to be? Fleurbaey stresses that 'it is important to distinguish disagreements about the social ranking from disagreements about the *reasons* supporting the social ranking. Only the former have practical implications and are directly relevant for the policy-maker' (Fleurbaey forthcoming). In other words, what is important for the policy-maker is the social ranking, and the minimal egalitarian statement says almost nothing about that.

Let me make two brief points. First, the Egalitarian Relational Claim is defined in terms of a ranking of outcomes, although a partial ordering in one respect only. Nevertheless, I am sure that Fleurbaey will insist that what is important for the policy-maker is the all things considered ordering. This brings me to my second point. There is of course a sense in which it is true that what has practical implications for individuals is the all things considered ordering. What should be implemented is the all things considered ordering and, in its implementation, the implications of each separate reason that supports this ordering are not felt by anyone.

But how is this supposed to show that the distinctions between distributive principles should be drawn (only) in terms of such orderings? The policy-maker, or anyone else for that matter, may want to know why she should implement a particular ordering, or whether an ordering that is already implemented is justified, and to answer these questions, she will need to invoke reasons. So distinctions in terms of reasons certainly have real importance.

Another implication of accounting for egalitarianism in terms of reasons is that we may find egalitarians and non-egalitarians endorsing identical orderings all things considered. In fact, as we shall see in the next section, we may even find egalitarians

⁹ This ordering is partial because the Egalitarian Relational Claim implies only that *equal* outcomes are in one respect intrinsically better than *unequal* outcomes. Thus, it does not tell us how to rank equal outcomes against other equal outcomes and nor does it tell us how to rank unequal outcomes against other unequal outcomes.

and prioritarians in agreement here (for a similar point, see Fleurbaey forthcoming; Jensen 2003: 101–3; Tungodden 2003: 30–1). The difference, of course, will consist in their reasons for endorsing a particular ordering.

But note that this talk of different reasons for accepting a particular ordering all things considered is in no way mysterious. For instance, several reasons have been given for accepting Rawls's difference principle. One of these reasons appeals to Rawls's contract argument, whereas another appeals only to the 'intuitive' case for giving priority to the worst off in cases where they are not themselves responsible for so being (Barry 1989: 213–14; Kymlicka 1990: 55). So do two people who both hold the difference principle, but each for their own reason, hold the same view? In a sense yes, and in another sense no. In order to bring out the full extent of their commitments, we shall have to refer to their reasons for holding the difference principle. Likewise, in order to describe the difference between egalitarians and other theorists, we shall have to refer to *their* reasons for ordering outcomes in the manner they do.¹⁰

On this account of egalitarianism, then, it may not always be possible to determine whether an ordering that a particular person holds is egalitarian or not. Furthermore, turning from the person who holds this ordering to the ordering itself, there may be no determinate answer to the question of whether it is egalitarian or not. This may seem to be a rather impractical implication.

However, first, there is nothing that this characterisation of egalitarianism prevents us from saying. Even if we cannot always determine whether a particular ordering is egalitarian, we can explain exactly why this is so and what the implications of this ordering are. We can, for instance, say that it is *compatible* with egalitarianism.

Second, as I have stressed, it seems to me that to be worthy of the predicate 'egalitarian', a principle must involve a commitment to Outcome Welfare Egalitarianism, the Perfect Equality Claim, and so the Egalitarian Relational Claim. Therefore, in order to be justified in calling a certain distributive view egalitarian, we must understand it as including a particular (egalitarian) reason for accepting it.

Finally, as I have already mentioned, the feature I have called the Egalitarian Relational Claim is what invites the Levelling Down Objection, and so drawing the distinction between egalitarianism and prioritarianism in terms of it ensures that this distinction has real theoretical interest.

My characterisation of egalitarianism, then, is limited both in that it does not enable us to order various patterns of inequality and in that it does not enable us to order

¹⁰ Incidentally, Fleurbaey seems to acknowledge that egalitarianism and prioritarianism rely on different normative reasons even if these reasons do not translate into any difference in the all things considered ordering (forthcoming). Broome also acknowledges that these two principles rely on different reasons, but claims that this difference *does* translate into a difference in the ordering (forthcoming).

outcomes all things considered. Nevertheless, as we shall now see, it does suffice to distinguish egalitarianism from prioritarianism.

3. Prioritarianism

According to prioritarians, roughly, a benefit morally matters more the worse off the individual to whom it accrues. How does this view differ from egalitarianism? Egalitarians value equality, and equality is a relation. Prioritarians, on the other hand, do not value a relation. Rather, what they are concerned with is *absolute* levels of individual welfare (Parfit 1991: 22–4). A benefit that falls at a particular level of welfare has the same moral value no matter what levels other individuals are at. And the lower this particular level, the greater the value of the benefit.

This description of the value commitments of the prioritarian needs to be made more precise. Like the egalitarian, we should characterise the prioritarian in terms of her commitments with respect to intrinsic value. We should require that she favours giving priority to the worse off not (only) because doing so will tend to further some other goal of hers, say, fraternity or political stability, but (at least in part) because she holds that, everything else being equal, an outcome in which a benefit falls at a lower level is *intrinsically* better than an outcome in which an equal benefit falls at a higher level.

I suggest that the best account of these value commitments is that the prioritarian ascribes intrinsic value to compound states of affairs, each consisting of the state that a benefit of a certain size befalls an individual and the state that the individual is at a particular welfare level, where this value increases when the size of the benefit increases but decreases when the level of welfare increases. Thus, the (compound) state that a benefit befalls an individual at a lower level is intrinsically better than the (compound) state that an equal benefit befalls an individual at a higher level.

Furthermore, like the egalitarian, a prioritarian may be a pluralist. That is, she may have concerns other than priority. But, *qua* prioritarian, she will hold that an outcome in which a benefit falls at a lower level is *in one respect* better than an outcome in which an equal benefit falls at a higher level. Combining these claims (and introducing an additive function) we get:

Outcome Welfare Prioritarianism. An outcome is in one respect intrinsically better, the larger a sum of weighted individual benefits it contains, where benefits are weighted such that they gain a greater value, the worse off the individual to whom they accrue.

I have suggested that egalitarianism satisfies the Egalitarian Relational Claim. Outcome Welfare Prioritarianism, on the other hand, does not let the value of a benefit to an individual depend on the welfare levels of others. The moral value of a further unit of

welfare to any individual depends only on the welfare level of that individual and the weight function. Therefore, prioritarianism does not satisfy the Egalitarian Relational Claim.

Furthermore, whereas Outcome Welfare Egalitarianism cannot reasonably be claimed to exhaust our axiological concerns, Outcome Welfare Prioritarianism may be claimed to do just that. For instance, unlike Outcome Welfare Egalitarianism, it favours equality at higher levels over equality at lower levels. After all, additional benefits increase the value of an outcome whatever the level at which they fall. So from now on and for simplicity, unless otherwise indicated, I shall simply assume that prioritarianism implies an all things considered ordering of outcomes equivalent to the ordering with respect to priority generated by Outcome Welfare Prioritarianism (although, strictly speaking, prioritarians are only committed to the latter principle *qua* prioritarians). That is, from now on I shall take ‘prioritarianism’ to imply

Overall Outcome Welfare Prioritarianism. An outcome is intrinsically better, the larger a sum of weighted individual benefits it contains, where benefits are weighted such that they gain a greater value, the worse off the individual to whom they accrue.

On this assumption, we can bring out the structure of prioritarianism and so the structural difference between egalitarianism and prioritarianism to which I referred above in greater detail. As John Broome has pointed out, prioritarianism can be formally represented by an additively separable function of the following form:

$$G = f(w_1) + f(w_2) + \dots + f(w_n),$$

where f is an increasing and strictly concave function of individual welfare, w_i (Broome 1991: 179). Thus, as I pointed out above, the contribution each individual makes to the value of an outcome depends only on her own welfare, not on the welfare of others.

Importantly, it is here assumed that individuals have identical weight functions in the sense that the moral value of a further unit of welfare is the same for *all* individuals who are at identical welfare levels. In a sense, the claim that individuals have identical weight functions can be said to express an ideal of the moral equality of individuals. A benefit to an individual at a given welfare level is exactly as important as a similar benefit to any other individual at the same level.¹¹ But of course, the claim that individuals are moral equals does not commit the prioritarian to *egalitarianism*, which encompasses the view that it is in one respect intrinsically better if individuals have more equal shares of welfare.

Furthermore, according to prioritarianism, moral outcome value is a strictly concave function of welfare, as illustrated in Figure 5.1. Thus, benefits *gradually* decrease in moral value, the higher the level at which they fall.

¹¹ More technically, the equal weight assumption is captured by the requirement of anonymity.

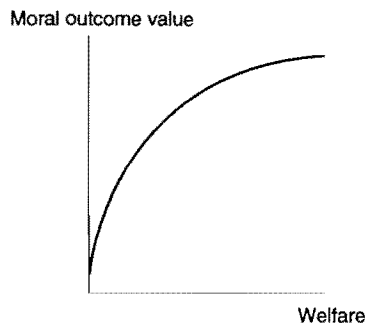


FIG. 5.1

Also, prioritarianism is an *aggregative* (indeed additive) principle, and I shall assume that it implies that for any finite sum of benefits that fall at a lower level, it can be outweighed by a sufficiently large sum of benefits that fall at a—indeed any—higher level. In this respect, it differs from a principle that would give absolute priority to the very worst off.¹²

Finally, note that prioritarianism satisfies the Pareto Principle. Since all increases in individual welfare, holding others constant, increase outcome value, it satisfies the Pareto Principle. And since it responds to transfers of a fixed sum of welfare from a better-off individual to a worse-off individual (that does not reverse their relative positions) by giving greater weight to this sum, there is a sense in which it satisfies the Pigou–Dalton Principle (although it does not assign value to equality).

Let us now return to the distinction between egalitarianism and prioritarianism. Note that my characterisation of the difference between these two principles is in terms of *reasons*.¹³ Whatever particular all things considered ordering an egalitarian accepts, she will do so at least in part because she accepts the Egalitarian Relational Claim. Prioritarianism, on the other hand, does not involve a commitment to the Egalitarian Relational Claim. Of course, as I pointed out above, some may want to combine prioritarian and other concerns, and so the acceptance of Outcome Welfare Prioritarianism is *compatible* with a commitment to this claim. But it is not something to which one is committed in virtue of one's *prioritarian* concerns.

Note also that, as I pointed out in the last section, this way of construing the difference between egalitarianism and prioritarianism does not rule out that egalitarians may in fact endorse a prioritarian ordering of outcomes all things considered. But since

¹² Thus, neither maximin nor leximin are versions of prioritarianism.

¹³ For the suggestion that the difference is to be explained in terms of reasons or justifications, see also Tungodden (2003: 24).

only egalitarians will endorse this particular ordering in part because they endorse the Egalitarian Relational Claim, egalitarianism and prioritarianism are *distinct* views.

Nevertheless, Fleurbaey denies that egalitarianism and prioritarianism are fundamentally different views. Thus, he argues that not only egalitarians, but also prioritarians, attach value to equality (Fleurbaey forthcoming). More precisely, assuming that the resulting ranking of outcomes is continuous, a prioritarian view can be represented by a function of the form

$$W = T(1 - IN),$$

where T is the total sum of welfare, and IN is an inequality index that satisfies the Pigou–Dalton Principle. Fleurbaey provides the following proof. We choose W as a function that represents the prioritarian ordering of outcomes and implies that whenever an outcome holds an equal distribution of welfare, W equals T . According to Fleurbaey, it can be shown that such a function exists. We then adopt the following inequality index:

$$IN = 1 - (W/T).$$

This is held to be a reasonable inequality index because, first, $IN = 0$ when the distribution is equal. And second, IN satisfies the Pigou–Dalton Principle. After all, a prioritarian view satisfies this principle and so any Pigou–Dalton transfer increases W (while T is held constant) and so decreases IN . The point, of course, is that by reversing IN , we get $W = T(1 - IN)$, which is what had to be proved.

In other words, just as Pareto Outcome Welfare Egalitarianism can be divided into parts some of which are egalitarian and others are not, so can prioritarianism. However, the egalitarian value to which Fleurbaey thinks the prioritarian is committed is instrumental, not intrinsic (Fleurbaey forthcoming). More precisely, Fleurbaey thinks the prioritarian is committed to the instrumental value of equality in a logical sense of 'instrumental value'. Prioritarians hold equality to be instrumentally good in that an equal distribution of welfare is held to be intrinsically better than an unequal distribution of the same sum. This logical sense of 'instrumental value' should be distinguished from its more ordinary causal sense, where an entity has instrumental value in so far as it causally produces some (intrinsic) value.

In fact, prioritarians hold that equality has instrumental value not only in the logical sense, but also in the causal sense. Equality has instrumental value in the causal sense because at least certain ways of increasing equality—giving to the worse off rather than the better off—will promote intrinsic prioritarian value.

But even if prioritarians are committed to the instrumental value of equality in both these senses, this does not make them *egalitarians*. Egalitarians hold that equality has *intrinsic* value. And as we shall see, this implies (amongst other things) that egalitarians

consider levelling down to be in one respect intrinsically good. Prioritarianism has no such implication.

Nevertheless, even if prioritarrians are not *committed* to egalitarianism, some may well be egalitarians. I have already conceded this much in the case of *pluralist* prioritarrians, who combine Outcome Welfare Prioritarianism and egalitarian concerns. But the point here is that even proponents of Overall Outcome Welfare Prioritarianism may in fact be egalitarians. This much seems to follow from the availability of Fleurbaey's inequality index, IN. According to IN, inequality reduces to zero when individuals have the same shares, in accordance with the Perfect Equality Claim. And so if an overall outcome welfare prioritarian holds Outcome Welfare Egalitarianism with respect to IN, she will indeed be an egalitarian. But, of course, she *need not* hold Outcome Welfare Egalitarianism with respect to IN. So once again, we reach the verdict that prioritarrians are not committed to egalitarianism *qua* prioritarrians.

While some prioritarrians may be egalitarians, unless otherwise indicated, I shall assume that prioritarrians hold their prioritarian views simply because they take benefits to have greater intrinsic value at lower levels and so not partly because they consider more equal outcomes in one respect intrinsically better than less equal outcomes.

4. The Levelling Down Objection

Outcome Welfare Egalitarianism and the Perfect Equality Claim imply that it is in one respect intrinsically better if inequality is eliminated, even if it does not involve making the worse off better off, but only involves making the better off worse off. After all, together, these two principles imply the Egalitarian Relational Claim, according to which an outcome in which everyone has the same share of welfare is in one respect intrinsically better than an outcome in which people have unequal shares. (For the claim that egalitarianism does in fact *not* have this implication, see Christiano, Chapter 2 in this volume, Sections 20–2.)

Of course, egalitarians need not claim that it is better *all things considered* to increase equality by harming some and benefiting none. Obviously, Pareto Outcome Welfare Egalitarianism implies that this would in fact be a change for the worse. Nevertheless, since Pareto Outcome Welfare Egalitarianism incorporates the three egalitarian principles mentioned above, it implies that an elimination of inequality that harms some and benefits none is *in one respect* intrinsically better. It is intrinsically better regarding equality. But how could this in any respect be better? It would benefit no one, not even the worse off. This constitutes the gist of the Levelling Down Objection (Parfit 1991: 17).




FIG. 5.2 Levelling down.

To illustrate the objection, consider the outcomes in Figure 5.2 (where the width of the columns represents the number of individuals in a group, the height of the columns represents their welfare, and the horizontal line connecting the columns represents the level where life ceases to be worth living). For egalitarians, B must in one respect be intrinsically better than A, namely regarding equality, although it is better for no one. Furthermore, C must in one respect be intrinsically better than A, although everyone suffers terribly in C. These judgements illustrate the force of the Levelling Down Objection.

Prioritarianism, on the other hand, is not vulnerable to the Levelling Down Objection. The feature of egalitarianism that invites this objection is the Egalitarian Relational Claim and, as I have argued, prioritarianism does not satisfy this claim. Prioritarrians value increases in welfare, especially to the worse off, and when we level down there are no increases.¹⁴ Therefore, the Levelling Down Objection serves to distinguish egalitarianism and prioritarianism.

Unsurprisingly, Fleurbaey disagrees. Consider again the prioritarian function *W*. Fleurbaey argues that when a distributional change implies that *W* is reduced in spite of a decrease in IN, the prioritarian must hold that the change worsens things *in spite* of the fact that something good happens on the IN side. Thus, Fleurbaey concludes that 'insofar as prioritarrians give instrumental value to equality, as shown above, they should also be subject to a *similar kind of criticism* [the Levelling Down Objection]' (Fleurbaey forthcoming; my emphasis).

¹⁴ Remember that I am assuming that prioritarrians do not hold this view (in part) because they are egalitarians. If they are *egalitarian* prioritarrians, they will of course be vulnerable to the Levelling Down Objection.

Presumably, the sort of case Fleurbaey has in mind here is a case of levelling down. So let us assume that the welfare level of the better off is reduced to the level of the worst off, resulting in decreases in W , T , and IN . T is reduced because there is a decrease in the sum of welfare and W is reduced because there is a decrease in the sum of *weighted* welfare. And since there is now perfect equality, IN reduces to zero. In other words, levelling down improves an outcome in one respect, namely regarding IN .

However, the *mere fact* that prioritarianism can be split into separate components such as to give rise to an extensionally equivalent ordering does not show that prioritarianism, or anyone who holds it, attaches any relevant kind of value to these separate components. And so the fact that IN decreases in the case described above does not imply that prioritarians are committed to the claim that there is, in any interesting sense, something good about levelling down.

It may be helpful to consider the following analogy. Total utilitarianism can be represented by:

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

Furthermore, U is equivalent to the following function of average welfare:

$$V = n \times AVE,$$

where

$$AVE = (w_1 + w_2 + \dots + w_n)/n.$$

Does this mean that total utilitarians are committed to the claim that an outcome that has a higher average welfare than another is in one respect better? Obviously not. Consider the following two outcomes (where the numbers represent individual welfare levels):

D: (-10, *)

E: (-10, -9).

In D, there is only one individual and he has a horrible life (the asterisk represents the non-existence of a second individual). In fact, he has a negative welfare level. In E, there is also another individual, and she has a life just slightly less horrible, but still clearly on the negative side.¹⁵ Importantly, the move from D to E decreases V in spite of the fact that AVE increases (people are, on average, slightly better off in E). So on Fleurbaey's line of reasoning, total utilitarians must admit that while V decreases, something good is happening on the AVE side.

But obviously, total utilitarians do not consider E to be in any respect better than D. They will hold that E is not intrinsically better in any respect. Furthermore, in

such cases, AVE does not have instrumental value in the causal sense. While certain ways of increasing AVE will also increase total welfare, the move from D to E does not do that. What about instrumental value in the logical sense, then? Admittedly, total utilitarians will take AVE to have instrumental value in the sense that, holding n constant, an increase in AVE increases the value of V . But this just means that if we held the number of individuals in D constant, any increase in AVE would intrinsically improve the outcome. And surely, the mere fact that there are such *other* outcomes that are intrinsically better than D does not imply that, in any interesting sense at least, E is better than D.

Likewise, prioritarians do not consider B to be in any respect better than A in Figure 5.2 (where B results from levelling down A). As we have seen, they will hold that B is not intrinsically better in any respect. Furthermore, in levelling down cases, decreases in IN do not even have instrumental value in the causal sense. While certain ways of decreasing IN will also increase prioritarian intrinsic value, the move from A to B does not do that. Of course, prioritarians will take decreases in IN to have instrumental value in the (logical) sense that, holding T constant, such a decrease increases the value of W . But this just means that if we held T constant, any decrease in IN would intrinsically improve the outcome. And again, the mere fact that there are such *other* outcomes that are intrinsically better than A does not imply that, in any interesting sense, B is better than A. A prioritarian will happily admit that there are such better outcomes, but surely this does not weaken her position. In any case, it is a far cry from claiming that B is in one respect intrinsically better than A, which is what the egalitarian claims. And so the 'similar criticism' Fleurbaey believes prioritarianism to be vulnerable to is really not similar at all.

Let us therefore now assess the Levelling Down Objection (only) in relation to egalitarianism. I believe that there are two distinct concerns from which this objection derives its force. The first is a somewhat general moral concern. There is a strong tendency to think that the value of outcomes must be tied to value for individuals. Roughly, an outcome cannot be better (worse) than another in any respect unless there is someone *for whom* it is better (worse). A problem that levelling down poses for the egalitarian, then, is that levelling down renders no one better off and yet it may involve an improvement in equality.

Therefore, part of the intuitive force of the Levelling Down Objection derives from the fact that the link between equality and improvements in welfare is purely contingent. Equality can come about in different ways, and some of these ways do not involve benefits to anyone. I shall discuss the basis of this intuition and the precise relation between outcome value and benefits it expresses in the next section.

The other concern from which the Levelling Down Objection derives its force is this. Cases of levelling down suggest that the concerns that motivate many of us to be egalitarians may not really be captured by egalitarianism. Part of what motivates our

¹⁵ This case is inspired by a very similar case described by Parfit (1984: 393).

concern for equality may be a concern for how those who are worse off—or worse off than others—fare. Our concern for equality, then, may in one way or another reflect a concern for the worse off. But we may wonder whether egalitarianism adequately captures this concern. Consider, for instance, a move from A to C in Figure 5.2 (where C is perfectly equal but everyone has a life worth not living). Equality is supposed to capture our concern for the worse off but, with respect to equality, C is better than A. The problem, of course, is that the move from A to C does not seem in any way to respond to our concern for the worse off. In fact, the worse off in A are much worse off in C.

Prioritarianism, on the other hand, nicely accommodates this second concern. In fact, it is simply designed to express a concern for the worse off. And while it does not assess the value of a change in the situation of a worse-off individual in terms of how her welfare level compares with that of others, it will tend to favour benefits to those who are worse off than others. If we can confer a benefit of a fixed size on one of two individuals and one of them is worse off than the other, then, everything else being equal, it is better to confer it on the one who is worse off.

5. The Strong Wide Person-Affecting Principle

I have suggested that one of the reasons why the Levelling Down Objection seems so persuasive is that we have a strong tendency to think that the value of outcomes must be tied to value for individuals. Roughly, an outcome cannot be better (worse) than another in any respect unless there is someone for whom it is better (worse). And since levelling down benefits no one, it cannot truly in any respect make an outcome better.

Outcome values, then, do not enter the scene out of nowhere; they crucially depend on how individuals are affected for better or worse. And the problem with egalitarianism is that the link between equality and benefits is purely contingent. However, in order for the idea that outcome values are tied to values for individuals to support the Levelling Down Objection, clearly this idea must be spelled out in greater detail. It is fine to say that outcome values depend *crucially* on values for individuals, but unless we are able to specify the nature of the relation with greater precision, we will not be able to tell just how supportive of the Levelling Down Objection this idea is.

Clearly, this is a point the egalitarian may want to press. He may argue that the sort of relation between outcome values and benefits we must assume does not stand up to closer scrutiny. In fact, this is exactly the strategy employed by Larry Temkin. Temkin suggests that the Levelling Down Objection presupposes a principle he calls the 'Slogan'. According to the Slogan: 'One situation *cannot* be worse (or better) than

another *in any respect* if there is *no one* for whom it is worse (or better) *in any respect*' (Temkin 1993a: 256). Or, to put it slightly differently, the objection presupposes

The Strong Narrow Person-Affecting Principle. An outcome, O_1 , cannot *in any respect* be better (worse) than another outcome, O_2 , if there is no one for whom O_1 is *in any respect* better (worse) than O_2 .

Obviously, this principle implies that there is no value to be found in levelling down. Since, in Figure 5.2, B is not better than A for anyone in any respect, it cannot be better in any respect.¹⁶ So, equality cannot be a value that renders an outcome better than another in even one respect.

However, Temkin believes that the Strong Narrow Person-Affecting Principle is implausible. One of his reasons is that he takes it to be incompatible with various moral ideals such as autonomy, freedom, and proportional justice (1993a: 258–77). But in the present context of axiological welfarism, we need not worry about this. After all, none of these ideals is a function merely of welfare.¹⁷ Another of his reasons for considering the Strong Narrow Person-Affecting Principle implausible is, however, highly relevant. This principle gets us into trouble when applied to the Non-Identity Problem. Consider Figure 5.3. p , we may suppose, is a group of people who presently exist. The people in this group may adopt the 'live for today policy', which involves having their children immediately and depleting natural resources for current uses.




FIG. 5.3 The Non-Identity Problem.

¹⁶ Strictly speaking, the people who are worse off in A might be in some respect better off in B than in A but not better off all things considered. For simplicity, however, let us assume that whenever a group is not better off in one outcome than in another, all things considered, then it is not in any respect better off.

¹⁷ In Holtug (2003a), I consider this reason of Temkin's in greater detail. And Temkin responds in Temkin (2003c).

In that case, F would occur: they would enjoy a great many benefits, but their children would enjoy far less. Or they may adopt the 'take care of tomorrow policy', which involves waiting a few years before they have children, thus conserving some of the resources. The children they would then have are numerically different from those they would have in F. Furthermore, with the 'take care of tomorrow policy', these presently existing people would be slightly worse off, but their children would be much better off than their children would be in F. In other words, G would result (for this particular version of the Non-Identity Problem, see Temkin 1993a: 255–6).

Clearly, F is worse than G. However, Temkin argues that according to the Strong Narrow Person-Affecting Principle it cannot be.¹⁸ According to this principle, in order for F to be worse in even one respect, there must be someone for whom it is worse in at least one respect. And, according to Temkin, there is none. Clearly, F is not worse for *p*. Nor is it worse for *q*. This group has lives worth living and would not exist if G came about. Finally, Temkin argues that F is not worse for *r*, since if F obtains, *r* does not exist and so cannot be worse off. As he puts it: 'one cannot harm or act against the interests of someone who will never exist and, more particularly, one does not harm someone by failing to conceive her' (Temkin 1993a: 255).

Nevertheless, since it seems fairly clear that F is worse than G, we should reject the Strong Narrow Person-Affecting Principle. And then this principle cannot support the Levelling Down Objection. Thus far, I agree with Temkin. However, it seems to me that the proper response to the Non-Identity Problem is not to abandon a person-affecting account of the value of outcomes, but to opt for a *wide* person-affecting account. Thus, I propose

The Strong Wide Person-Affecting Principle. An outcome, O_1 , cannot in *any respect* be better (worse) than another outcome, O_2 , if there is no one for whom, were O_1 to obtain, O_1 would be in *any respect* better (worse) than O_2 and no one for whom, were O_2 to obtain, O_2 would be in *any respect* worse (better) than O_1 .¹⁹

This principle allows us to claim that F is worse than G if G would be better for some, were it to obtain. And it may be argued that, if G were to obtain, then this would be better for *r*. This does not contradict Temkin's claim that 'one cannot harm or act against the interests of someone who will never exist'. Rather, since *r* exists in G, these individuals can have properties there, including the property of being

¹⁸ Strictly speaking, what Temkin says is this: 'Most believe that the "take care of tomorrow policy" should be adopted. But this is incompatible with [the Narrow Person-Affecting Principle] . . .' (Temkin 1993a: 255). However, this principle is in fact quite compatible with the claim that the 'take care of tomorrow policy' should be adopted. After all, there is more to morality than axiology. Thus, perhaps the 'live for today policy' would be wrong, even if it would be worse for no one.

¹⁹ See also Holtug (1998: 17; 2003a: 8). My wording of a principle similar to this in Holtug (1998) was unfortunate, as Glen Newey has pointed out.

better off than in F. More precisely, if G obtains, the triadic relation—G is better for *r* than F—may also obtain. Obviously, if G obtains, then the first two relata obtain and while the third relatum does not obtain, the non-obtaining of an outcome does not prevent it from figuring in a betterness relation (if it did, then the relation 'the outcome that the allies win the war is better than the outcome that the Nazis win the war' could not obtain either).²⁰ Therefore, G can be better for *r* than F, and so the Strong Wide Person-Affecting Principle is compatible with the claim that F is worse than G.²¹

Furthermore, the Strong Wide Person-Affecting Principle implies that it cannot be in any respect better to level down. Consider Figure 5.2 once more. B cannot in any respect be better than A since, were B to obtain, there would be no one for whom, in any respect, this outcome would be better than A and, were A to obtain, there would be no one for whom this outcome, in any respect, would be worse than B.

Since, then, the Strong Wide Person-Affecting Principle implies that there is no value to be found in levelling down, Temkin's claim that the Levelling Down Objection presupposes the Strong Narrow Person-Affecting Principle cannot be correct. Since, furthermore, the Strong Wide Person-Affecting Principle does not get us into trouble with the Non-Identity Problem, it seems to be just the principle needed to meet Temkin's objection (see also Holtug 1998, 2003a). Incidentally, Temkin has recently acknowledged both that the Strong Wide Person-Affecting Principle can be invoked to support the Levelling Down Objection and that it handles the Non-Identity Problem (Temkin 2003b: 74).

Furthermore, prioritarianism does not contradict the Strong Wide Person-Affecting Principle. In order for prioritarianism to imply that an outcome is intrinsically better than another, in any respect or all things considered, it would have to be better for someone, were it to obtain, or the other outcome would have to be worse for someone, were that outcome to obtain. Therefore, the Strong Wide Person-Affecting Principle enables the prioritarian to provide support for the Levelling Down Objection in a manner that is quite consistent with her own commitments.²²

²⁰ For a much more elaborate argument to the effect that existence can be better for an individual than never existing, see Holtug (2001).

²¹ Unlike the Strong Narrow Person-Affecting Principle, the Strong Wide Person-Affecting Principle also allows us to avoid what I call the Problem of Suffering; see Holtug (1998: 170).

²² Note, however, that since some prioritariums may hold prioritarianism partly in virtue of egalitarian concerns, some prioritariums may hold views that do not satisfy the Strong Wide Person-Affecting Principle. Furthermore, while Outcome Welfare Prioritarianism satisfies the Strong Wide Person-Affecting Principle, if we combine this principle and other concerns, the resulting principle *need not* satisfy it. For instance, someone may want to combine prioritarian and egalitarian concerns, and such a combined principle will, obviously, not satisfy the Strong Wide Person-Affecting Principle. (Temkin and Tungodden both seem to accept such a combined view, see Temkin 1993a: 24 and Tungodden 2003: 26.)

6. Impersonal and Person-Affecting Values

While prioritarianism is thus compatible with the Strong Wide Person-Affecting Principle, there is nevertheless a sense in which it is not a person-affecting principle. It implies that a benefit of a fixed size has a higher moral value if it falls at a low level of welfare than if it falls at a high level. But this extra value that is realised at the low level is not a value *for anyone*. Thus, prioritarians seem committed to the existence of impersonal values (see also Persson 2001: 28–9; Temkin 2000: 151–3).

What I have suggested is that what prioritarians ascribe intrinsic value to is compound states of affairs, each consisting in the state that a benefit of a certain size befalls an individual and the state that she is at a particular welfare level. Obviously, both these constituent states concern individual welfare.²³ But when we focus on the *size* of the intrinsic value ascribed to the compound state, we realise that it includes an impersonal element. The difference in value between a compound state that includes a lower welfare level, at which a benefit falls, and a compound state that includes a higher welfare level, at which an equal benefit falls, is not a difference in value for anyone. And, of course, it is the fact that equal benefits can give rise to unequal values that makes it possible for the prioritarian to give priority to the worse off.²⁴

While prioritarians are thus committed to the existence of an impersonal element in intrinsic value, they are not committed to the existence of intrinsic values that contain *only* impersonal elements. In fact, they are committed to the claim that if an outcome is intrinsically better than another, then, compared to this other outcome, the better outcome *necessarily* includes improvements in individual welfare. In this respect, prioritarians differ from egalitarians. Egalitarians are committed to the intrinsic value of states of affairs that are only contingently linked to welfare improvements, namely states of equality.

²³ I do not mean to suggest that prioritarians cannot ascribe intrinsic value to anything but the compound states of affairs under consideration. For instance, they may hold that besides such compound states, one of the constituents, namely the state that a benefit of a certain size befalls an individual, has intrinsic value. This would explain why they are concerned with the *welfare* of the worse off. My point here is merely that, *qua* prioritarians, they ascribe intrinsic value only to the compound state of affairs.

²⁴ However, Dennis McKerlie has suggested a different account of prioritarianism that does not assign impersonal value to benefits (personal communication). He suggests that an equal gain in the lives of two individuals may have an unequal impact on the value of their lives and that this is what is reflected in the judgement that benefits at lower levels have more value than benefits at higher levels. This account relies on a distinction between a gain in welfare and its contribution to the value of the life in which it occurs and I am not quite sure how this distinction should be made. Also, it is incompatible with the claim that increases in life-value that fall at lower levels should have priority and this is a claim that some prioritarians will want to make.

Nevertheless, does the impersonal element in prioritarianism somehow compromise this position? Prioritarianism, as Ingmar Persson points out, is incompatible with the following principle (2001: 28):

The Person-Affecting Improvement Principle. If an outcome, O_1 , is in some respect better (worse) than another outcome, O_2 , then the betterness (worseness) of O_1 for some (collective) is greater than the betterness (worseness) of O_2 is for any (collective).

This principle provides a much stronger link between outcome values and individual welfare than does the Strong Wide Person-Affecting Principle. It links the value of outcomes to the *size* of benefits. Consider the two outcomes in Figure 5.4 (where the slashed line represents the average level of welfare in both H and I). Prioritarians and egalitarians will hold that I is better than H, but this judgement is ruled out by the Person-Affecting Improvement Principle. There is no individual or group of individuals for whom I is better to a greater extent than H is better for any individual or group. Thus, while I is better for q , H is to the same extent better for p . More precisely, the sum of benefits that accrues to q when we go from H to I equals the sum of benefits that accrues to p when we go from I to H.

In fact, the Person-Affecting Improvement Principle implies that just as I cannot be better than H, H cannot be better than I. More generally, while it does not directly imply axiological utilitarianism, it is closely related to this principle.

Now, Persson's point is that since prioritarianism is incompatible with the Person-Affecting Improvement Principle, the prioritarian cannot invoke this principle in defence of the Strong Wide Person-Affecting Principle. Thus, the prioritarian will have to reject what is a conclusive reason for accepting the latter principle.²⁵ And the




FIG. 5.4

²⁵ Persson (2001: 29). Strictly speaking, Persson does not consider the Strong Wide Person-Affecting Principle, but a principle similar to the Slogan. As we have seen, the Strong Wide Person-Affecting Principle is a revised version of the Slogan that allows us to reach the right answer when applied to the Non-Identity Problem. Furthermore, while the Person-Affecting Improvement Principle does imply the Strong Wide Person-Affecting Principle, the former principle gets us into trouble with the Non-Identity

Strong Wide Person-Affecting Principle, of course, is what is supposed to explain at least part of the force of the Levelling Down Objection.

However, it seems to me that the Strong Wide Person-Affecting Principle is quite capable of standing on its own when explicating the force of the Levelling Down Objection. I have suggested that this objection trades on the fact that, in egalitarianism, the link between outcome value and improvements in welfare is entirely *contingent*. In other words, one problem with egalitarianism is that equality can be improved in the absence of increases in welfare. And as I have pointed out, this is exactly the feature that is ruled out by the Strong Wide Person-Affecting Principle. This principle claims that an outcome cannot be improved in any respect in the absence of increases in welfare. Therefore, the Strong Wide Person-Affecting Principle provides a plausible diagnosis of (at least part of) what is troubling about levelling down. And so prioritariums do not need to invoke the Person-Affecting Improvement Principle to defend their commitment to the Strong Wide Person-Affecting Principle.

Now, Temkin suggests that there is in fact a stronger relation to be found between the Strong Wide Person-Affecting Principle and the Person-Affecting Improvement Principle—a relation that I have so far ignored. He claims that reflections on these two principles suggest that those who are attracted to the former should also be attracted to the latter, and that those who reject the latter should also reject the former (Temkin 2000: 152). Hence, since prioritariums reject the Person-Affecting Improvement Principle, they should also reject the Strong Wide Person-Affecting Principle.

However, I have argued that the relation between these two principles is not as strong as suggested by Temkin. A prioritarian may have reasons to hold the Strong Wide Person-Affecting Principle that are quite independent of the Person-Affecting Improvement Principle. For instance, she may hold the former principle because it nicely explains her worries about levelling down as well as about various values that can be realised in the absence of increases in welfare. For instance, it nicely explains the view that there is no value to be found in, say, respecting people's rights or the flourishing of ecological systems in cases where it benefits no one.²⁶ And worries about such values need not hinge on the Person-Affecting Improvement Principle.

In conclusion, prioritarianism not only avoids the Levelling Down Objection, it fully accommodates both the concerns from which this objection derives its force.

Problem and does so for the same reason as does the Slogan. Nevertheless, although I shall not elaborate on this point here, the Person-Affecting Improvement Principle can be revised along the lines of the Strong Wide Person-Affecting Principle, such as to avoid these problems *and* imply this principle.

²⁶ For an extensive list of arguments in which a principle such as the Strong Wide Person-Affecting Principle seems to be implicitly invoked, see Temkin (1993a: 249–55).

7. Relational Justice

I have now argued that prioritarianism handles levelling down cases better than does egalitarianism. This, I believe, gives us a reason to prefer the former view to the latter, although this reason is not conclusive. Nevertheless, it may be suggested that there are also important aspects of our moral thinking that would be lost if we were to abandon egalitarianism in favour of prioritarianism. Thus, Temkin invites us to consider the following case of a 'typical' poor person in the United States:

Ruth is not wretched, but she is a single parent of four, works at two jobs, drives an old car, wonders how she will meet the payments on her two bedroom apartment, and has no idea how her children will afford college on her \$20,000 income. Many are deeply moved by the plight of people like Ruth in a land where so many others live in half million dollar homes, own fancy new cars, send their children to private schools, take expensive vacations, and have annual household incomes well over \$100,000. Is it not clear that the extent to which many are moved by Ruth's situation is heavily influenced not merely by how she fares in *absolute* terms, but by how she fares *relative to the other members of her extraordinarily well-off society?* (Temkin 2003b: 70)

If I understand him correctly, Temkin's point is that the motivational pull Ruth exercises on us is so strong in part because she lives in a society where so many people are much better off than she is. Had her circumstances been the same, but everyone else been much worse off than she is, then we would not have been equally motivated to help her (Temkin 2003b: 70–1).

I am inclined to think that Temkin is right about this. But the question is whether it really shows that, in our moral thinking, we are committed to the value of equality. Let me point to two other features that may help explain our motivational patterns in this sort of case. First, since we are told that there are people (indeed, many of them) who are much better off than Ruth, we tend to focus on how grand it must be for them to be this well off. Being presented with the further details of such lives (half million dollar homes, fancy new cars, etc.), we imagine this rather vividly! We then think of poor old Ruth, worried sick about how to make ends meet, and imagine how grand it would be if *she* were this well off.

If, on the other hand, we are presented with a case in which Ruth is better off than everyone else, we are not as inclined to focus on the sort of life the rich lead, and thus not as inclined to imagine what it would be like for Ruth to lead such a life. Nevertheless, if someone were to prompt us to vividly imagine how nice such a life would be for Ruth (not having to worry about house and car payments, education for her kids, etc.), arguably, the thought of helping Ruth would start to exercise a greater pull on us. The point, of course, is that this explanation of the difference

in motivational pull in the two cases does not rely on the thought that *equality* matters.

Second, and more importantly, motivational pulls may have a tendency to weaken when they are confronted with stronger motivational pulls. By way of illustration, suppose that I feel like having a drink and happen to remember that there is an ice-cold beer in the fridge. Strongly motivated, I open the fridge. I then notice that sitting on the shelf just next to the beer is a delicious bottle of Chardonnay. Realising that I would much rather have the wine, my motivation to have the beer instantly drops. However, it then hits me that the wine is for the guests I am having over. Slightly annoyed, I grab the beer. I know that I shall not enjoy *the beer* any less just because I cannot have the wine, but my motivation to have it has nevertheless dropped. I cannot help thinking of the wine I could have had instead.

Now consider the situation of Ruth. Since Ruth is worse off than most others, she is, if not first in line, then at least in the front part of the queue with respect to whom we most want to help. As prioritarrians, we have a stronger urge to help her than we have to help all those who are better off than she is. So, in terms of our motivation in this case, she does not face much superior competition. If, instead, we imagine that Ruth is better off than everyone else, she is no longer in the front part of the queue. Rather, she is now at the very back. And so, in terms of our motivation to help, she does indeed face superior competition, and lots of it. Therefore, even if she is the only person we *can* help, we may be less motivated to do so. Not because we believe that it is morally less important to help her than it was in the former case, but because we cannot help thinking of how much *better* it would be if we could help some of those people who are much more unfortunate than she is. Importantly, this explanation of our motivational structure does not rely on the thought that *equality* matters, but merely on the thought that priority does.

I do not mean to suggest that the account of our motivational patterns I have just sketched is clearly right and that Temkin's is clearly wrong. In fact, it seems to me that it would be difficult to determine which explanation is best. Nevertheless, I believe that I have provided a somewhat plausible alternative explanation. And the fact that such an explanation can be given weakens Temkin's case for claiming that, in our moral thinking, we are committed to the value of equality.²⁷

²⁷ However, perhaps Temkin is not appealing to our *motivations* but to our *reasons* to help Ruth. Perhaps he takes us to believe that there is a stronger reason to help Ruth if she is worse off than most others than if she is better off than most others. But it is not clear to me that this is what we (or most of us) believe. Alternatively, we may believe that whereas the reason to help Ruth does not change, this reason's relative position in our total pattern of reasons *does* change. That is, it may go from being one of our strongest moral reasons to being one of our weakest. Again, I do not want to claim that this account is superior to Temkin's; merely that it does not seem inferior.

8. Compassion

Let me now turn to a final objection to prioritarianism. According to this view, the worse off always have a higher priority than the better off, no matter how well off the worse off are. Roger Crisp has recently complained that this is implausible. He suggests that our concern for the worse off is based on the virtue of *compassion*. Furthermore, he suggests that our compassion is limited to individuals who are, in absolute terms, poorly off. Once an individual reaches a certain level of welfare, compassion runs out (Crisp 2003: 757).

To accommodate this point about compassion, Crisp proposes a principle with the following structure:

Outcome Welfare Sufficiency. An outcome is in one respect intrinsically better, the larger a sum of weighted individual benefits it contains, where benefits are weighted such that they (a) have zero-value when they fall at or above a specific threshold of individual welfare, l , and (b) have positive value when they fall below l , where this value decreases the closer to l they fall.

Outcome Welfare Sufficiency expresses the idea that benefits become increasingly less important, morally speaking, the better off an individual is, and that there is in fact some level of welfare where further benefits simply have no moral value at all. Of course, a benefit of a fixed size has the same value for the individual to whom it accrues no matter what level of welfare she is at, but its moral value decreases the better off she is. At some point, our compassion ceases and this is where the threshold is.²⁸ Graphically, this view can be represented as in Figure 5.5 (where the slashed line represents the welfare threshold). Note that this graph is similar to that of prioritarianism in Figure 5.1, up to the point where the threshold is reached (this first part of the curve is strictly concave). Thus, Outcome Welfare Sufficiency is in a sense prioritarianism, but with a narrower scope. But while intimately related to prioritarianism, Outcome Welfare Sufficiency nevertheless is not a version of prioritarianism as I have defined it. I have claimed that, in prioritarianism, priority has maximal scope.

It seems to me dubious that compassion runs out above a certain level of welfare.²⁹ Suppose that a given individual is at the threshold level. Let us assume, for the sake

²⁸ There are other proponents of the principle of Outcome Welfare Sufficiency, besides Crisp; see Frankfurt (1987); Rosenberg (1995); Tungodden (2003: 37–40). Frankfurt, it should be noted, is concerned with monetary rather than welfare priority.

²⁹ Perhaps I should qualify this claim. In the dictionary sense of 'compassion', compassion is a feeling directed at the *suffering* of others (see e.g. *The Oxford Advanced Learner's Dictionary of Current English*). If this is what we mean by 'compassion', clearly there are increases in welfare that cannot be objects of this virtue.

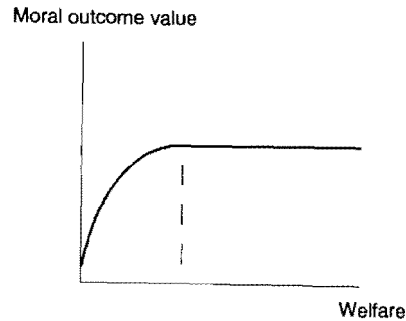


FIG. 5.5

of argument, that this level is quite high. So she is quite well off. Suppose also that, owing to some (very fortunate) changes in the world economy, everyone else becomes much better off than she is. So everyone but her now has a life of extreme luxury and happiness. It seems to me that, in such a case, we would not only regret the fact that she was 'left behind', and so attach value to further benefits to her. We would in fact attach greater value to benefits to her than to (further) benefits to others. If we were in a position to redistribute, we would want to favour her, everything else being equal.³⁰

Even, then, if we were to grant Crisp that compassion is at the heart of priority, we need not agree that there is a level at which compassion sets out. It seems to me more plausible to claim that our compassion is simply greater, the worse off an individual is. Given this, the virtue of compassion is better captured by prioritarianism.

Admittedly, when individuals are very well off, we may *feel* no compassion for them at all, but I suggest that this is because two potentially misleading features affect our emotions. One is that when some people are much worse off than others, we are inclined to focus only on *their* misfortune. However, when there are no such others to 'steal' our attention, we realise that our compassion reaches out even to those who are very well off. The other relevant feature is that what I am suggesting is that compassion is relative to welfare levels; we should thus not expect to feel a *great deal* of compassion for individuals who are doing very well with respect to welfare. In fact,

However, this conception of compassion cannot be what Crisp has in mind. After all, an individual need not actually *suffer* in order to have a low welfare level. Perhaps her life just consists of a few simple joys and nothing else.

³⁰ It may be objected that, in this example, it is the fact that she is now worse off *than others* that is doing the intuitive work. Therefore, I am appealing to egalitarian rather than prioritarian concerns. However, the fact of the matter is that both egalitarians and prioritarians can provide an explanation of why we should redistribute in this case, albeit different such explanations. What the presence of others who are *better* off does, for the prioritarian, is to push forward the worse-off individual in the queue with respect to whom it is most urgent to help.

when individuals are very well off, our compassion may be so reduced that it almost *seems* as if we have none. And when both of these two features are at work at once, we may *feel* no compassion at all.

Let me note a further worry about Crisp's account of distributive justice. Outcome Welfare Sufficientarianism does not assign any weight to benefits above the threshold. Therefore, it does not satisfy the Pareto Principle. In fact, it does not even satisfy the *Weak* Pareto Principle, according to which an outcome is better than another if it is better for everyone. After all, if everyone is at or above the threshold, Outcome Welfare Sufficientarianism does not assign any value to increases to anyone. So even if we could increase everyone's welfare level immensely, it would not make for a better outcome. Surely, this cannot be right.³¹

However, Outcome Welfare Sufficientarianism orders outcomes in one respect only and so may be combined with one or more other principles that do assign value to increases in welfare above the threshold. Along such lines, Crisp suggests that it should be combined with a version of axiological utilitarianism. On the resulting view, priority is given to increases in welfare below the threshold, whereas equal increases have equal value from the threshold up (Crisp 2003: 758). This view satisfies the (Strong *and* Weak) Pareto Principle.

Now, there are different ways in which the two principles can be combined. Crisp's suggestion is that non-trivial benefits below the threshold have absolute priority over benefits above the threshold (2003: 758–9). Presumably, the purpose of the requirement that benefits below the threshold should be *non-trivial* to have absolute priority is to rule out that even the *timiest* benefits below outweigh even the *greatest* benefits above. And indeed, such an implication would be implausible. However, if one is troubled by the claim that trivial benefits below have absolute priority over huge benefits above, then why think that it completely changes the picture if we substitute 'non-trivial' for 'trivial' benefits in this claim? The difference between a trivial and a non-trivial benefit may be very slight indeed.

In reply to this objection, perhaps it may be suggested that the notions of 'trivial' and 'non-trivial' benefits appealed to here are vague. Thus, if we have a spectrum that covers all the possible sizes a benefit may have, there is a range in which it is neither determinately true that the benefits are trivial, nor determinately true that they are non-trivial. And so, between the largest (determinately) trivial and the smallest (determinately) non-trivial benefit there is a range of indeterminacy. Owing

³¹ Nevertheless, perhaps Frankfurt endorses such a view. He claims that 'We tend to be quite unmoved, after all, by inequalities between the well-to-do and the rich; our awareness that the former are substantially worse off than the latter does not disturb us morally at all' (Frankfurt 1987: 32). While, in this passage, Frankfurt addresses only *inequalities* above the threshold, perhaps his concern (or the lack thereof) for the well off can be generalised. At least, he does not in any way suggest that benefits above the threshold matter morally.

to this range of indeterminacy the difference between the two benefits in question cannot be too slight. Accordingly, it may seem less awkward that although a trivial benefit below the threshold does not have absolute priority over even the greatest benefit above, a non-trivial benefit below still does have just that.

However, this reply introduces a new problem, namely what to say about trade-offs between benefits in the indeterminacy range that fall below the threshold and benefits above the threshold. Do indeterminate benefits below have absolute priority over benefits above? If we answer yes, then our original problem reappears at another level: while the largest possible trivial benefit below the threshold does not have absolute priority over benefits above, a slightly bigger benefit in the indeterminacy range does. If, on the other hand, we answer no, then the original problem again reappears. The largest possible benefit in the indeterminacy range that falls below the threshold does not have absolute priority over benefits above, even though a slightly greater non-trivial benefit below the threshold does have such priority.

Perhaps, then, it is better to say that it is neither determinately true nor determinately false that benefits in the indeterminacy range that fall below the threshold have absolute priority over benefits above. However, this leaves us with the problem that Crisp's combined principle does not always provide us with determinate answers, not because the empirical reality to which it is applied is complex, but for purely theoretical reasons. Of course, if the indeterminacy range is rather narrow, so is the range of cases in which the principle does not yield determinate answers. But if the indeterminacy range is narrow, then the difference between (determinately) trivial and (determinately) non-trivial benefits is correspondingly slight. Thus, again, it may seem strange that while a non-trivial benefit below the threshold has absolute priority over benefits above, a trivial benefit below does not.

I also have another worry about Crisp's suggestion. Consider the following outcomes (the numbers represent individual welfare levels):

J: (8, 100) L: (8, 8, 8, 8, 8, 100)
K: (10, 10) M: (9.9, 9.9, 9.9, 9.9, 9.9, 10)

Assume that the threshold level is 10 and that an increase in welfare has to contain at least two units in order to amount to a non-trivial benefit. Clearly, then, Crisp's principle implies that K is better than J. After all, in the move from J to K, there is a (non-trivial) gain of two welfare units for the first individual, which falls below the threshold. Therefore, it has absolute priority over the loss of 90 units by the second individual above the threshold. Now consider L and M. Intuitively, it would seem that if K is better than J, then M is better than L. After all, in the move from L to M five people each gain *almost* as much as *one* person does in the move from J to K, whereas one person loses 90 units in both moves. However, Crisp's principle does not entail that M is better than L. In fact, it says nothing about how to compare them.

Crisp may of course invoke the utilitarian element in his principle to deal with trivial benefits below the threshold, but note that L has a higher total of welfare than M does, entailing that this utilitarian element cannot explain why M should be better.

Finally, since it implies that equal benefits above the threshold have equal value. Crisp's account remains at variance with at least my intuitions about conflicts above the threshold. Consider Figure 5.6 (where the slashed line represents the average level of welfare in N and O, and the continuous line represents the welfare threshold). Suppose that the differences between the two groups do not in any way correspond to differences in desert or responsibility. In N, the first group just happens to be much luckier than the second group, although the second group is itself quite well off. In their forests, the first group has discovered some very special fruit trees. The fruits contain vitamins that not only vastly increase people's health, but also allow them to live, say, 200 years. Having heard the good news, the second group has searched their forests but only to find that they have not been blessed with any of those marvellous trees.

Nevertheless, the first group could decide to share their fruits with the second group. Unfortunately, there are not enough fruits such that both groups can obtain the maximal effect of the vitamins. But even so, if the fruits were shared equally, both groups would experience a significant increase in welfare (perhaps the individuals in both groups would experience some improvements in their health and live, say, 120 years). In fact, the increase that would be experienced by the second group would exactly equal the loss that would be experienced by the first group. In other words, if the fruits were shared equally, O would result.

Now, according to Crisp's account, N and O are equally good. This is because, above the threshold, an extra unit of welfare always counts the same. However, intuitively, O seems to be better than N. Indeed, I suggest that O is better for the same reason that an equal distribution of a particular sum is better below the threshold. Welfare matters more at lower levels. All in all, then, it seems to me that prioritarianism will handle cases of conflict both across and above the threshold better than Crisp's account.

Crisp has, in response to my fruit tree scenario, argued that if we are inclined to prefer O to N, then it just shows that the threshold has been set too low (personal communication). Thus, had the more unfortunate group in N been at a suitably *high*




FIG. 5.6

level of welfare, we would no longer be inclined to believe that it would be better to distribute the fruits equally. In fact, no matter where we place the threshold, Crisp may reply that if we take an equal distribution above the threshold to be better, this is because the threshold is still too low.

I have some worries about this line of argument. First, the proponent of sufficientarianism cannot just dismiss cases in which this principle seems to have implausible implications by saying that the threshold has been set too low. Surely it is up to him to point to a threshold and tell us why it is plausible.

Second, it seems to me that the fruit tree scenario contradicts Crisp's suggestion even at very high levels. Suppose that everyone in N lives for at least 10,000 years but that the lucky first group gets an additional 10,000 years. In O, on the other hand, the welfare is more evenly spread out, since each group gets an equal share of the fruits. Suppose also that N is the actual outcome. It seems to me that, standing at the end of their lives and yearning for more excellent years to live, the second group may reasonably claim that it would be fairer and better if the first group had shared their fruits with them. Why should *they* (the second group) have to die when there is an alternative outcome in which *everyone* has, say, 3,000 happy years ahead of them (with no loss in total welfare)?

Finally, it is worth pointing out that if the threshold is set at a sufficiently high level, sufficientarianism will have no distinct practical significance whatsoever. If, say, the threshold is reached only after 200 years of pure ecstasy, no one will ever be above the threshold and so, for all practical purposes, sufficientarianism will coincide with prioritarianism. It may be replied that, for practical purposes, prioritarianism does not differ much from utilitarianism either, and that a prioritarian would therefore be ill advised to push this point too far. However, it is hardly obvious that prioritarianism and utilitarianism do not differ much on a practical level. Furthermore, if they do not, this will primarily be because it is very difficult to assess welfare precisely and so to know where utilitarianism and prioritarianism come apart in their practical recommendations. The practical implications of sufficientarianism, on the other hand, would coincide with the practical implications of prioritarianism even if we had full knowledge of the possible welfare distributions.

9. Conclusion

I have provided an account of the distinction between egalitarianism and prioritarianism and argued that it renders the former but not the latter principle vulnerable to the Levelling Down Objection. I have also suggested that there are two distinct concerns from which this objection derives its force and that prioritarianism, unlike

egalitarianism, is compatible with both of them. This, I have suggested, gives us a reason to consider prioritarianism superior to egalitarianism, albeit not a conclusive reason.

Nevertheless, it has been suggested that there are important aspects of our moral thinking concerning relational justice that would be lost if we were to abandon egalitarianism in favour of prioritarianism. However, I have suggested that these aspects are in fact explained equally well in prioritarian terms. Finally, I have considered the suggestion that priority is based on the virtue of compassion, and that we only have compassion for individuals who are, in absolute terms, poorly off. Against this I have argued that there is no such limit to compassion and pointed to some further problems with this suggestion.

Thus, I have defended prioritarianism against various objections. And in the process of doing so, I have suggested that prioritarianism is superior to egalitarianism. However, I have not provided anything like a positive case for prioritarianism. So, for all I have said, prioritarianism may be more plausible than egalitarianism and yet not plausible enough to be justifiable. Nevertheless, many political theorists are attracted to views that favour redistributing welfare from the better off to the worse off and, if my arguments are sound, perhaps these theorists should go for priority rather than equality.

References

- ARNESON, R. J. (2000), 'Luck Egalitarianism and Prioritarianism', *Ethics*, 110/2: 339–49.
- BARRY, B. (1989), *Theories of Justice* (Berkeley: University of California Press).
- BLACKORBY, C., DONALDSON, D., and WEYMARK, J. A. (1984), 'Social Choice with Interpersonal Utility Comparisons: A Diagrammatic Introduction', *International Economic Review*, 25/2: 327–56.
- BROOME, J. (1991), *Weighing Goods* (Oxford: Basil Blackwell).
- (forthcoming), 'Equality Versus Priority: A Useful Distinction', in D. Wickler and C. J. L. Murray (eds.), *'Goodness' and 'Fairness': Ethical Issues in Health Resource Allocation* (World Health Organization).
- BROWN, C. (2003), 'Giving Up Levelling Down', *Economics and Philosophy*, 19/1: 111–34.
- CRISP, R. (2003), 'Equality, Priority, and Compassion', *Ethics*, 113/4: 745–63.
- FLEURBAEY, M. (forthcoming), 'Equality Versus Priority: How Relevant Is the Distinction?', in D. Wickler and C. J. L. Murray (eds.), *'Goodness' and 'Fairness': Ethical Issues in Health Resource Allocation* (World Health Organization).
- FRANKFURT, H. (1987), 'Equality as a Moral Ideal', *Ethics*, 98/1: 21–43.
- HAUSMAN, D. M. (forthcoming), 'Equality Versus Priority: A Badly Misleading Distinction', in D. Wickler and C. J. L. Murray (eds.), *'Goodness' and 'Fairness': Ethical Issues in Health Resource Allocation* (World Health Organization).
- HOLTUG, N. (1998), 'Egalitarianism and the Levelling Down Objection', *Analysis*, 58/2: 166–74.

- HOLTUG, N. (1999), 'Utility, Priority and Possible People', *Utilitas*, 11/1: 16–36.
- (2001), 'On the Value of Coming into Existence', *Journal of Ethics*, 5/4: 361–84.
- (2003a), 'Good for Whom?', *Theoria*, 69/1–2: 4–20.
- (2003b), 'Welfarism—the Very Idea', *Utilitas*, 15/2: 151–74.
- (forthcoming), 'Equality, Priority and Levelling Down', *International Journal of Ethics*.
- HOOGER, B. (2000), *Ideal Code, Real World* (Oxford: Clarendon Press).
- JENSEN, K. K. (2003), 'What Is the Difference Between (Moderate) Egalitarianism and Prioritarianism?', *Economics and Philosophy*, 19/1: 89–109.
- KYMLICKA, W. (1990), *Contemporary Political Philosophy* (Oxford: Clarendon Press).
- MCKERLIE, D. (1994), 'Equality and Priority', *Utilitas*, 6/1: 25–42.
- (1996), 'Equality', *Ethics*, 106/2: 274–96.
- (1997), 'Priority and Time', *Canadian Journal of Philosophy*, 27/3: 287–309.
- (2003), 'Understanding Egalitarianism', *Economics and Philosophy*, 19/1: 45–60.
- NAGEL, T. (1991), *Equality and Priority* (New York: Oxford University Press).
- NG, Y.-K. (1990), 'Welfarism and Utilitarianism: A Rehabilitation', *Utilitas*, 2/2: 171–93.
- PARFIT, D. (1984), *Reasons and Persons* (Oxford: Clarendon Press).
- (1991), *Equality or Priority? The Lindley Lecture* (Lawrence: University of Kansas).
- PERSSON, I. (2001), 'Equality, Priority and Person-Affecting Value', *Ethical Theory and Moral Practice*, 4/1: 23–39.
- RAZ, J. (1986), *The Morality of Freedom* (Oxford: Oxford University Press).
- ROSENBERG, A. (1995), 'Equality, Sufficiency, and Opportunity in the Just Society', *Social Philosophy and Policy*, 12/2: 54–71.
- SCANLON, T. M. (1982), 'Contractualism and Utilitarianism', in A. Sen and B. Williams (eds.), *Utilitarianism and Beyond* (Cambridge: Cambridge University Press), 103–28.
- SCHEFFLER, S. (1982), *The Rejection of Consequentialism* (Oxford: Clarendon Press).
- SEN, A. (1987), *On Ethics and Economics* (Oxford: Basil Blackwell).
- TEMEIN, L. S. (1993a), *Inequality* (New York: Oxford University Press).
- (1993b), 'Harmful Goods, Harmless Bads', in R. G. Frey and C. W. Morris (eds.), *Value, Welfare, and Morality* (Cambridge: Cambridge University Press), 290–324.
- (2000), 'Equality, Priority, and the Levelling Down Objection', in M. Clayton and A. Williams (eds.), *The Ideal of Equality* (London: Macmillan), 126–61.
- (2003a), 'Egalitarianism Defended', *Ethics*, 113/4: 764–82.
- (2003b), 'Equality, Priority or What?', *Economics and Philosophy*, 19/1: 61–87.
- (2003c), 'Personal Versus Impersonal Principles: Reconsidering the Slogan', *Theoria*, 69/1–2: 21–31.
- TUNGODDEN, B. (2003), 'The Value of Equality', *Economics and Philosophy*, 19/1: 1–44.
- VALLENTYNE, P. (2000), 'Equality, Efficiency, and the Priority of the Worse Off', *Economics and Philosophy*, 16/1: 1–19.
- WEIRICH, P. (1983), 'Utility Tempered with Equality', *Nous*, 17/3: 423–39.