

INEQUALITY, INJUSTICE AND LEVELLING DOWN

Thomas Christiano and Will Braynen

Abstract

The levelling down objection is the most serious objection to the principle of equality, but we think it can be conclusively defeated. It is serious because it pits the principle of equality squarely against the welfares of the persons whose welfares or resources are equalized. It suggests that there is something perverse about the principle of equality. In this paper, we argue that levelling down is not an implication of the principle of equality. To show this we offer a defence of, and partial elaboration of, what we call a *common good conception* of the principle of equality, which principle favours states in which everyone is better off to those in which everyone is worse off. We contrast this with what we call a *purely structural conception* of the principle of equality. The common good conception of equality involves two basic components: (1) in each circumstance there exists an ideal egalitarian distribution, which distributes equally all the available good in the distribution with the highest average welfare and (2) in evaluating how just the world is, it will matter how far the actual distribution is from the ideal distribution. The ideal egalitarian distribution in the circumstance is Pareto optimal and the approximation rule implies that Pareto superior states are less unjust than Pareto inferior states.¹

Introduction

The levelling down objection is the most serious objection to the principle of equality, but we think it can be conclusively defeated. It is serious because it pits the principle of equality squarely against the welfares of the persons whose welfares or resources are equalized. It suggests that there is something perverse about the principle of equality. Some of the usual examples of the objection should give the flavour of this perversity. Probably the

¹ We thank Geoffrey Brennan, David Chalmers, G. A. Cohen, John Deigh, Andrei Marmor, Calvin Normore, Mark Schroeder, Andrew Williams, and Gideon Yaffe for helpful discussion.

most notorious example is the lifeboat case in which there are, say, three people in a boat that is at risk of sinking unless one person is removed. All will die if no one is removed; only one will die if one person is removed. The common sense approach to this case is to distribute lots to each and select the person with the shortest (or longest) lot to be removed from the boat. But many argue that the principle of equality recommends that all die rather than that only one dies. The former outcome is egalitarian while the latter is inegalitarian. The claim is that the principle of equality favours levelling down the welfares of persons over the inequalities even though some are better off and no one is worse off under inequality. For another example, suppose two persons contribute in an equally deserving manner to the production of two goods. But the two goods cannot be divided and are of unequal worth. Some claim that the principle of equality would recommend throwing away the two goods in order that the two persons are not unequally rewarded. They say that equality may require levelling down rather than unequal rewards. For a final important example suppose that productive efficiency requires a significant use of markets and markets bring about inequalities that cannot be entirely effaced while a more centralized system of production is more equal in its rewards but far less productive. Some claim that the principle of equality implies that we ought to choose the centralized system of production even if everyone is better off under the partially decentralized system. Once again, the principle of equality is thought to imply the levelling down of the welfares of persons so as to ensure equality.

To be sure, those who embrace the idea that equality implies levelling down are careful to say that it is only one consideration among many. They may agree with most people that the one ought to die rather than all or that the goods ought to be distributed unequally. They simply say that the principle of equality implies levelling down but that there are other considerations that tell against levelling down, which often defeat the levelling down implication. Still, the idea is that the principle of equality itself implies levelling down.

Before continuing we want to point out that levelling down may apply for the same reasons to some non-egalitarian conceptions of distributive justice as they do to egalitarian ones. Consider comparative desert theories. Suppose that two people produce two things. One of these things is very valuable and the other is only a little bit and the things cannot be divided up. Now we must

distribute these things between the two of them. Now suppose that one of them is a bit more deserving than the other but the two objects are very different in value so that their distribution cannot fit the differential merits. Indeed, the only way to fit the merits would be to give the small object to the more deserving person and destroy the large object. This is because the more deserving person is only a little more deserving and the smaller object has only a little value. Would such a desert theory require the destruction of the larger, more valuable object?

Here we can see pressures in the direction of distributing the more valuable object as well as pressures in the direction of destroying the more valuable object. On the one hand, some will say, the less deserving person deserves something. And if we distribute in accordance with the exact merits, the less deserving will not get anything. On the other hand, if we make sure the less deserving gets something, the less deserving will get much less than he deserves at least by comparison with the more deserving. Presumably similar reasoning could apply to principles that require distribution in accordance with need. So the worry about levelling down is much more general than an objection to equality.

One way to respond to the levelling down objection is to say that it is not an objection at all. This response accepts that equality does imply levelling down but that this is not a problem. Partly this response is predicated on the observation above that equality is only one consideration among many. We will argue in what follows that this is not a desirable way to go. We will show that there are considerations internal to the principle of equality (and many other comparative principles of justice) that suggest that if levelling down is a genuine implication of the principle of equality, then this is a real problem for the principle.

In this paper, we argue that levelling down is not an implication of the principle of equality. Indeed, we argue that the principle of equality, once properly understood, is opposed to levelling down. It always favours states in which everyone is better off, even if unequally so, to ones in which everyone is worse off even if equally well off. We will show that though the principle of equality implies that all inequalities are unjust, some inequalities are less unjust than some equalities, namely those in which everyone is worse off. To the extent that the levelling down objection is the principal objection to the principle of equality, we think this argument should go a long way to undermining opposition to the principle of equality.

To show this we offer a defence of, and partial elaboration of, what we call a *common good conception* of the principle of equality, which principle favours states in which everyone is better off to those in which everyone is worse off. We contrast this with what we call a *purely structural conception* of the principle of equality. The purely structural conception of equality evaluates distributions of goods exclusively in terms of the comparisons across persons. On this conception the only information relevant, from the point of view of the principle of equality, to the evaluation of distributions of goods are comparisons across persons in terms of their welfare or opportunities for welfare. A purely structural conception of equality does imply levelling down and if the arguments of this paper are sound this is reason to prefer the common good conception over the purely structural conception.

To give a brief preview, the common good conception of equality will involve two basic components: (1) in each circumstance there exists an ideal distribution (or the ideally just distribution), and (2) in evaluating how just the world is, it will matter how far your actual distribution is from the ideal distribution. Thus, the first step is to identify the ideally just distribution. On our theory, this ideal distribution is an equal distribution of the feasible goods available in the circumstances. We note that the general strategy we outline in this paper should be of interest to any distributive justice theorist who is concerned with the justice of patterns.² A comparative desert theorist for example, might instead argue that the ideally just distribution is something else. The second step is to use one's ideal theory (the ideal distribution) to evaluate any actual or current state of affairs (some arbitrary distribution) by measuring the difference from the ideal distribution to the actual distribution. In other words, to know exactly how unjust some arbitrary distribution is, we want to know by how much that distribution diverges from the ideally just distribution. It is in these two steps that the levelling down objection is defeated. The ideal distribution equalizes all the feasible gains in a circumstance and thus does not waste any goods. The rule of approximation then says that while inequalities are unjust, they

² Not all distributive theories are of this sort of course. Nozick's libertarian theory of property entitlement that is governed by a just process denies that patterns or distributions occupy a central place in a theory of distributive justice because on Nozick's view, a just process will upset any distribution of goods. See *Anarchy, State and Utopia* (New York: Basic Books, 1974), chapter 7.

are not as unjust as equalities in which everyone is worse off. It is this idea that we defend and elaborate in this paper.³

In what follows we present a version of the levelling down objection using cardinal and interpersonally comparable utilities. We then give an argument for why egalitarians (and by implication many other comparative principles of distributive justice) should take the levelling down objection seriously. Then we show that there is a fundamental gap in the argument that alleges that levelling down is an implication of equality. Finally we elaborate and defend a conception of equality that rejects levelling down. And we elaborate a rule for comparing distributions of goods.

The levelling down objection

It has been said by some that the principle of egalitarian justice is subject to a fatal intuitive objection. The objection is that the principle of equality has an extremely implausible implication. Suppose two alternative states (2, 2) and (7, 3) in which two persons are represented in terms of the amount of welfare they each enjoy in that state. The numbers are cardinal and represent interpersonally comparable utilities. To some, the principle of equality appears to say that (7, 3) is worse than (2, 2) at least in the respect relevant to the principle of equality. This is because (2, 2) is egalitarian while (7, 3) is not. (7, 3) represents a departure from equality while (2, 2) does not. Thus the principle of equality appears to imply that, at least as far as equality is concerned, we ought to make everyone worse off. Of course, other principles may contend with equality and override its recommendation in this case. But the worry is that to the extent that the principle of egalitarian justice makes the recommendation that everyone be made worse off, that is a strike against the principle.⁴

³ A positive argument for the principle of equality understood according to the common good conception is in Thomas Christiano, 'A Foundation for Equality', *Egalitarianism: New Essays on The Nature and Value of Equality*, ed. Nils Holtug and Kasper Lippert-Rasmussen (Oxford: Oxford University Press, 2007) and in Chapter 1 of Thomas Christiano, *The Constitution of Equality: Democratic Authority and Its Limits* (Oxford: Oxford University Press, 2008).

⁴ See Derek Parfit, 'Equality or Priority?', The Lindley Lecture (University of Kansas, 1991), p. 23 for an account of the objection. See also Jan Narveson, 'On Dworkinian Equality', *Social Philosophy and Policy*, Volume 1, issue 1, pp. 1–22, Harry Frankfurt, 'Equality

Why the levelling down objection matters

The levelling down objection would be a real problem for the principle of equality, were it true that levelling down was an implication of the principle of equality.⁵ For most egalitarians seem to hold both to the idea that equality is important and to the idea that well-being is important (or at least opportunity for or access to well-being or capacities for functioning). And these two judgments seem to come together in their egalitarianisms. For one thing, many egalitarians think that the promotion of well-being or at least the opportunities or access to well-being is important.⁶ This suggests that these egalitarians cannot be indifferent between two egalitarian states (2, 2) and (5, 5), which are such that everyone is better off in the latter. Egalitarians must prefer the Pareto superior equality to the Pareto inferior equality, and that preference derives from the correct understanding of the principle of equality, or so we will argue.

Here is the basic argument for this thesis. In what follows we will usually speak for simplicity's sake as if well-being is the relevant metric for equality. We think the same considerations will hold for any other fundamental good that serves as a metric of equality such as resources (as Dworkin defines them) or opportunities for welfare or access to advantage. There is an internal connection between the rationale for equality and the value of the relevant fundamental good that is equalized. If it were not true that more well-being is better than less, then there would be no point to equality. There would be no reason to care about equality. Since the importance of well-being seems to be built in to the rationale for the principle of equality – it is the reason for the principle taking the shape that it does – an egalitarian cannot be indifferent between these two states.

as a Moral Ideal', in his *The Importance of What We Care About* (Cambridge: Cambridge University Press, 1998), for other sympathetic accounts of this objection.

⁵ Not everyone agrees that this is a strike against the principle of egalitarian justice. Larry Temkin argues that there are other principles that may have strongly Pareto inefficient implications. Temkin cites principles of retributive desert that imply in some circumstances that everyone ought to be worse off than they could because they deserve to be worse off. He argues that many accept such a view despite its welfare diminishing character. See Larry Temkin, 'Equality, Priority and the Leveling Down Objection', in *The Ideal of Equality*, ed. Mathew Clayton and Andrew Williams (New York: St. Martin's Press, 2000), pp. 126–161, esp. p. 138.

⁶ See G. A. Cohen, 'The Currency of Egalitarian Justice', *Ethics*, October 1989 for a thorough discussion of some different options for understanding the metric of equality.

To take an example, suppose that we are concerned with distributing bread among persons and we have much more than is needed. In this context there is a definite level at which each person has enough and beyond which more does not matter. In this context we would only be concerned that each receive enough of the bread and beyond that we would not care how much each gets. Beyond the satisfaction of need and culinary taste whether we distribute equal amounts of bread or not would be, in and of itself, a matter of indifference.⁷

To be sure, if we did not have enough for everyone, we might be concerned with equal distribution. So if we had enough for everyone to survive but not enough to satisfy everyone we would be concerned with the correct distribution of bread. But in this case, it is precisely because the amount of bread we have is such that more is better than less for everyone that we concern ourselves with its distribution.

Another example is a concern over how many letters people have in their last names. For the most part people are indifferent to how many letters they have in their names. As a consequence, they will be indifferent to the quantitative distribution of letters in each person's name. Equality could not be important in this context.

So a necessary condition for equality mattering is that the thing being equalized is such that more is better than less. We want to argue that since the truth of the proposition that more substantial good is better than less substantial good is a necessary condition for there being a rationale for the principle of equality in the substantial good, the right account of the principle of equality must somehow include the idea that equalities in which everyone is better off are better than equalities in which everyone is worse off.

Clearly this proposition is not a conceptual one. One can imagine a principle of equality of shares in substantial good that does not concern itself with whether people have more of that substantial good. But the question is: does such a principle make

⁷ To be sure, if each person could sell his excess bread, then equality above the level of enough bread would make sense. The observation about the distribution of bread not mattering above the level of sufficiency is the central insight of defenders of sufficiency theories such as Harry Frankfurt in his 'Equality as a Moral Ideal'. For critiques of sufficiency as a general account of distributive justice see Paula Casal, 'Why Sufficiency Is Not Enough', *Ethics*, 117 (January 2007), pp. 296–326 and Thomas Christiano, 'A Foundation for Egalitarianism'.

sense? The things of which we care about egalitarian distribution are things that we want more of rather than less for everyone. And it is because we want more of rather than less of these things for everyone that we think that egalitarian distribution of these things matters.

Moreover, the values over which the principle of equality ranges do not have any components that justify lowering the welfares of persons (or the opportunity for welfare). To appreciate this, note the contrast with some other conceptions of justice. Suppose one holds a non-comparative desert principle that states that the vicious ought to suffer and the virtuous ought to thrive. If every person in a society is vicious, then they all deserve to be badly off. This would favour a Pareto inefficient outcome in which everyone is badly off. The introduction of a new set of resources that could make everyone better off would not be an improvement; indeed, on the desert theory if everyone were to be made better off than they deserved to be, that would be worse. On these views, it is better that a person who deserves it be worse off. This is part of the value theory of these conceptions of desert. We may reject this value theory, but it makes some sense and levelling down makes sense in the context of this kind of principle. But there is no analogous feature in egalitarian principles. The egalitarian theorists are not in the position of the non-comparative desert theorist. There is nothing in the value theory that says that it is better for a person to be worse off.

Because there is an internal connection between the importance of equality and the idea that it is better to have more rather than less of the thing being equalized, the levelling down objection is an objection that ought to be taken seriously by egalitarians, if it works. But we do not think that it works. That is we do not think that levelling down is an implication of the principle of equality. And the fact that it does not work can be seen in a number of ways.

A gap in the levelling down objection

The levelling down objection derives its apparent strength from the claim that an egalitarian must think that something is lost when there is some inequality. This claim follows from the central egalitarian claim that all inequalities are unjust. And from this it is inferred that any egalitarian distribution must be better than any

non-egalitarian distribution, at least in one respect. And from this it is inferred that there is one important respect in which an egalitarian distribution is better than a Pareto superior state (i.e. a state in which at least one person is better off and no one is worse off).

But, the inference above is problematic. From the observation that there is loss from inequality it does not follow that every egalitarian distribution is better in respect of the principle of equality than every non-egalitarian distribution. What the egalitarian must say is that every non-egalitarian distribution is unjust because it is not equal. Only equality is fully just. But the egalitarian need not say that every egalitarian state is better than every non-egalitarian state.

How can it be that only equality is just but not that every equality is better than every inequality? The basic idea is that the egalitarian can and should say that for every circumstance, there is an ideal egalitarian distribution such that failure to realize that ideal egalitarian distribution in the circumstances is unjust. Instead of identifying a purely structural feature of distributions that can be satisfied just as much by Pareto inferior states as by Pareto superior states, the present suggestion is that the principle of equality identifies first and foremost an ideal distribution in each relevant circumstance. It is this ideal distribution that embodies what is fully just in the circumstance. Any departure from this ideal distribution in the circumstance is unjust. A complete principle of equality then defines a rule of approximation to this ideal egalitarian distribution. It determines which departures from the ideal egalitarian distribution are more unjust and which are less unjust. It gives a ranking of states in terms of how close to the ideal egalitarian distribution they are. Our thought is that this rule of approximation need not entail that all equalities, no matter how inefficient, are better than all inequalities. This is a substantive matter on which there must be argument for one rule over another.

To see this, think of the division of a pie between two equally deserving people and where nothing else is to be taken into account, so this is a complete description of the circumstances of justice for these two people. The ideal egalitarian distribution of this pie is to cut it in two pieces of equal size. That is what egalitarian justice requires in the circumstance. Any other division is, let us suppose, unjust. Now ask how other divisions of this pie would rank in terms of closeness to the ideal distribution.

Suppose, for some crazy reason, that the pie could not be cut into equal parts. It can only be cut into two somewhat unequal parts or the whole thing could be thrown away. The question for the principle of equality as we understand it here is which of these two alternative distributions is less unjust? And the way to answer this question is to determine which of the two alternative distributions is closest in terms of justice to the ideal egalitarian distribution. Once the egalitarian has identified the ideal egalitarian distribution and has asserted that all other distributions are unjust in the circumstance, there is no reason to suppose that the egalitarian must think that the distribution closest in justice to the ideal distribution is the one in which the pie is thrown away. It is a nontrivial task to determine what rule of approximation to use in deciding which distribution is closest to the ideal distribution. We may think that a rule of approximation to the ideal that ranks the Pareto superior but unequal division as less unjust than the Pareto inferior but equal division is the best one. Indeed we will vindicate this thesis in what follows. But for now it is worth noting that holding to a principle of equality does not require that one choose the approximation rule that favours the Pareto inferior equality. Hence, levelling down need not be an implication of a principle of equality.

Consider the three states we have discussed so far in abstract terms. (2, 2) and (5, 5) are egalitarian and (7, 3) is non-egalitarian. (5, 5) and (7, 3) are both strongly Pareto superior to (2, 2). (7, 3) and (5, 5) are Pareto non-comparable.⁸ But (5, 5) is egalitarian and (7, 3) is not. The difference is that in (7, 3) at least one person is better off and another is worse off than in (5, 5). All the egalitarian is committed to asserting is that there is something lost in (7, 3) because all the people are not equally well off. This may merely imply that (5, 5) is better or more just than (7, 3). It is compatible with this to say that (7, 3) is better or more just than (2, 2), just as (5, 5) is. And this set of claims is sufficient to make sense of the claim that for every inequality, there is something lost with respect to equality.

Of course, (5, 5) may not be feasible but this need not be a reason to think that it is not what justice requires in the circumstance or that (7, 3) is defective precisely because it diverges from

⁸ Two states are 'Pareto non-comparable' when they have different distributions but neither is Pareto superior to the other.

(5, 5). As we see it, a state of affairs can be unjust even if there is no way to improve it and we can have a conception of what is just under the circumstances even if we cannot bring about that justice. Justice does not obey the 'ought implies can' principle entirely. In part, it serves as an ideal to be approximated.

We can see this in a variety of contexts. We know that even the best penal system is likely to convict some innocent persons and let some guilty persons go free and we are inclined to think that these are injustices even though we cannot improve the system that makes these errors beyond a certain point. We think we can see it in a distributive case as well. If two persons have contributed equally to the production of a pair of unequal but indivisible substantial goods, we may think that there is some injustice in giving the better good to one and the lesser good to the other. Each person in the circumstance will feel some regret at the imperfect justice of the distribution, even if they use a lottery to determine who gets what. No one can be blamed for the inequality and the use of an egalitarian lottery goes some way to alleviating the sense of injustice but it does not go away entirely.

All that is needed for the egalitarian is to say that the unequal state is missing something or is unjust because it is not equal. This is still an importantly egalitarian theory but it does not say that any equality is better in some respect than any inequality. It says there is an ideal egalitarian distribution in each circumstance and that every inequality is unjust because it is not equal.

What our argument so far has shown is that there is a significant gap in the levelling down objection against equality. It has shown that a crucial inference is unsupported by the objector and that alternatives to that inference are logically possible. What we want to show in the next section is that the logically possible alternative missed by the proponents of the levelling down objection is in fact the best account of equality.

What we argue for in the following is not a complete principle but a strategy for constructing such a principle. We do not have the time to construct a complete principle and we are not sure we can do it at the moment anyway. But it is a defensible strategy, which implies that there is some principle that satisfies the strategy, which is the correct principle of equality. If this is a legitimate strategy, then it is sufficient to defeat the levelling down objection.

An argument for Pareto optimal inequalities

There are a number of cases in which inequality is Pareto superior to feasible equality. In one sort of case, some goods are lumpy, so we cannot achieve a completely egalitarian distribution. The second case involves production. In some cases, the complexities and uncertainties of production require that incentives be offered to those who are most suited to the tasks to be performed. In this kind of case, inequalities arise as a kind of by product of the process of production.⁹ A third case is the lifeboat case in which it is not possible that every person live equally good lives. Our thesis is that if all persons can be made better off than under the best feasible equality, then the principle of equality we have defended implies that we should choose that state in which some are better off and none are worse off than under the best feasible equality.

First, we will consider whether it is just to bring about a strong Pareto improvement over feasible equality. So we will compare just two states and we will just consider this for two persons: A and B. This is a narrow idealization but it is hard enough to grasp it properly. The argument proceeds by comparing two states: (2, 2) in which A and B are equally well off and (7, 3) in which both are better off than in the first but A is better off than B.

The strong Pareto improvement in (7, 3) pushes us in a direction that does not allow us fully to satisfy the constraint stated by the principle of equality. There is a failure of justice in cases of unequal distribution. In our example, either A is being treated better than the reasons allow in his case or B is being treated worse than the reasons allow in his case or both of these claims are true.

If we compare (2, 2) (feasible equality) and (7, 3) (feasible strong Pareto improvement) the egalitarian principle of justice says that there is something wrong in (7, 3) because it is not equal.

⁹ G. A. Cohen has argued quite persuasively in his *Rescuing Justice and Equality* (Cambridge, Mass.: Harvard University Press, 2008), chapter 1, that there is something problematic from the standpoint of justice in allowing more talented persons to receive more as an incentive to work than others receive. This argument casts serious doubt on Rawls's use of incentives in this context in *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971). But incentives may still be required for informational purposes and it may be impossible to wipe the inequalities entirely away without significant cost. See Christiano's 'Cohen on Incentives and Inequality', in *Ethics and Economics*, ed. Christi Favor, Gerald Gaus and Julian Lamont (Palo Alto, CA: Stanford University Press, 2009) for this argument.

But we want to say that there is something even worse in (2, 2) with regard to justice. And so we have reason to prefer the feasible strong Pareto improvement to the feasible equality, from the point of view of justice. What is our reason for this?

The reason is grounded in a common good conception of equality or a conception of equality that includes a concern to advance the interests of each and every person. This conception is the basis of a plausible strategy for constructing a principle of equality that gives us a plausible account of greater and smaller departures from equality. Both (2, 2) and (7, 3) may be unjust and thus departures from justice, but the question is, which one is a greater departure from justice? We will propose a strategy for constructing a principle that gives us a schematic account of how to assess the size of the departure from ideal equality.

The strategy for constructing the principle comes in two stages: first the strategy calls for a rule identifying the ideal egalitarian distribution in a given circumstance. Second, the strategy calls for a rule for assessing greater and lesser departures from the ideal egalitarian distribution.

The ideal egalitarian distribution

The first step is to identify that distribution that would be ideally just in the circumstances. On an egalitarian theory that would be an egalitarian distribution of the substantial goods involved in the circumstances. And this is the distribution relative to which all states are evaluated in the circumstances by justice. Our proposal is that the ideal egalitarian distribution in a given circumstance is a potential Pareto optimal distribution. That distribution is better for the worst off person in the feasible Pareto optimal inequality but not as good for the better off person in the feasible inequality. Since the ideal egalitarian distribution may not be feasible, it is only a potential Pareto optimal point. In the abstract numerical example we are discussing, the ideal egalitarian distribution is (5, 5) in which B is better off than in (7, 3) but A is worse off. And A and B are better off than in the feasible egalitarian state.

How do we determine the ideal egalitarian distribution in a particular situation? First, we need to say how we identify a circumstance. We propose to identify a circumstance or situation in terms of the highest level of average utility that is feasible for the persons involved. So in the numerical example the state that gives

the highest feasible level of average utility of 5 per person is (7, 3) by hypothesis. The circumstance for which we are trying to define the ideal egalitarian distribution is identified by this highest available average utility. If we find that more average utility is possible than previously because of increased production or more manna from heaven then we have moved to a new circumstance.

Second, in order to define the ideal egalitarian distribution for the circumstance we propose to take the highest average level of well-being between A and B in the circumstance and say that the ideal distribution consists in each getting that level of well-being. So for a circumstance in which (7, 3) has the highest feasible average utility, (5, 5) is the ideal egalitarian distribution. To be sure, sometimes this ideal egalitarian distribution is not feasible and sometimes it is.

Combining these two elements, the numerical example we have given tells us that (5, 5) is the ideal egalitarian distribution in the circumstance. Should the circumstances change and a higher average become available, then the ideal egalitarian distribution would change as well. The thought is that there is an ideal egalitarian distribution for every circumstance. To be sure, in most circumstances, there are a number of possible Pareto optimal states. The above rule says that the ideal egalitarian distribution in a particular situation is the mean level of well-being of the best (understood in terms of the total well-being) feasible Pareto optimal state.

A number of considerations favour this way of identifying the ideal egalitarian distribution. One, this conception of the ideal egalitarian distribution accommodates the idea that there is an internal connection between the rationale for the principle of equality and the importance to each person of the thing being equalized. It does this in two ways. First, it does this by using the device of an ideal distribution against which all other distributions in a circumstance are to be evaluated as opposed to evaluating distributions merely on the basis of their structural features. The use of the ideal distribution implies that there is a certain level of provision of goods that is important to justice in a circumstance. A purely structural approach to equality, in contrast, states that all that matters to egalitarians is the shape of the distribution. Such an approach seems to ignore altogether the value of the things being distributed. Second, the ideal egalitarian distribution implies that each should have the highest average level of good that is available.

Two, it seems like a natural egalitarian distribution for the circumstance because it equalizes all the gains from the best feasible Pareto optimal state available. And it seems that this is a natural interpretation of the grounds of regret people might have when two equally deserving people produce two unequally valued and indivisible objects. In our pie example, supposing for some strange reason that it cannot be cut equally, the two equally deserving contributors to the making of the pie will not want to throw it away but they may still experience some regret at the fact that the pie is not distributed equally.

In the lifeboat example we have discussed, in which one person out of the five must be tossed out in order to preserve anyone at all, the ideal egalitarian distribution might involve an equal distribution of the total number of years to live available to the group. In this case, supposing that the remaining four expect to live another thirty years each, the ideal egalitarian distribution would be that all five live twenty four years each. When one person is thrown out of the boat, each has reason to regret that the burden falls entirely on that one person even if the choice of person is made by lot. A natural way to interpret this regret is that the burden did not fall equally on each in such a way that each could have an equal life span. Presumably if this option had been available, they would have chosen this were they egalitarians. They seem to be regretting what has happened because they are evaluating it against a standard of ideal egalitarian distribution.

Three, this conception of the ideal egalitarian distribution seems to strike the right balance between considerations of realism and the ideal in thinking about principles of justice. We have defended the ideal egalitarian distribution against the charge that it is sometimes infeasible. Feasibility is not a necessary condition for justice and the infeasibility of better alternatives is not a sufficient basis for saying that there is nothing wrong with the feasible state. We can see this in the cases of necessarily imperfect penal systems, and with the kinds of regrets we experience when a fully just distribution is not feasible. At the same time, though, this conception of the ideal egalitarian distribution represents a kind of ideal that is suited to the circumstances at hand. The circumstances are defined by the available amount of good to the group in question and in terms of the best feasible Pareto optimal state. Though the ideal equality may not itself be achievable in the circumstances it is nevertheless connected in a realistic way to the circumstances.

So the ideal egalitarian distribution is not merely pie in the sky. Indeed it has real bite. This is because it says that the best off in the Pareto optimal inequalities are unjustly advantaged and the worst off are unjustly disadvantaged.¹⁰ We will also see that it has bite to it when we elaborate a plausible rule for approximating to the ideal egalitarian distribution.

Approximations to justice

The second stage is grounded in the idea that when ideal justice is not achievable, we may still have a conception of what is closer to the ideal and what is further away. In particular, we may still have a conception of what is more just and what is less just in relation to the ideal. And for this we develop a rule for determining approximations to justice. An egalitarian theory will determine which departures from the ideal egalitarian distribution are greater departures from justice and which departures are smaller ones. The smaller departures from justice are more just or less unjust than the greater departures from justice. In this way, even if the ideal egalitarian distribution is not feasible itself, it sets an ideal to which we can approximate and the principles governing the measurement of approximation are what we are concerned with in this second stage. In this way, the ideal egalitarian distribution has practical significance even if it cannot be realized.

We will canvass a couple of proposals in this paper to give a sense of how to think this problem through before we provide what seems to us to be an acceptable rule. First we will discuss what we call the difference summing method of measuring approximations to justice. Then we will propose a more complex method of assessing approximations that is meant to resolve the difficulties in the first approach.

¹⁰ One way that we would argue we should not specify the ideal egalitarian distribution is to say that the equality at the highest level of well-being is the only fully just arrangement. The reason for this rejection is that such a principle would not be distinctively egalitarian. This is because it would say what non-egalitarians could readily agree to, which is that Pareto improvements are to be pursued. We doubt if many want to disagree with that. What makes the principle of equality distinctive is that it recommends making some people worse off if that will make others better off and it recommends this up to the point of equality.

Criteria for a good rule

First, we will describe what we take to be some plausible constraints on the construction of such a rule. These constraints mostly are implied by the ideas underlying the principle of equality.

The first and clearest constraint on the construction of such a rule is that the most just state of affairs in the circumstances must be the ideal egalitarian distribution. So the rule will recommend the ideal equality whenever it is feasible. Because this is an egalitarian theory, we assume that the ideally just distribution in any given circumstance is egalitarian (but not that any egalitarian distribution is ideal or even desirable since, as we have argued above, such a distribution could be very far from the ideal).

The second constraint on the approximation rule is that the rule must never favour Pareto inferior states over Pareto superior states. This aspect of the rule follows from the importance of the well-being of all persons that is essential to the principle of equality. And this aspect of the rule is what ensures that the principle of equality never implies levelling down.

The third constraint on the rule of approximation is that the rule does not track average or total utility because it is not a utilitarian rule. There will be cases in which some states are closer approximations to the ideal egalitarian distribution than others even though the others have a greater amount of average utility. This aspect of the rule of approximation follows from the fact that the principle at issue is not a utilitarian principle and so will prefer many egalitarian states that have lower average utility. But remember this condition is to be combined with the second condition that says that the rule never selects Pareto inferior states as more just than Pareto superior states. Thus the rule will select some states that have less average utility than others on condition that they are all Pareto non-comparable.

The fourth and final constraint on the rule of approximation tells us that the rule must not depart too much from the difference principle. This does not follow directly from any of the aspects of the principle of equality that we have defended. It is an intuitive idea about approximating equality. But it requires some explanation. One, we say that the rule must not depart too far from the difference principle because we think that some realizations of the difference principle get very close to levelling down and are what we call quasi-levelling down. For example if we imagine two states of affairs such that in one each is equally well

off and in the other one person is much better off while another is only a tiny bit worse off than in the equal state, the difference principle will favour the equal distribution. And we can make it so that the difference between the worst off person in the unequal state of affairs and those in the equal state of affairs is arbitrarily small though still significant. The difference principle will always prefer the equal state even if the advantage to the worst off is extremely small while the advantage to the better in the unequal state is much greater. This seems to us to suggest that the difference principle involves something like levelling down in some circumstances, what we call quasi-levelling down. Still, the difference principle in many circumstances does seem to give an intuitively plausible answer to the question of what the best approximation to ideal equality is and so we say that the approximation rule should not be the same as the difference principle but it should not depart too far from it. Admittedly, this is a vague constraint, but our intention here is just to give a sense of what a rule of approximation to ideal equality is that does not imply levelling down. It is worthwhile noting here that this explains two features of the difference principle. One, it is not a principle of justice in the sense that it defines ideal justice in a circumstance. Hence, this approach departs from Rawls's conception of the difference principle as a principle of justice. Nevertheless, the difference principle does stand out in many circumstances as the best way to approximate full justice and so it can be thought of as a principle of justified injustice or even as a principle of the least unjust outcome (with the caveat noted above that it does not always pick out the least unjust state).

The difference summing rule of approximation

Our initial proposal is to measure the departures from justice in terms of the sum of differences in well-being between the state being evaluated and the ideal egalitarian distribution. And the state in which the sum of differences from the ideal egalitarian distribution is smallest is closer to ideal equality and so is the more just state. More formally, this principle is given by the following function: $d(S_{actual}, S_{ideal}) = |x^* - x| + |y^* - y|$ between the ideal distribution (x^*, y^*) and the actual distribution (x, y) . The hope is that d will be closer to zero as the actual distribution gets closer to the ideal distribution. This too seems like a natural method, but

as we will see, it will require some modification. Let us call this the *difference summing principle*.

For (7, 3) we add the difference in well-being for person B between (7, 3) and (5, 5) to the difference in well-being for person A between (7, 3) and (5, 5). And for (2, 2) we add the difference in well-being for B between (2, 2) and (5, 5) to the difference in well-being for A between (2, 2) and (5, 5). If the sum of differences is greater for (2, 2) than for (7, 3), then (2, 2) is a greater departure from justice than (7, 3). And we can see that this holds in this case. There the sum of differences between (2, 2) and (5, 5) is 6. While the sum of differences between (7, 3) and (5, 5) is 4. In general, for the two person cases, this will imply that feasible strong Pareto improvements will be more just than the Pareto inferior feasible equality. This is because the two persons' welfares in the strong Pareto improvement are closer to their welfares in the ideal equality than are the welfares in the Pareto inferior feasible equality. Therefore, the principle of equality that accords with this strategy will say that the strong Pareto improvements are superior to the Pareto inferior feasible equality. To be sure both are unjust, but one is more unjust than the other.

The strategy accords with the kind of principle of equality that values both well-being and a non-arbitrary way in which it is distributed because the ideal egalitarian distribution is the product of a concern for well-being and a non-arbitrary distribution. The requirement of non-arbitrariness ensures that the ideal distribution is one of equality; the concern for well-being ensures that the Pareto superior inequality is more just than the Pareto inferior equality. And the method of assessing departures from equality in terms of the sum of differences of well-being from the ideal equality point seems a highly intuitive way of capturing all these concerns and the idea of departures from equality.

Furthermore, note that the difference summing rule does not give exclusive priority to well-being. Some states that are egalitarian will be superior to some states that are unequal even if the unequal state has more average utility. You can see this in the following numerical example. Here (5, 5) is the ideal egalitarian distribution. (7, 3) has the same average utility as (5, 5) and (4, 4) has a lower average utility than the ideal or (7, 3). Nevertheless, by the difference summing rule, (4, 4) is more just than (7, 3).

This seems like a good result to us. First, the equality that is Pareto inferior to the ideal equality is unjust. It is, however, less

unjust than the best Pareto non-comparable inequality. And, second, the equality that is Pareto inferior to the inequality is less just than the inequality.

It makes sense that a principle of equality would have the implication of sometimes choosing a state in which the total amount of well-being is not maximized because the principle of equality does not track the principle of average utility and the principle of approximation does not track it either. So the thought is that equality can be more just than inequality even if the inequality includes more total utility. This is a welcome result if, like us, you want to reject utilitarianism and so do not want to simply track total or average utility.

But there are serious difficulties with the difference summing rule. To see these, suppose that we have three states $(10, 0)$, $(5, 5)$ and $(1, 1)$. $(10, 0)$ is the best feasible Pareto Optimal state for person A and person B because it has the highest average well-being. $(5, 5)$ is the ideal egalitarian distribution but is not feasible. $(1, 1)$ is an egalitarian point that is inferior to $(5, 5)$ but it is non-comparable to $(10, 0)$. So the situation is as described above so far. The sum of differences of $(1, 1)$ from $(5, 5)$ is 8. The sum of differences of $(10, 0)$ from $(5, 5)$ is 10. So here it looks like $(1, 1)$ is more just than $(10, 0)$. Maybe one isn't disturbed by this implication but this experiment can be repeated to the point where the worst off of the equality is only slightly better off than the worst off in the Pareto optimal inequality. This would be the case if in the feasible equality A and B both had 0.1 as in $(0.1, 0.1)$. Yet it is more just than $(10, 0)$. And we can of course push this further to say $(0.0001, 0.0001)$. It is still better than $(10, 0)$. We are getting asymptotically close to a levelling down objection. No matter how small the difference is between the worst off in the equality from that in the Pareto optimal inequality, the equality is superior. This seems like a very troubling result. Let us call this the *quasi levelling down objection* to the difference summing proposal.

But there is another levelling down objection to the difference summing proposal in the near vicinity. Suppose we take $(10, 0)$ again as the best feasible Pareto optimal inequality and we take $(5, 5)$ as the ideal equality. Now suppose we find that there is a state $(5, 0)$ in which the worst off is the same as the worst off in $(10, 0)$ but the better off is at or near $(5, 5)$ or only slightly better off. So $(5, 0)$ is Pareto inferior to $(10, 0)$ but the sum of differences between $(10, 0)$ and $(5, 5)$ is greater than the sum of differences

between (5, 0) and (5, 5). Here the rule minimizing the sum of differences from ideal equality seems to say (5, 0) is better than (10, 0) even though it is Pareto inferior. And note that the rule implies that the Pareto inferior distribution is better than the Pareto superior one all the way to the point in which the better off person in the Pareto inferior distribution is only slightly better off than the worst off person. Hence (0.1, 0) is, on the difference summing rule, superior to (10, 0). The difference summing rule of approximation seems to rank some Pareto inferior states above some Pareto superior states, which violates our concern with well-being. These are cases of weakly Pareto inferior distributions being preferred to weakly Pareto superior distributions (in which some persons are better off and no one is worse off). And we think that the concern for well-being implicit in the principle of equality should imply that the weakly Pareto superior distribution should be preferred. Furthermore, we can see that the difference summing rule will also prefer some strongly Pareto inferior distributions to the strongly superior. Just consider the example of (9, 1) and (5, 0). Here the sums of the differences from ideal equality amount to 8 and 5 respectively. These are instances of what we call the *nested levelling down objection*.

These two levelling down worries seem to us to show that we need another principle of approximation to equality if we are to have a satisfactory conception of equality.

The divergence rule of approximation

We think it is intuitively clear that we can construct a rule that satisfies the four constraints stated above and thereby avoids the worries that apply to the difference summing rule. In this section we will construct such a rule.

This problem of having a rule for comparing an actual distribution to an ideal distribution is in fact formally similar to a class of problems outside political philosophy. One example is evaluating how good a given scientific theory is. Perhaps the most famous way to pick between competing theories is Occam's razor. This, however, is not the only way. Another way to evaluate how well a given theory describes the data is to see how close it gets to the data. The data (or observations) is represented as a probability distribution. So is a theory. The theory is then thought of as an arbitrary distribution (because we can have more than one theory

describing the same set of data), whereas the data is thought as the unique ideal distribution. Whichever theory (as a probability distribution) diverges less from the observed data (also represented as a probability distribution) better describes the data. This can be thought of as information divergence or relative entropy. The same technique seems appropriate in the case of ideal political theory when one deals with distributions. Instead of probability distributions, however, a welfare egalitarian will have to evaluate and compare welfare distributions (i.e. distributions of well-being across persons).¹¹

For contrast, consider the utilitarian measure of evaluating states of affairs by tracking total or average utility. A measure of this sort will have nothing to say about the *shape* of the distribution and so will not tell you by how much the actual distribution diverges from the ideal distribution. For example, consider the following two distributions: (5, 5) and (7, 3). If (5, 5) is the ideal distribution, then (5, 5) is clearly closer to the ideal than (7, 3). Average or total utility has no way of tracking this, however. Both have the same average utility and both have the same total utility, yet the shapes of the two distributions are very different. Because egalitarianism is a comparative notion, it must take shape into account in some way. (And so will theories of comparative desert for example.) In fact, going by total or average utility does not single out one unique distribution as the ideal. If the total amount of well-being in the ideal case to be divided up among persons is 10, then both (5, 5) and (7, 3) look identical from the point of view of justice, unless some other principle is brought in to break the tie.¹² Divergence measures, on the other hand, can measure how far a given actual distribution is from *the ideal distribution* (where the ideal distribution is identified by some other prior method) while paying attention to the shape of a distribution when they measure by how much one distribution diverges

¹¹ If we thought of welfare probabilistically, then we can even easier plug it into the developed toolkit for measuring divergence between probability distributions. It seems appropriate for a politician, for example, to think probabilistically when choosing between policies. In this paper, however, we will be dealing with welfare distributions that obtain with certainty and so will not be dealing with probability distributions.

¹² Sidgwick for example uses equality to break ties in such cases, giving preference to the more egalitarian of the two distributions in cases when both distributions have the same total utility. See Henry Sidgwick, *The Methods of Ethics* 7th edition [1907] (Indianapolis: Hackett Publishers, 1981), p. 417.

from another.¹³ This will give due attention to each point in that distribution (i.e. give due consideration to each *person* since each point in the distribution represents the amount of welfare a given person has).

There is, however, more than one possible way to measure divergence between two distributions and which divergence measure is the right one is an open question. One requirement is that the measure satisfies the four constraints we listed above. We find divergence measures attractive not only because they offer a way to measure the injustice of a distribution relative to the ideal distribution (and because they seem to give due attention to each point in the distribution and hence to each person, instead of, say, aggregating the distribution), but also because divergence measures seem automatically to satisfy our four constraints as we will show below in the example below.

One standard measure of divergence between distributions used in information theory is the Kullback–Leibler divergence.¹⁴ This measure, however, was designed to deal with probability distributions and so is not properly suited for welfare distributions. This is because probability distributions always sum up to 1, which means that the sum of probabilities across one distribution will be equal to the sum across any other distribution. For welfare distributions, however, total welfare in one distribution will not necessarily be equal to the total welfare in another distribution. Recall the levelling down objection. In particular, recall the life boat example. The point of the objection is that while it would be nice for everyone to live, regrettably this is not possible. One person must get thrown overboard or everyone dies. It might also be more just if, instead of drawing straws and throwing

¹³ Divergence measures are not necessarily metrics and so do not measure the distance between two distributions because they lack symmetry: by how much distribution S_1 diverges from distribution S_2 is not necessarily equal to how much distribution S_2 diverges from distribution S_1 .

¹⁴ KL-divergence is used on probability distributions, which means that points in the distributions it is applied to add up to 1. KL-divergence is often interpreted as measuring the entropy of one probability distribution relative to another. It can be used, for example, to measure the difference from the data or scientific observations (the ideal distribution) to a scientific theory that tries to account for that data (the actual/ arbitrary distribution). In the case of ideal theories of distributive justice that concern themselves with patterns, e.g. in our case an egalitarian theory, we are trying to do something very similar in spirit: measure the difference from the ideally just distribution (e.g. the ideal egalitarian distribution that maximizes average utility) to the actual distribution (i.e. the current state of affairs).

someone overboard, we could instead subtract the number of years that person had left to live from the total sum of years everyone in the boat has left to live and then divide those years equally. In other words, it might be fairer for everyone in the boat to share the burden equally instead of having one person bear *all* of it. But this isn't possible either. So our only two choices are: everyone drowns or only one person drowns. If we think of the ideally just distribution as one in which everyone lives and, for the case when there are only three people in the boat, represent it as (3.3, 3.3, 3.3), then the case when everyone drowns is (0, 0, 0) and the case when only one person drowns is (5, 5, 0). The latter distribution clearly diverges from the ideal less than the levelled-down distribution. What is even more striking, however, is that the sums across the distributions aren't equal. Total welfare in the levelled-down distribution is 0, whereas in the ideal case it is 10 (and similarly, in the case when only one person drowns, total welfare is 10). This is important because it sets welfare distributions apart from probability distributions and makes standard measures of divergence between probability distributions (e.g. KL-divergence) not suited for our present purposes. Furthermore, normalizing a welfare distribution so it would look like a probability distribution (i.e. so that it would sum up to 1) would miss the whole point we have been labouring to make. If the ideal case in which everyone lives and the levelled down case in which everyone drowns are normalized so that the two distributions both add up to 1, they would look identical! But this, we have argued, is a mistake because what motivates our egalitarian intuitions is a concern with the welfare of each person, whereas normalization and levelling down only pay attention to the shape of the distribution, which violates our constraint that, all other things being equal, pareto superior welfare distributions ought to be preferable to pareto inferior distributions.

Thus, as a working hypothesis, we would instead like to suggest a simpler measure of divergence for the two-person case (which does not demand that distributions be normalized to sum up to 1):

$$\text{Divergence } D(S_{\text{actual}} \| S_{\text{ideal}}) = \left(\frac{1}{x} + \frac{1}{y} \right) - \left(\frac{1}{x^*} + \frac{1}{y^*} \right)$$

between the ideal distribution (x^*, y^*) and the actual distribution (x, y) , where x and y represent portions instead of actual welfare: in other words, x and y are fractions of the total amount of welfare to be

divided up under ideal circumstances.¹⁵ In the life boat example, with two people in the boat, the ideal distribution (5, 5) would be $(5/10, 5/10) = (0.5, 0.5)$ because the total amount of welfare (or maybe years) to be divided up in the ideal circumstances is 10. The case (10, 0) when one person lives and the other drowns would then become (1, 0), and the levelled down case would, of course, remain (0, 0). Formally, $x = A/W$ and $y = B/W$, where W is the amount of total welfare to be divided up in the ideal circumstances, A is the amount of welfare the first person actually gets and B is the amount of welfare the second person actually gets. If the ideal distribution were (6, 6) instead of (5, 5), the numbers in any actual distribution would have to be adjusted accordingly before they could be used in this divergence measure. An actual distribution (5, 5) for example would no longer be $(5/10, 5/10) = (0.5, 0.5)$ when represented as percentages of total welfare, but instead would be represented as $(5/12, 5/12) \approx (0.4, 0.4)$ because total welfare in an ideal distribution (6, 6) is obviously 12 because $6 + 6 = 12$.

This measures how far each person’s well-being is from the ideal. The closer to zero, the better because the ideal distribution will not diverge from itself at all (i.e. will have zero divergence from the ideal). The following table ranks different distributions using this divergence measure, where the ideal distribution $S_{ideal} = (x^*, y^*) = (0.5, 0.5)$. The distributions are listed from most just to least just in table 1:

Distribution (0.5, 0.5) coincides with the ideal distribution (5, 5) in our initial example and so does not diverge from it at all, which is why its divergence value is zero.

Note that (0.5, 0.6) is not possible because it would mean that the total amount of welfare has exceeded 100%. Another way of

¹⁵ The more general form of this divergence measure for the n-person case and which does not presuppose that the ideal distribution is egalitarian is

$$D = \sum_{i=1}^n \frac{x_i^*}{x_i} - \sum_{i=1}^n \frac{x_i^*}{x_i^*} = \sum_{i=1}^n \left(\frac{x_i^*}{x_i} - \frac{x_i^*}{x_i^*} \right) = \sum_{i=1}^n \left(\frac{x_i^*}{x_i} - 1 \right) = \left(\sum_{i=1}^n \frac{x_i^*}{x_i} \right) - n, \text{ such that } \sum_{i=1}^n x_i^* = 1, \text{ but}$$

not necessarily $\sum_{i=1}^n x_i = 1$ since leveling down might be an (inferior) option. The total

amount of well-being in the *ideal* distribution is then normalized to sum up to 1 (albeit not for all actual distributions). Divergence D measures the difference from the ideal welfare distribution $(x_1^*, x_2^*, \dots, x_n^*)$ between n persons to some arbitrary/actual distribution (x_1, x_2, \dots, x_n) , where x_i is the amount of well-being person i actually gets and x_i^* is the amount of well-being person i should get in ideal (i.e. ideally just) circumstances.

Table 1

$S_{\text{actual}} = (x, y)$	$D(S_{\text{actual}} S_{\text{ideal}})$	
(0.5, 0.5)	0.00	← Ideally just distribution.
(0.45, 0.45)	0.44	
(0.4, 0.5)	0.50	
(0.44, 0.44)	0.54	
(0.7, 0.3)	0.76	
(0.4, 0.4)	1.00	
(0.2, 0.2)	6.00	
(0.99999, 0.00001)	999997	
(0.5, 0.00001)	999998	
(0.1, 0.00001)	1000006	
(0.1, 0.00001)	1000096	← Least just distribution of those listed in this table.

stating the same point is that were the corresponding welfare distribution (5, 6) possible, then the ideal welfare distribution would be at least (5.5, 5.5) since the total amount of welfare under circumstance would then be at least $5 + 6 = 11$. If the total amount of well-being in the circumstance turned out to be even greater than 11, say 12, then the ideal egalitarian distribution would be (6, 6) and (5, 6) would indeed be closer to the ideal than would (5, 5).

The distribution (0.99999, 0.00001) is an approximation to (10, 0), which would result in division by zero and so, using our divergence measure, would result in an undefined divergence value.

Let us see how the measure we have discussed satisfies the four constraints on a measure of approximation to ideal equality. We can see from the above table that the ideal equality is the best given the divergence measure since the divergence is 0. And we can also see that the measure prefers some equalities over some inequalities even though they have less average utility such as the case in (0.45, 0.45) or even (0.44, 0.44) over (0.7, 0.3). And the measure prefers equal distributions over unequal distributions that have the same average utility. But the measure also prefers (strongly and weakly) Pareto superior distributions to Pareto inferior distributions.

We are not sure if (0.7, 0.3) getting ranked as more just than (0.4, 0.4) is a troubling result. This might, of course, not be the

right divergence measure for egalitarian justice and so, as we have noted above, we offer it merely as a working hypothesis because we think it provides a good illustration of how one might go about measuring the divergence of an actual distribution from the ideal distribution and thus having a procedure for choosing the least of all evils when comparing and evaluating non-ideal distributions.

Notice that the principle of equality articulated in the above way will give us some practical guidance even in those circumstances where the ideal equality point is infeasible. It does this by telling us what feasible states are the most just approximations of the ideal equality point and thereby telling us what the least unjust outcome is in the circumstances. So the principle of equality we have elaborated does give us practical guidance even when the ideal cannot be achieved.

The strategy accords with the principle of equality defended in this paper because both the ideal equality and the rule of approximation take structural and well-being considerations into account. The principle of equality articulated here ensures that only equality is the ideal distribution; the importance of well-being ensures that the Pareto superior equality is more just than the Pareto inferior equality. The method also prefers Pareto superior inequalities over Pareto inferior equalities. And the method of assessing departures from equality in terms of divergence from the ideal equality seems highly intuitive.

This seems to us to give exactly the right result. The idea is that there is injustice in an unequal condition when there are no relevant differences between the persons between whom the inequality holds. But, there is less injustice in such an unequal condition than in an equal condition that is strongly or even weakly Pareto inferior.

But this shows how one can think that there is something lost and problematic in efficient inequality, and also think that it does not follow that egalitarians are committed to the proposition attributed to them in the levelling down objection. For that objection to work against an egalitarian principle of justice it is not enough that the egalitarian is committed to the injustice of inequality. The levelling down objection applies to equality only if the egalitarian is also committed to the claim that inequality is worse from the point of view of the principle of equality than a Pareto inferior equality. But it is this last claim that the egalitarian need not and indeed ought not be committed to, as we have argued.

Have we simply changed the question? We have responded to the levelling down objection by defending a principle that seems to avoid it. But is this really what egalitarians are after? Or is it a new principle? We want to defend the claim that it is the right principle of equality. Remember that an essential part of the rationale for the principle of equality involves the claim that more substantial good is better than less. The argument for equality involves this essential premise. What would be the point of equality if this were not so? Given the internal connection between equality and the idea that more is better than less when it comes to the thing equalized and given the argument for equality that we have given, it seems reasonable to hold that this is the proper conception of equality.

Cohen and the difference principle

In chapter 4 of his book *Rescuing Justice and Equality* G. A. Cohen argues against the difference principle on the grounds that it allows inequalities on the basis of arbitrary differences among persons. This argument goes beyond his famous arguments against Rawls's inclusion of incentives to the talented as legitimate bases for inequality.¹⁶ He argues that even if it is not possible to help the worst off without inequality the inequality is unjust. In this paper, we believe we have defended a position that accepts Cohen's conclusion because it argues that all inequalities are unjust. But it also qualifies that conclusion in an interesting way, making some sense of the idea that the difference principle is connected with justice. For we have argued that feasible Pareto improvements over feasible equalities are unjust but they are nevertheless more just than the feasible Pareto inferior inequalities. The reason for this, we have argued, is that the feasible Pareto superior inequalities are closer to the ideal equality in the circumstances than the feasible Pareto inferior inequalities.

One way of describing many of the distributions favoured by the difference principle is as a principle of justified injustice. They are unjust for the reasons Cohen describes. They favour some over others for morally arbitrary reasons. But they can be justified nevertheless because they are superior to the feasible equality. But

¹⁶ In chapter 1 of *Rescuing Justice and Equality*.

our point is stronger than this. Our point is also that many of the distributions favoured by the difference principle may be required by principles of justice when we take into account not only the principle of ideal justice but we also take into account principles of how best to approximate justice. On these principles, not only might many distributions favoured by the difference principle be justified; it may also be the case that the distributions favoured by the difference principle are more just than the feasible alternatives. And in this respect many distributions favoured by the difference principle may be required by justice.

We think this addendum to Cohen's position is necessary for the reasons we have stated above. It avoids the levelling down objection while endorsing a pure conception of egalitarian justice. It does so, we believe, in a theoretically satisfying way.

Thomas Christiano
Department of Philosophy
University of Arizona
Tucson, Arizona 85721-0027
thomasc@u.arizona.edu

Will Braynen
Department of Philosophy
University of Arizona
wbraynen@gmail.com