

This is a post-peer-review, pre-copyedit version of an article published in *Synthese*. The final authenticated version is available online at: <http://dx.doi.org/10.1007/s11229-018-1736-5>

Is Maximin Egalitarian?

Jacob Barrett

1. Introduction

According to the Maximin principle of distributive justice, one outcome is more just than another if the worst off individual in the first outcome is better off than the worst off individual in the second. This is often seen as a highly egalitarian principle, and, more specifically, as a highly egalitarian way of balancing a concern with equality against a concern with efficiency. John Rawls, for example, argues that his version of Maximin, the Difference Principle, is “strongly egalitarian” yet “compatible with the principle of [Pareto] efficiency” (1999: 65, 69).¹ Brian Barry explains further that Maximin “picks out the most egalitarian of all the Pareto-optimal arrangements satisfying the requirement that everyone should gain from inequality” (1989: 227). And Russell Hardin contends that, in being the first to suggest a variant of Maximin, Rawls makes “the cleverest move in all of contemporary political philosophy” since he “puts two incommensurable values ([efficiency] and egalitarianism) together” (2003: 105).²

¹ Rawls’s Difference Principle is concerned more precisely with the distribution of primary goods. Despite its fame, it plays a rather limited role in his broader theory: it applies only when evaluating the basic structure of society (rather than when evaluating particular outcomes), and, even then, it is subordinate to his other two principles of justice. In both of these ways, it differs from Maximin as I define it here, and as it is commonly defined in the philosophical and economic literature on egalitarianism.

² Hardin says “welfare” instead of “efficiency,” but it is clear from the context that he has efficiency in the production of welfare (or, better, Rawlsian primary goods) in mind.

On the other hand, some theorists deny that Maximin shows any concern for equality at all. As Douglas McKerlie puts the challenge:

Any argument from equality to [Maximin] faces the same problem. If we care about equality it is plausible to think that we object to inequality between any two groups for its own sake. How can we get from this starting-point to the conclusion that we should assess inequality only in terms of its effect on the worst off? (1994: 33)

To this, egalitarian proponents of Maximin do, however, have a powerful reply: that it is possible to derive Maximin from surprisingly weak conditions relating efficiency and equality to justice. Formal derivations of this sort have been put forward by a number of theorists, most recently by Bertil Tungodden and Peter Vallentyne (2005). Even if these derivations do not show that Maximin is the best way to balance equality against efficiency, they do appear to show that it is one such way, thereby providing a response to McKerlie's challenge.

In this paper, I argue that, despite appearances, this reply fails. I begin by reconstructing Tungodden and Vallentyne's recent proof of Maximin in a way that makes explicit how it aims to show that Maximin balances a concern with efficiency against a concern with equality. Then, I defend a condition on any reasonable conception of equality that I call "Maximin Failure," and demonstrate that the truth of Maximin Failure refutes Tungodden and Vallentyne's derivation of Maximin by transforming it into an impossibility result. Finally, I show that the truth of Maximin Failure also undermines every other egalitarian proof of Maximin in the literature, since each of these proofs relies on conditions that, when conjoined with Maximin Failure, also yield an impossibility result.

If successful, my argument shows that even if it possible to defend Maximin on other (say, prioritarian) grounds, there is no known way to defend it on the grounds that it balances a concern with equality against a concern for efficiency. But since Tungodden and

Vallentyne prove that a set of rather weak and plausible egalitarian-friendly conditions entail Maximin, and I demonstrate that these same conditions conjoined with any reasonable conception of equality lead to an impossible result, my argument also implies that—unless they can resist my claim that all reasonable conceptions of equality satisfy Maximin Failure—egalitarians must reject some other condition in 'Tungodden and Vallentyne' proof. In the conclusion, I briefly consider what options this leaves open to egalitarian theorists of distributive justice, as well as the extent to which my argument generalizes to theories of distributive justice that deem further considerations beyond equality and efficiency—such as rights and desert—to be relevant to justice, too.

2. Tungodden and Vallentyne's Derivation of Maximin

Following Tungodden and Vallentyne (2005), let us concern ourselves with outcomes in which generic, homogenous, interpersonally comparable “benefits” are distributed among some group of individuals.³ These benefits may be interpreted as units of welfare, income,

³ More specifically, I will assume that the benefits in question are *ratio* comparable. This allows us to make *level* comparisons of the form “*A* has more benefits than, less benefits than, or just as many benefits as *B*,” *unit* comparisons of the form “*A* gained or lost *n* times as many benefits as *B* gained or lost,” and *ratio* comparisons of the form “*A* has *n* times as many benefits as *B*.” Assuming the first two sorts of comparability is warranted in the present context because, as we will see in the next section, all conceptions of equality require us to make them (see fn. 13). Assuming the third is furthermore warranted because, as we will also see, doing so allows me to capture a wider range of conceptions of equality, and only makes things harder for me anyway. For helpful discussions of these different sorts of interpersonal comparability, and more generally of the social choice-inspired framework

primary goods, or so on, but for ease of exposition, I adopt welfarist language and refer to an individual who has more benefits than another as “better off” than the other. This allows us to define Maximin as follows:

Maximin: for all outcomes x, y , if the worst off individual in x is better off than the worst off individual in y , then x is more just than y ⁴

Maximin entails, for example, that $\langle 3, 6 \rangle$ —an ordered pair which represents an outcome in which one individual has a benefit level of 3, and a second a benefit level of 6—is more just than $\langle 2, 2 \rangle$ and $\langle 2, 10 \rangle$. Since $\langle 3, 6 \rangle$ is less equal than $\langle 2, 2 \rangle$ (more on which shortly), Maximin clearly does not track equality by itself, but it might nonetheless be thought somehow to balance equality and efficiency. This is what Tungodden and Vallentyne try to show, and they do so by relying on just a few more conditions.

The first two conditions are these:

Domain Richness: the set of possible outcomes includes all logically possible benefit distributions

Quasi-Transitivity: for all outcomes x, y , and z , if x is more just than y , and y is more just than z , then x is more just than z

employed throughout, see Sen (1977), d’Aspremont and Gevers (2002), and Bossert and Pfingstein (2004).

⁴ Maximin is sometimes defined to include the further claim that if the worst off individuals in x and y are equally well off, then x is equally just as y . I omit this claim here because I wish to show that one cannot even justify the weaker version of Maximin that I have defined in the main text on egalitarian grounds, and because some of the proofs I consider only concern it.

Domain Richness is less controversial than it might initially seem. It does not say that all logically possible benefit distributions are feasible, nor that every logically possible benefit distribution can be ranked *vis-à-vis* every other logically possible benefit distribution across any particular normative dimension (such as justice, equality, or efficiency). It says only that all logically possible benefit distributions are admissible candidates for ranking across these dimensions, or, in other words, that when a principle like Maximin quantifies over “all outcomes,” the set of outcomes it quantifies over includes all logically possible benefit distributions. Quasi-Transitivity—a weakening of the more familiar condition of Transitivity (discussed in section 4 below), which adds to Quasi-Transitivity the claim that if x is equally just as y , and y equally just as z , then x is equally just as z —is perhaps more controversial, but requires less explanation. It entails, for example, that if $\langle 3, 6 \rangle$ is more just than $\langle 3, 5 \rangle$, and $\langle 3, 5 \rangle$ is more just than $\langle 2, 2 \rangle$, then $\langle 3, 6 \rangle$ must be more just than $\langle 2, 2 \rangle$, too.

Now we come to some more substantive conditions. First, an efficiency condition:

Pareto: for all outcomes x, y , if everyone is at least as well off in x as in y , and one person is furthermore better off in x than in y , then x is more just than y ⁵

Pareto says that if at least one person benefits and nobody is made worse off, then this improves justice. So, for example, it entails that $\langle 3, 6 \rangle$ is more just than $\langle 3, 3 \rangle$, but it does not entail anything about the relative justice of $\langle 3, 3 \rangle$ and $\langle 2, 6 \rangle$. Following convention, we may put this by saying that $\langle 3, 6 \rangle$ is “Pareto superior” to $\langle 3, 3 \rangle$ and that $\langle 3, 3 \rangle$ and

⁵ This condition is often called “Strong Pareto,” in contrast with “Weak Pareto,” which states that x is more just (or more efficient) than y if *everyone* is better off in x than y . For ease of exposition, I drop the “Strong” here. Actually, both Strong and Weak Pareto are usually defined to include the further claim that if everyone is just as well off in x as in y , then x is equally just as y . I drop that here, too.

$\langle 2, 6 \rangle$ are “Pareto incomparable.” This allows us to define a slightly stronger efficiency condition:

Anonymous Pareto: for all outcomes x, y , if x is Pareto superior to y or to some permutation of the benefit distribution in y , then x is more just than y ⁶

Anonymous Pareto entails Pareto, but adds that even when two benefit distributions are Pareto incomparable, the first is more just than the second if it is Pareto superior to some permutation of the second, where two distributions are permutations of each other when they differ from each other only in which benefit levels are assigned to which individuals (as $\langle 1, 2, 3 \rangle$ differs from $\langle 2, 3, 1 \rangle$). Thus, whereas Pareto entails that $\langle 6, 3 \rangle$ is more just than $\langle 5, 3 \rangle$ but that it is incomparable with $\langle 3, 5 \rangle$, Anonymous Pareto entails further that $\langle 6, 3 \rangle$ is more just than $\langle 3, 5 \rangle$, because $\langle 6, 3 \rangle$ is Pareto superior to $\langle 5, 3 \rangle$ and $\langle 3, 5 \rangle$ is a permutation of $\langle 5, 3 \rangle$. On the other hand, Anonymous Pareto still says that $\langle 3, 7 \rangle$ and $\langle 4, 5 \rangle$ are incomparable because $\langle 3, 7 \rangle$ is neither Pareto superior to $\langle 4, 5 \rangle$ nor to $\langle 5, 4 \rangle$.

I call Pareto and Anonymous Pareto “efficiency conditions” because imposing either on the justice relation amounts to giving absolute priority to a certain criterion of efficiency over all other considerations, at least in cases where that criterion judges one outcome to be superior to another. Each can be factored into this condition:

Efficiency Trumps: for all outcomes, x, y , if x is more efficient than y , then x is more just than y

in conjunction with one of the following conditions:

Pareto Efficiency: for all outcomes x, y , if everyone is at least as well off in x as in y , and one person is furthermore better off in x than in y , then x is more efficient than y

⁶ This condition derives from Patrick Suppes (1966), and is often referred to as “Suppes’ Grading Principle.”

Anonymous Pareto Efficiency: for all outcomes x, y , if x is Pareto superior to y or to some permutation of the benefit distribution in y , then x is more efficient than y

Thus, Pareto says that, in determining the justice of outcomes, efficiency, interpreted as Pareto Efficiency, trumps all other considerations. And Anonymous Pareto says the same where efficiency is interpreted as Anonymous Pareto Efficiency. Both do leave open, however, that in cases where two outcomes are incomparable with respect to Pareto or Anonymous Pareto, other, non-efficiency conditions may be relevant to justice. This is where equality enters the picture.

Given that one endorses either Pareto or Anonymous Pareto, the strongest role one may give to equality is respectively limited to one of the following conditions:

Moderate Egalitarianism: for all outcomes x, y such that x and y are Pareto incomparable, if x is more equal than y , then x is more just than y

Weak Egalitarianism: for all outcomes x, y such that x and y are Anonymous Pareto incomparable, if x is more equal than y , then x is more just than y

Moderate Egalitarianism entails Weak Egalitarianism. Both say that, when x and y are Anonymous Pareto incomparable, considerations of equality trump all other considerations in determining which is more just. Moderate Egalitarianism adds that the same is true when x and y are merely Pareto incomparable. So each entails, for example, that $\langle 4, 4, 4 \rangle$ is more just than $\langle 5, 4, 3 \rangle$, since the two are Anonymous Pareto incomparable (and therefore Pareto incomparable) and the former is more equal. But while Weak Egalitarianism entails nothing about the relative justice of $\langle 4, 10 \rangle$ and the Pareto incomparable (but not Anonymous Pareto incomparable) $\langle 5, 4 \rangle$, if we assume that $\langle 5, 4 \rangle$ is more equal than $\langle 4, 10 \rangle$, then Moderate Egalitarianism entails that $\langle 5, 4 \rangle$ is more just.

The final condition, which has already been implicit in what we have said so far, concerns the equality rather than the justice or efficiency relation:

Perfect Equality: for all outcomes x, y if everyone is just as well off as each other in x , but the same is not true in y , then x is more equal than y

Perfect Equality is self-explanatory. It entails, for example, our earlier judgments that $\langle 2, 2 \rangle$ is more equal than $\langle 3, 6 \rangle$ and that $\langle 4, 4, 4 \rangle$ is more equal than $\langle 5, 4, 3 \rangle$. This is entirely uncontroversial. Of course, it would be very controversial to claim that Perfect Equality is all that we can say about equality, and later we will consider some further conditions that should be imposed on the equality relation as well. But the strength of Tungodden and Vallentyne's proof is that it does not rely on any of these further conditions; it requires only that, whatever else is true of equality, perfectly equal outcomes are judged more equal than imperfectly equal ones.

We are now ready to provide a statement of that proof:

Result 1: Domain Richness, Quasi-Transitivity, Pareto, Weak Egalitarianism, and Perfect Equality entail Maximin

Proof: see Tungodden and Vallentyne's (2005) "Result 3"⁷

This is a very powerful result. It shows that it is possible to derive Maximin from weak and plausible conditions relating equality and efficiency to justice. In particular, it appears to vindicate something very close to Barry's claim that Maximin "picks out the most egalitarian of all the Pareto-optimal arrangements satisfying the requirement that everyone should gain

⁷ Tungodden and Vallentyne appeal to a further condition called "Benefitism" in their proof, but I have omitted it here since the proof does not require it. They also prove some related results that rely on weaker consistency conditions than Quasi-Transitivity that I put aside for now.

from inequality” (1989: 227).⁸ After all, if “picking out” an outcome is judging it most just, then the conjunction of Pareto and Moderate Egalitarianism instruct us to do just what Barry claims. First, Pareto requires us always to pick from the set of Pareto optimal outcomes (those outcomes that are either Pareto superior to or Pareto incomparable with all other outcomes). Second, Moderate Egalitarianism requires us to pick whichever of these Pareto optimal outcomes is most equal. The only difference between Barry’s interpretation and Tungodden and Vallentyne’s, then, is that the latter rely on Weak rather than Moderate Egalitarianism. Otherwise, Tungodden and Vallentyne just make explicit the other assumptions (Domain Richness, Quasi-Transitivity, and Perfect Equality) that are needed to derive Maximin in this way.

In fact, however, we should see Tungodden and Vallentyne as providing a refinement of rather than a deviation from Barry’s interpretation of Maximin. For, in the first place, Moderate Egalitarianism entails Weak Egalitarianism, so the proof would still go through if we replaced the latter with the former. And in the second, Tungodden and Vallentyne have a good reason for opting for the weaker condition: they prove another result showing that if one accepts the rest of the conditions in Result 1 but replaces Weak Egalitarianism with Moderate Egalitarianism (as Barry might suggest), then, if one imposes another extremely plausible condition on the equality relation called “Strong Conditional Contracting Extremes,” one is led instead to an impossibility result (2005: Result 2). Their results therefore show that if one gives the strongest role possible to equality that is compatible with Pareto, one must accept Maximin. McKerlie’s objection therefore appears to be defused: Maximin is a “strongly egalitarian” position after all (Rawls 1999: 65). It puts

⁸ I thank an anonymous reviewer for urging me to clarify the relation between Tungodden and Vallentyne’s proof and Barry’s interpretation of Maximin.

the “two incommensurable values” of efficiency and equality together (Hardin 2003: 105)—not by commensurating or weighing them against each other, but by giving absolute priority to the value of (Pareto) efficiency over the value of equality, and absolute priority to the value of equality over all other considerations in cases where two outcomes are (Anonymous) Pareto incomparable.

3. The Failure of Maximin

To begin to see where the trouble with Tungodden and Vallentyne’s argument lies, consider the pair of Anonymous Pareto incomparable outcomes $\langle 1.99, 2.01, 2.01, 2.01, 2.01 \rangle$ and $\langle 2, 2, 100, 1000, 10000 \rangle$. Perfect Equality is silent on this case, but the former (which hardly departs from perfect equality) is intuitively much more equal than the latter (which greatly departs from it), from which it follows by Weak Egalitarianism that the former is more just than the latter. But since the worst off individual in the latter is better off than the worst off individual in the former, it follows by Maximin that the former is instead *less* just than the latter. So if $\langle 1.99, 2.01, 2.01, 2.01, 2.01 \rangle$ is indeed more equal than $\langle 2, 2, 100, 1000, 10000 \rangle$, then Weak Egalitarianism and Maximin are inconsistent: Weak Egalitarianism judges $\langle 1.99, 2.01, 2.01, 2.01, 2.01 \rangle$ to be more just than $\langle 2, 2, 100, 1000, 10000 \rangle$, while Maximin judges the reverse. In order to preserve their consistency, such that one may derive Maximin from a set of conditions including Weak Egalitarianism, one therefore must deny that $\langle 1.99, 2.01, 2.01, 2.01, 2.01 \rangle$ is more equal than $\langle 2, 2, 100, 100, 1000 \rangle$. But this is intuitively implausible—or so, I suspect, many would think.

The trouble does not just end with intuitive implausibility, however: it ends in an impossibility result. We will get there in two steps. First, what the above pair of outcomes suggests is that any reasonable conception of equality should, at least intuitively, satisfy the following condition (which, it is important to note, applies to the equality rather than to the

justice relation):

Maximin Failure: there are some outcomes x, y , such that x and y are Anonymous Pareto incomparable, the worst off individual in x is better off than the worst off individual in y , but y is more equal than x

And to generalize the argument in the last paragraph, if this extremely weak condition—which just says that there is *some* case where even though two outcomes are Anonymous Pareto incomparable and Maximin prefers the first to the second, the second is more equal than the first—holds, then we have the first step in our impossibility result:

Result 2: It is impossible to jointly satisfy Weak Egalitarianism, Maximin, and Maximin Failure

Proof: By Maximin Failure, there is some x and some y such that x and y are Anonymous Pareto incomparable, the worst off individual in x is better off than the worst off individual in y , and y is more equal than x . Thus, by Weak Egalitarianism, y is more just than x . But since the worst off individual in x is better off than the worst off individual in y , it follows by Maximin that x is more just than y . This is a contradiction, so the result follows.

Although this proof is trivial, the result is not: it shows that if Maximin Failure holds, then Tungodden and Vallentyne's derivation of Maximin fails, since it relies on a condition that is itself inconsistent with Maximin. In other words, Result 2 says that if we add Maximin Failure to the conditions of Result 1, we transform it into an impossibility result. Tungodden and Vallentyne's proof of Maximin can therefore only go through if they deny Maximin Failure.

But why shouldn't they simply deny Maximin Failure? So far, we have seen that Maximin Failure follows from certain intuitive examples, for example, from the judgment that $\langle 1.99, 2.01, 2.01, 2.01, 2.01 \rangle$ is more equal than $\langle 2, 2, 100, 1000, 10000 \rangle$, or, to add another, that $\langle 8, 10, 10, 10 \rangle$ is more equal than $\langle 9, 99, 999, 999 \rangle$. But perhaps intuitive

examples are not enough; perhaps we need more principled grounds for claiming that all reasonable conceptions of equality satisfy Maximin Failure. The second step of our argument will be to provide these grounds. To do so, we will turn to the literature on axiomatic approaches to equality. Although this literature is concerned primarily with equality in the distribution of income, the conditions it invokes may be generalized to a concern with benefits more generally (where income is one possible type of benefit). I propose to do just that.

In the literature just mentioned, there is widespread agreement that any reasonable conception of equality must satisfy a number of standard conditions. Most prominently, they must satisfy the following condition:

Pigou-Dalton: for all outcomes x, y , any two individuals A, B , and any number $r (r > 0)$, if A is at least well off as B in both x and y , and the only difference between x and y is that A 's benefit level is r greater in y than in x while B 's benefit level is r greater in x than in y , then x is more equal than y ⁹

Pigou-Dalton is a weak condition, which says that if (i) A “transfers” some benefits to B , such that A loses by exactly as much as B gains, (ii) A is better off than B before the transfer, and remains at least as well off as B after, and (iii) nobody else is affected, then the transfer improves equality. For example, it says that $\langle 4, 6 \rangle$ is more equal than $\langle 3, 7 \rangle$, and that $\langle 3, 5, 7 \rangle$ is more equal than $\langle 2, 5, 8 \rangle$. Pigou-Dalton is therefore equivalent to the highly compelling condition that if some fixed quantity of benefits can be added to either a better or worse off person (without switching their rank), then giving it to the worse off person bring about a more equal arrangement than giving it to the better off one.

⁹ Pigou-Dalton, also known as the “Transfer Axiom,” was first suggested by Pigou (1912) and then popularized by Dalton (1920).

Pigou-Dalton is the cornerstone of the contemporary economic literature on income inequality. It is satisfied by all standard equality metrics (such as the Gini, Atkinson, and Theil indices), and commands nearly universal assent. Indeed, the popular “intersection” approach to equality—which states that one distribution is more equal than another at least when all reasonable conceptions of equality agree that it is—always proceeds under the assumption that each of the conceptions in this intersection satisfies Pigou-Dalton. There are overlapping grounds for holding this view, but it would take me too far away from my present purposes to go through them all here.¹⁰ Instead, I offer the following alternative condition to those who would rather reject Pigou-Dalton than accept my conclusions:

Not Just The Bottom Two: there are some outcomes x, y , such that the two worst off members of x are equally well off as each other, and the two worst off members of y are not, but y is more equal than x .

Although Not Just the Bottom Two has not been explored in the literature, it is hard to see how it could generate any controversy. To deny it is to claim that, no matter how many

¹⁰ For a helpful overview of this literature, as well as a seminal statement of the intersection approach, see Sen (1997). Note that both Tungodden (2003) and Vallentyne (2009) themselves accept Pigou-Dalton, although they take different attitudes toward it: whereas Vallentyne holds that Pigou-Dalton follows directly from the concept of equality, Tungodden seems to think that it stands in need of more substantive justification (which he takes to be available). This same difference is found throughout the rest of the literature. For example, Broome’s (2015) view is closer to Vallentyne’s, while Kolm’s (1999) is closer to Tungodden’s. Finally, for a criticism of the intersection approach, see Temkin (1993). There, Temkin criticizes Pigou-Dalton as well, but he later retracts this criticism on the grounds that it relies on a mistaken interpretation of it (Temkin 2003).

other people there are and how their benefits are distributed, the fact that the two worst off members in one distribution are equal entails that no other distribution that does not have this same property can be more equal than it—that, for example, $\langle 2, 2.01, \dots \rangle$ cannot be more equal than $\langle 2, 2 \dots \rangle$ no matter how we fill in the rest of either distribution. Not Just The Bottom Two therefore captures a weakened version of McKerlie’s claim that “if we care about equality it is plausible to think that we object to inequality between any two groups” (1994: 32-33)—namely, that we do not object to inequality only between the two very worst off individuals in a population. For this reason, it seems safe to assume that any reasonable conception of equality must satisfy at least one of Pigou-Dalton and Not Just The Bottom Two. The former has the virtue of being both intuitive and widely accepted by experts on the formal properties of equality; the latter has the virtue of being even more intuitively obvious than the first.

The next condition we need is:

Equality Transitivity: for all outcomes, x, y, z , if x is at least as equal as y , and y is at least as equal as z , then x is at least as equal as z

Equality Transitivity is a familiar condition, and within the literature on economic inequality it is usually taken to be an unquestionable assumption (e.g., Sen 1997). Nonetheless, some philosophers have recently questioned it on the basis of certain alleged counter-examples.¹¹ These counter-examples are all very contentious, and moreover, they all involve the comparisons of outcomes that differ in population size, whereas a restricted version of Equality Transitivity that holds only in cases where we are comparing the equality of distributions with fixed populations would suffice for my purposes:

¹¹ The case against Equality Transitivity has been made most prominently by Temkin in a series of publications, culminating in Temkin (2011).

Restricted Equality Transitivity: for all outcomes, x, y, z such that x, y , and z all contain the same population of individuals, if x is at least as equal as y , and y is at least as equal as z , then x is at least as equal as z

To avoid controversy, I will therefore assume that all reasonable conceptions of equality satisfy only Restricted Equality Transitivity, rather than its unrestricted version.

With these general conditions on all reasonable conceptions of equality stated, we are ready to turn to competing conceptions of equality. Such conceptions come in three types: they may be absolute, relative, and intermediate.¹² To understand the difference between these, consider the outcomes $\langle 3, 3 \rangle$, $\langle 1, 5 \rangle$, and $\langle 10, 20 \rangle$. All conceptions of equality agree that the first is more equal than the latter two, but they differ over whether the second or third is more equal, because they disagree about how to compare the size of the gaps between individuals. Absolute conceptions are concerned with “absolute gaps,” or with the absolute differences between individuals’ benefit levels: they judge that the $\langle 1, 5 \rangle$ is more equal than $\langle 10, 20 \rangle$ because in the former one person has 4 more than the other, while in the later one person has 10 more than the other. Relative conceptions are concerned with “relative gaps,” or with the proportionate differences between individuals’ benefit levels: they judge that $\langle 10, 20 \rangle$ is more equal than $\langle 1, 5 \rangle$ because in the former one person has twice as much as the other, while in the latter one person has five times as much as the other. Finally, intermediate conceptions are concerned both with absolute gaps and relative gaps: they are neither completely absolute, nor completely relative. So different intermediate conceptions will disagree about whether $\langle 1, 5 \rangle$ or $\langle 10, 20 \rangle$ is more equal, depending on

¹² These distinctions were first drawn by Kolm (1976), and have since become standard fare. For helpful discussion, see Vallentyne (2000).

the relative importance they assign to absolute and relative equality.¹³

Absolute conceptions of equality are standardly defined as those that satisfy the following property:

Translation Invariance: for all outcomes x, y , if y is a translation of x (that is, y can be obtained from x by adding to each individual's benefit level the same number r), then x is just as equal as y

For example, they judge that $\langle 1, 2, 3 \rangle$ is just as equal as $\langle 2, 3, 4 \rangle$, or, more generally, that $\langle 1, 2, 3 \rangle$ is just as equal as $\langle 1+r, 2+r, 3+r \rangle$. This follows from an absolute conception of equality because in each distribution the absolute gaps between each individual remain the same: the third individual has 1 (or r) more than the second, who has 1 (or r) more than the first. In other words, Translation Invariance simply formalizes the thought that since adding the same quantity of benefits to each individual's benefit level can never affect the size of absolute gaps between individual, it can never affect equality—and this is exactly what absolute conceptions of equality claim.

On the other hand, relative conceptions of equality (and, as it happens, the majority of equality metrics) are standardly held to satisfy the following property:

Scale Invariance: for all outcomes x, y , if y is a scale of x (that is, y may be obtained from x by multiplying each individual's benefit level by the same number r ($r > 0$)), then x is just as equal as y

For example, they judge that $\langle 1, 2, 4 \rangle$ is just as equal as $\langle 2, 4, 8 \rangle$, or, more generally that

¹³ Judgments about absolute gaps involve level and unit interpersonal comparisons of benefits; judgments about relative gaps furthermore involve ratio interpersonal comparisons. Thus, absolute conceptions of equality require these first two sorts of interpersonal comparability, and relative and intermediate conceptions require all three. See fn. 3 above.

$\langle 1, 2, 3 \rangle$ is just as equal as $\langle 1r, 2r, 4r \rangle$. This follows from a relative conception of equality because multiplying each individual's benefit level by a positive number does not change the relative gaps between individuals: in each distribution the third individual has two (or r) times as much as the second, who has two (or r) times as much as the first. Again, it holds as a perfectly general matter that multiplying each individual's benefit level by the same positive number cannot have any effect on the size of the relative gaps between individuals, and therefore cannot have any effect on relative equality.

Finally, intermediate conceptions of equality are concerned with both absolute and relative inequality, so they do not satisfy either Translation or Scale Invariance. Instead, they are typically held to satisfy the following property:

Compromise Condition: for all outcomes x, y such that neither is perfectly equal, (i) if y is an increasing translation of x (that is, y can be obtained from x by adding to each individual's benefit level the same number r ($r > 0$)), then y is more equal than x , and (ii) if y is an increasing scale of x (that is, y may be obtained from x by multiplying each individual's benefit level by the same number r ($r > 1$)), then x is more equal than y ¹⁴

Thus, Compromise Condition says that if we take $\langle 1, 2, 3 \rangle$ and add to each individual's benefit level some positive number, say, 2, then this improves equality, and if we multiply each individual's benefit level by some number greater than 1, say, 2, this decreases it. So it

¹⁴ Although there is disagreement about how best to understand intermediate conceptions of equality, it is common ground that such conceptions must at least satisfy Compromise Condition. For example, Compromise Condition was already suggested by Kolm (1976) when he introduced the distinction between the three conceptions, and Bossert and Pfingstein (1990) take satisfying Compromise Condition to be a success condition on defining an intermediate conception of equality.

entails that $\langle 1, 2, 3 \rangle$ is less equal than $\langle 3, 4, 5 \rangle$ but more equal than $\langle 2, 4, 6 \rangle$. This is because part (i) of Compromise Condition—which entails that $\langle 1, 2, 3 \rangle$ is less equal than $\langle 3, 4, 5 \rangle$ —describes an unambiguous case of relative equality increasing (the relative gaps between every pair of individuals diminish) without absolute equality changing (the absolute gaps between every pair of individuals remains the same), and part (ii) of Compromise Condition—which entails that $\langle 1, 2, 3 \rangle$ is more equal than $\langle 2, 4, 6 \rangle$ —describes an unambiguous case of absolute equality decreasing (the absolute gaps between every pair of individuals grow) without relative equality changing (the relative gaps between every pair of individuals remain the same).¹⁵ Thus, any intermediate conception of equality—any conception that gives positive concern to both relative and absolute equality, and that therefore gives sole concern to either relative or absolute equality in cases where the other is silent—must satisfy Compromise Condition.

In the literature, there is no fourth alternative to absolute, relative, and intermediate conceptions of equality, and there does not seem to be any conceptual space for such a view. Thus, I will define a “reasonable conception of equality” as one that satisfies Restricted Equality Transitivity, at least one of Pigou-Dalton and Not the Bottom Two, as well as at least one of Translation Invariance, Scale Invariance, or Compromise Condition. We may now prove the following result:

Result 3: Given Domain Richness, all reasonable conceptions of equality—that is, all conceptions of equality that satisfy Restricted Equality Transitivity, at least one of Pigou-

¹⁵ This explanation of Compromise Conditions relies on the uncontroversial assumption that an increase in the absolute or relative gaps between every pair of individuals respectively increases absolute or relative equality. Kolm (1999) suggests a similar condition.

Dalton and Not Just The Bottom Two, and at least one Translation Invariance, Scale Invariance, and Compromise Condition—satisfy Maximin Failure.

Proof: See the Appendix.

This result simply states that all reasonable conceptions of equality, as just defined, satisfy Maximin Failure. The significance of this is obvious: since Result 2 claims that it is impossible to jointly satisfy Maximin Failure, Weak Egalitarianism, and Maximin, Results 2 and 3 combine to entail that, given any reasonable conception of equality, it is impossible to jointly satisfy Weak Egalitarianism and Maximin. And from this, it follows in turn that Tungodden and Vallentyne's derivation of Maximin is unsuccessful. One cannot derive Maximin from a set of conditions that includes Weak Egalitarianism, because, given any reasonable conception of equality (at least as I have argued we should characterize such conceptions), Weak Egalitarianism and Maximin are jointly inconsistent.

4. The Failure Extended

Tungodden and Vallentyne are not the only theorists who attempt to derive Maximin from conditions that relate efficiency and equality to justice, so the argument just provided might seem to have a rather limited reach. In this section, I will argue to the contrary that Maximin Failure also makes trouble for every other derivation of Maximin in the literature, since each either relies on Weak Egalitarianism (or a variant of it that is similarly inconsistent with Maximin and Maximin Failure), or else turns out not to be an egalitarian derivation of Maximin at all, in the sense that none of the conditions it relies on relate the equality relation to the justice relation. Before we get to these derivations, however, we will need to define two conditions, and prove one preliminary result.

The conditions are:

Transitivity: for all outcomes x, y , if x is at least as just as y , and y is at least as just as z , then x is at least as just as z

Weaker Egalitarianism: for all outcomes x, y such that x and y are Anonymous Pareto incomparable, if x is more equal than y , then x is at least as just as y

Transitivity is a stronger consistency condition than Quasi-Transitivity, which differs from it insofar as it applies not only to the “more just than” relation but also to the “at least as just as” (or to the “more just than or equally just as”) relation. I mention it here only because the following proofs all rely on it. Weaker Egalitarianism is like Weak Egalitarianism, but weakened to say that when x and y are Anonymous Pareto Incomparable, and x is more equal than y , then x is *at least as* just as rather than *more* just than y . It, too, is incompatible with Maximin and Maximin Failure, as the following result shows:

Result 4: It is impossible to jointly satisfy Weaker Egalitarianism, Maximin, and Maximin Failure

Proof: By Maximin Failure, there is some x and some y such that x and y are Anonymous Pareto incomparable, the worst off individual in x is better off than the worst off individual in y , and y is more equal than x . Thus, by Weaker Egalitarianism, y is at least as just as x . But since the worst off individual in x is better off than the worst off individual in y , it follows by Maximin that x is more just than y . This is a contradiction, so the result follows.

With these preliminaries out of the way, we may now to these alternative proofs.

4.1 Tungodden (i) and (ii)

In a paper that is a precursor to his paper with Vallentyne that I have been discussing so far, Tungodden (2000) provides three proofs of Maximin. The first two are these. First, Tungodden shows that Domain Richness, Transitivity, Anonymous Pareto, Weaker Egalitarianism, and another plausible condition on the equality relation called Strong

Conditional Contracting Extremes entail Maximin (2000: Theorem 1).¹⁶ And, second, he shows that if we replace Strong Conditional Contracting Extremes with another condition on the equality relation called Strong Absolute Priority of Those Below The Mean, this furthermore entails Leximin: a view that implies Maximin but furthermore that when the worst off individuals in two outcomes are just as well off as each other, then the more just outcome is that in which the second worst off individual is better off (such that, for example, $\langle 1, 5 \rangle$ is more just than $\langle 1, 4 \rangle$) or, if the second worst off individuals are also tied, that the more just outcome is the one in which the third worst off individual is better off (such that, for example $\langle 1, 4, 6 \rangle$ is more just than $\langle 1, 4, 5 \rangle$), and so on (2000: Theorem 5).¹⁷ However, since both of these proofs assume Weaker Egalitarianism, it follows by Result 4 that, given the truth of Maximin Failure, the conditions that the proofs rely on are

¹⁶ For a statement of Strong Conditional Contracting Extremes, see Tungodden (2000: 233). Technically, Tungodden does not identify Weaker Egalitarianism as its own condition when giving his proof. But after defining Strong Conditional Contracting Extremes over the equality relation, he then extends it to hold over the justice relation by suggesting that “we assign absolute priority to equality promotion in distributive conflicts,” where, for Tungodden, a “distributive conflict” occurs whenever two alternatives are Anonymous Pareto incomparable (232). So Tungodden implicitly relies on Weaker Egalitarianism when giving his proof, and he relies on it explicitly in later papers, such as Tungodden (2003), when he discusses it. Vallentyne (2009) provides the same interpretation of Tungodden’s proof.

¹⁷ For a statement of Strong Absolute Priority of Those Below the Mean, see Tungodden (2000: 238).

inconsistent with Maximin. Maximin Failure therefore transforms each of these proofs into an impossibility result.

4.2 Hammond and Tungodden (iii)

The third proof Tungodden develops is closely related to an earlier, famous proof of Leximin (and therefore Maximin) due to Peter Hammond (1976: Theorem 7.2). Since Tungodden's proof builds on Hammond's, it is worth looking at Hammond's first.

Hammond's key condition is:

Hammond Equity: for all outcomes x, y , and any two individuals A, B , if A is at least as well off as B in both x and y , and the only difference between x and y is that A is better off in y than in x while B is better off in x than in y , then x is at least as just as y

Hammond Equity is similar but in two ways importantly distinct from Pigou-Dalton. First, unlike Pigou-Dalton transfers, Hammond transfers (as I will call them) do not require that the transfer has no effect on the sum of benefits: they allow that A may lose by more than B gains, or that B may gain by more than A loses. Second, Hammond Equity says that if y can be reached from x by a Hammond transfer—if, for example, x is $\langle 1, 4, 10 \rangle$ and y is $\langle 1, 5, 6 \rangle$ —then y is at least just as x , not merely that it is more equal. And Hammond proves that in the presence of Domain Richness, Transitivity, and Anonymous Strong Pareto, Hammond Equity entails Leximin.

This is an interesting result, but it is not obvious how to square it with our discussion so far. It is clear that Hammond Equity is unconcerned with efficiency, but its relation to equality is less straightforward than it might initially seem. The most obvious way to explain this relation would be to point out that Hammond Equity follows from Weaker Egalitarianism in conjunction with the claim that Hammond transfers improve equality. This would entail Hammond Equity, but given that we would then be appealing to Weaker

Egalitarianism, the truth of Maximin Failure would once again (by Result 4) lead to an impossibility result. A different justification of Hammond Equity is therefore needed before we can claim that Hammond's proof justifies Maximin on egalitarian grounds.

Tungodden's third proof is, in effect, an attempt to provide such a justification (2000: Theorem 3). Tungodden shows that Weaker Egalitarianism, another plausible condition on the equality relation called "Weak Conditional Contracting Extremes,"¹⁸ and a condition called "Strong Separability"—which states that adding or removing individuals who are indifferent between outcomes cannot affect the ranking of those outcomes (such that, for example $\langle 3, 6, 15, 15 \rangle$ is more just than $\langle 1, 10, 15, 15 \rangle$ if and only if $\langle 3, 6 \rangle$ is more just than $\langle 1, 10 \rangle$)—entail Hammond Equity.¹⁹ Given Hammond's proof, it follows that these conditions, in conjunction with Domain Richness, Transitivity, and Anonymous Pareto entail Leximin. But because this proof, too, relies on Weaker Egalitarianism, we again have failed to derive Maximin on egalitarian grounds, for the familiar reason that, if Maximin Failure holds, then Weaker Egalitarianism and Maximin are inconsistent.

¹⁸ For a statement of Weak Conditional Contracting Extremes, see Tungodden (2000: 236). Again, Tungodden combines Weaker Egalitarianism and Weak Conditional Contracting Extremes into a single condition, but for the same reasons mentioned in fn. 16, his proof is properly interpreted in the way I present in the main text.

¹⁹ While it may seem intuitive, Strong Separability is actually quite controversial in the present context. This is because some theorists—such as Broome (2015) and Vallentyne (2009)—claim that what distinguishes egalitarian views from prioritarian ones is exactly that prioritarians but not egalitarians accept Strong Separability. For an argument that egalitarians can indeed accept Strong Separability, see Tungodden (2005).

4.3 Sen

In another famous proof, Amartya Sen (1977: Theorem 8) shows that if we accept “2-Leximin,” or Leximin in cases where only two individuals’ interests are at stake—that is, in cases where exactly two people are not indifferent between a pair of outcomes—then, so long as we accept Domain Richness and Transitivity, this entails Leximin in all cases. But Sen provides no further argument for 2-Leximin, so, as it stands, this proof does not demonstrate that Leximin has any connection with equality. To explore this connection further, let us consider outcomes x and y such that only two individuals are not indifferent between x and y , since these are the only sort of outcomes that 2-Leximin concerns. 2-Leximin says that x is more just than y in two such cases: when both non-indifferent individuals are better off in x than in y , and when the worse off non-indifferent individual is better off in x than y but the better off non-indifferent individual is worse off in x than y . It can therefore be factored into two conditions: a weakened version of Anonymous Pareto that holds when only two individuals’ interests are at stake, and a strengthened version of Hammond Equity that says that Hammond transfers make things more just (rather than only at least as just). These two conditions respectively cover the first and second cases, and their conjunction is equivalent to 2-Leximin. As we have seen, however, there is no known way to justify Hammond Equity that does not rely on Weaker Egalitarianism, so it follows that there is no known way to justify this strengthened version of Hammond Equity that does not rely on Weak Egalitarianism. But if we invoke this justification in order to show that 2-Leximin is egalitarian, then our proof relies on Weak Egalitarianism, which by Result 3 and by the truth of Maximin Failure is inconsistent with Maximin. In the presence of Maximin Failure, this proof therefore transforms into an impossibility result, too.

4.4 Deschamps and Gevers

The final formal proof in the literature, this time due to Robert Deschamps and Louis Gevers (1978: Theorem 2), once again relies on Domain Richness, Transitivity, Anonymous Strong Pareto, and Strong Separability, but also on a restriction on the sort of interpersonal comparisons we can make called “Cardinal Full Comparability,” which says that while it is possible to make level and unit interpersonal comparisons, it is not possible to make ratios (for an explanation of these different types of interpersonal comparability, see fn. 3). Together, these conditions entail that only three views are possible: Leximin, Leximax—which is the reverse of Leximin, in that it gives absolute priority to the best off individual, then to the second best off individual, then to the third best off individual, and so on—or a Utilitarian-Type view, which states that if x has a greater sum of benefits than y , then x is more just than y , but makes no judgment about the relative justice of x and y when the two have the same sum of benefits in them.

The problem with this proof is that none of the conditions it invokes have anything to do with equality. Now, one condition that Deschamps and Gevers invoke to eliminate Leximax is called “Minimal Equity,” which says, essentially, that it is not only the benefit level of the best off individual that determines the ranking of outcomes. But even still, we are left with either Leximin or a Utilitarian-type view. And, unfortunately for the egalitarian proponent of Leximin, there is no known way to complete Deschamps and Gevers’ proof by appealing to an egalitarian condition that selects Leximin over a Utilitarian-Type view.²⁰ This

²⁰ It is possible to complete the proof by limiting the sort of interpersonal comparisons we can make to level comparisons, since Utilitarian-Type views require unit comparisons. But this is not an egalitarian condition, and it makes it impossible to coherently formulate a conception of equality, because, as I pointed out in fn. 13, absolute conceptions of equality require unit comparisons, and both relative and intermediate conceptions furthermore

is especially problematic because Utilitarian-Type views are compatible with “Quasi-Egalitarianism”—the view that, when x and y have the same sum of benefits in them, and x is more equal than y , then x is more just than y —and so have a real claim to being (very weakly) egalitarian views (Tungodden 2005: 21). For this reason, one cannot simply take Deschamps and Gevers’ proof as a license to select Leximin on the grounds that it shows that Leximin and Quasi-Egalitarianism are the only two possible egalitarian views that satisfy its other conditions, and that the latter is not egalitarian enough: whether or not Leximin is egalitarian at all is exactly what is at issue. So while if we did have an argument that Leximin was egalitarian, Deschamps and Gevers’ proof might provide us with some reason to select it over other egalitarian views, that argument is exactly what I have shown that we lack, and Deschamps and Gevers’ proof does not provide it.

5. Conclusion

We have seen that every existing proof of Maximin in the literature either relies on Weak or Weaker Egalitarianism, or else is unrelated to equality. But Weak and Weaker Egalitarianism are, given the truth of Maximin Failure, inconsistent with Maximin. So if I am right that all reasonable conceptions of equality satisfy Maximin Failure, then every proof in the literature fails to deliver an egalitarian justification of Maximin. Egalitarians are therefore left without any reply to McKerlie’s challenge. If Maximin is indeed an egalitarian view, we have been given little reason to think so as of yet.

This result also suggests the following, more far-reaching point. Recall Tungodden and Vallentyne’s original proof, which showed that the plausible conditions Domain

require ratio comparisons. So this does not provide an egalitarian derivation of Leximin (or Maximin) either.

Richness, Quasi-Transitivity, Pareto, Weak Egalitarianism, and Perfect Equality entail Maximin. Since I have shown that these conditions, in conjunction with Maximin Failure, lead to an impossibility result, it turns out that one must reject at least one of these conditions. Otherwise, one is left with an inconsistent view.²¹

Perfect Equality is presumably unassailable, and I have argued at some length that we should endorse Maximin Failure. Thus, while one way for an egalitarian to avoid this conclusion would be to reject Maximin Failure (and, with it, the plausible characterization of reasonable conceptions of equality that I have shown entail it), I will not explore this option further here. Instead, I will assume that an egalitarian must reject at least one of Domain Richness, Quasi-Transitivity, Pareto, or Weak Egalitarianism. Let us begin with Quasi-Transitivity, which says that if x is more just than y , and y is more just than z , then x must be more just than z . Though many hold that this is a formal constraint on justice, an egalitarian might attempt to avoid our impossibility result by weakening Quasi-Transitivity to the following condition:

Acyclicity: if, for any set of outcomes, each outcome is more just than the last (such that the first is more just than the second, the second more just than the third, and so on until the second last is more just than the last), then the last is not more just than the first

Acyclicity claims that if x is more just than y , and y more just than z , then z is not more just than x ; it is weaker than Quasi-Transitivity since it does not claim that x must furthermore be *more* just than z . As its name suggests, it therefore prevents the justice relation from problematically “cycling” as it would if it claimed x was more just than y , y more just than z , but z more just than x . Unfortunately, though, Tungodden and Vallentyne prove that if we adopt the other conditions in Result 1 and replace Quasi-Transitivity with Acyclicity, we are

²¹ I thank an anonymous reviewer for suggesting I address this implication of my argument.

still led to the result that justice can never yield judgments that contradict Maximin (Tungodden and Vallentyne: Result 3), and by a glance at the proof of our own Result 3, we can see that Weak Egalitarianism and Maximin Failure do indeed imply that justice contradicts Maximin in this way. So weakening Quasi-Transitivity to Acyclicity does not avoid the result; instead, one can accept Domain Richness, Pareto, Weak Egalitarianism, Perfect Justice, and Maximin Failure only if one allows the justice relation to cycle. This is a price that I suspect few egalitarians would be willing to pay; though there are always exceptions (e.g., Temkin 2011).

Second, one might reject Domain Richness, which claims that the set of possible outcomes over which principles like Pareto and Weak Egalitarianism apply include all logically possible benefit distributions. For example, one might take inspiration from Rawls's "chain connection" assumption (1999: 71) and claim, say, that one's principle of distributive justice only applies when making comparisons over ranges of benefit distributions such that, for all x , if the worst off individual in x is better off than the worst off individual in y , then so too must everyone else in x be better off than they are in y . Any such domain restrictions amounts to a substantial limitation on the application of one's theory of distributive justice, but one that could perhaps be justified by empirical results pertaining to the sorts of distributions we find in actual societies. It is therefore worth investigating the possibility of justifying Maximin or some other egalitarian principle by adopting the rest of Tungodden and Vallentyne's conditions in conjunction with some plausible domain restriction.

The next option available to an egalitarian is to weaken Pareto or Weak Egalitarianism. One way to do this would be to let other considerations besides equality and efficiency enter into one's theory of distributive justice: for example, one might claim that efficiency and equality are only relevant so long as certain basic rights are satisfied, or that

they must be weighed against considerations of desert when evaluating justice. But this does not by itself solve the problem, since we may then simply reinterpret Pareto and Weak Egalitarianism as applying only to cases where equality and efficiency are the only relevant considerations: where, say, everyone's rights are satisfied and everyone gets exactly what they deserve. Still, this strategy may be worth exploring in conjunction with that of applying a domain restriction. Rawls, for example, suggests that “[c]hain connection may often be true, provided the other principles of justice [besides the Difference Principle] are fulfilled” (1999: 70). Adapting this thought, one might argue that only certain distributions of benefits are compatible, say, with respect for people's rights, and so argue that there are only a limited range of benefit distributions over which equality and efficiency can be the only relevant considerations. This would then provide one with a normative rather than an empirical basic for rejecting Domain Richness, now interpreted as a claim about the range of cases in which efficiency and equality can be the only justice-relevant considerations in play.

Assuming, however, that an egalitarian does not wish to allow cycles in the justice relation or restrict the application of her principle of distributive justice in this way, the only options left to her are either to reject Pareto or reject Weak Egalitarianism (even when reinterpreted as principles applying only in cases where efficiency and equality are the only justice-relevant considerations). For example, she might claim that efficiency is entirely unrelated to justice, and therefore reject Pareto in favor of a pure egalitarian view on which one distribution is more just than another whenever it is more equal than it. Now, if by this she means that equality is all that *justice* is concerned with, but that justice must still be balanced against efficiency (e.g., Cohen 2008), then this is a non sequitur: everything I have said until this point will extend to whichever relation—say, a generic “better than” relation—one thinks balances justice (interpreted as equality) against efficiency. But if she is making a

more substantive claim that we ought to be pure egalitarians about whatever normative criterion balances equality against efficiency (I will continue to call it “justice”), then a widely noted problem for her view is that it commits us to a very strong sort of leveling down: to claiming, for example, that $\langle 1, 1, 1 \rangle$ is more just than $\langle 2, 5, 8 \rangle$ even though everyone is worse off in the former than the latter (e.g., Parfit 1997). And even if one does not endorse pure egalitarianism per se, the mere rejection of Pareto at least commits one to a weaker sort of leveling down, in the sense that it entails that there is at least some case where even though some people are better off and no one is worse off in x than in y (such that x is Pareto superior to y), x is not more just than y . This, again, is an implication that many egalitarians have been loathe to accept—most egalitarians have gladly accepted Pareto, and attempted to defend views that are consistent with it—though, again, there are always exceptions (e.g., Temkin 2003).

Finally, one might reject Weak Egalitarianism itself. Weak Egalitarianism, recall, says that in Anonymous Pareto incomparable cases, the more equal distribution is the more just one. One way to weaken this is to adopt Weaker Egalitarianism, and claim only that, in such cases, the more equal distribution is *at least as just* as the less equal one. But this condition is, as our discussion of other proofs in the literature has shown, inconsistent with various other sets of plausible egalitarian-friendly conditions, and furthermore never allows considerations of equality to rank one outcome as *more* just than another. Other, more interesting ways of weakening Weak Egalitarianism are therefore worth investigating.

Now, despite its name, Weak Egalitarianism is in fact a rather strong condition, since it insists that considerations of efficiency are entirely irrelevant when deciding between Anonymous Pareto incomparable alternatives. This would follow were Pareto and Anonymous Pareto the only interpretations of efficiency available, but they are not: if we

interpret efficiency in, say, utilitarian terms—such that x is more efficient than y if the sum of the benefits in x are greater than the sum of the benefits in y —then it is possible to give efficiency a greater role in determining justice. For example, we might accept with Quasi-Egalitarians that equality only serves as a “tie-breaker” when comparing outcomes with the same sum of benefits in them, or strike a middle ground where, in Anonymous Pareto incomparable cases, equality and efficiency (interpreted in utilitarian terms) are “relevant but only pro tanto considerations” that must somehow be weighed against each other (Tungodden and Vallentyne 2005: 130). The former possibility is unlikely to satisfy many egalitarians, but the latter is certainly a live option, and one that, I suspect, many egalitarians would gladly accept. Exercising it would, however, require egalitarians to address the worry that equality and efficiency are incommensurable, and so cannot be weighed against each other after all—a worry that Maximin appeared to avoid by giving absolute priority to Pareto efficiency over equality rather than weighing them against each other.

Our reflection on various attempts to derive Maximin from weak and plausible conditions relating equality and efficiency to Maximin has therefore led us to the following conclusions. First, given any reasonable conception of equality (in particular, one that satisfies Maximin Failure), each such proof either collapses into an impossibility result or turns out not to involve equality at all, leaving egalitarians without a good defense of Maximin. Second, egalitarians must therefore either (a) refute my argument that all reasonable conceptions of equality satisfy Maximin Failure, (b) allow the justice relation to cycle, (c) restrict the application of their principle for weighing equality and efficiency to a limited domain of benefit distributions (on either empirical or normative grounds), (d) reject the principle of Pareto efficiency, or (e) weaken the role of equality by allowing considerations of efficiency to play some role in determining justice even when comparing

Anonymous Pareto incomparable alternatives. In my own view, the last option is the most promising. But having laid out what it at stake, I will leave it to others to decide which route to pursue.

6. Appendix

Result 3: Given Domain Richness, all reasonable conceptions of equality—that is, all conceptions of equality that satisfy Restricted Equality Transitivity, at least one of (a) Pigou-Dalton and (b) Not Just The Bottom Two, and at least one of one of (i) Translation Invariance, (ii) Scale Invariance, and (iii) Compromise Condition—satisfy Maximin Failure.

Proof: We will first prove the result for (a) Pigou-Dalton, and then for (b) Not Just The Bottom Two.

(a) Suppose that the relevant conditions do not entail Maximin Failure. Then there is no case where outcomes x and y are Anonymous Pareto incomparable, the worst off individual in x is better off than the worst off individual in y , and y is more equal than x . For example, $\langle 9, 15, 15 \rangle$ cannot be more equal than $\langle 14, 14, 50 \rangle$. We now show, just to take this one example, that according to each of (i) Translation Invariance, (ii) Scale Invariance, and (iii) Compromise Condition, $\langle 9, 15, 15 \rangle$ is indeed more equal than $\langle 14, 14, 50 \rangle$. First, by Pigou-Dalton, $\langle 9, 15, 15 \rangle$ is more equal than $\langle 1, 15, 23 \rangle$ and $\langle 1, 15, 23 \rangle$ is more equal than $\langle 1, 1, 37 \rangle$. By Restricted Equality Transitivity, $\langle 9, 15, 15 \rangle$ is therefore more equal than $\langle 1, 1, 37 \rangle$. Furthermore, by Translation Invariance $\langle 1, 1, 37 \rangle$ is just as equal as $\langle 14, 14, 50 \rangle$. So by Restricted Equality Transitivity, $\langle 9, 15, 15 \rangle$ is more equal than $\langle 14, 14, 50 \rangle$. This proves (i). Second, by Pigou-Dalton $\langle 9, 15, 15 \rangle$ is more equal than $\langle 7, 15, 17 \rangle$ and $\langle 7, 15, 17 \rangle$ is more equal than $\langle 7, 7, 25 \rangle$. By Restricted Equality Transitivity, $\langle 9, 15, 15 \rangle$ is therefore more equal than $\langle 7, 7, 25 \rangle$. Furthermore, by Scale Invariance $\langle 7, 7, 25 \rangle$ is just as equal as $\langle 14, 14, 50 \rangle$ and by Compromise Condition $\langle 7, 7, 25 \rangle$ is more equal than $\langle 14,$

14, 50>. Thus, given either condition, it follows by Restricted Equality Transitivity that <9, 15, 15> is more equal than <14, 14, 50>. This proves (ii) and (iii).

(b) By Not Just The Bottom Two, we may stipulate that there is some x and some y such that the worst off pair of individuals in x are equally well off, and the worst off pair of individuals in y are not, but y is more equal than x . Let x_w be the benefit level of the two worst off individuals in x , and y_{w1} and y_{w2} the respective benefit levels of the two worst off individuals in y . Since $y_{w1} < y_{w2}$, there will always be some number r such that $y_{w1} + r < x_w < y_{w2} + r$. We now add r to the benefit level of all the individuals in y to obtain y' . By Translation Invariance y' is just as equal as y , and by stipulation y is more equal than x , so it follows by Restricted Equality Transitivity that y' is more equal than x . But by design, x and y' are Anonymous Pareto incomparable, and the worst off individual in y' is worse off than the worst off individual in x ($y_{w1} + r < x_w$). This entails Maximin Failure and proves (i). Next, since $y_{w1} < y_{w2}$, there will always be some positive number q such that $y_{w1}q < x_w < y_{w2}q$. We now multiply q by the benefit level of all the individuals in y to obtain y'' . By Scale Invariance y'' is just as equal as y , and by stipulation y is more equal than x , so it follows by Restricted Equality Transitivity that y'' is more equal than x . But by design, x and y'' are Anonymous Pareto incomparable, and the worst off individual in y'' is worse off than the worst off individual in x ($y_{w1}q < x_w$). This entails Maximin Failure and proves (ii). Finally, since $y_{w1} < y_{w2}$, there are three cases: $y_{w1} < x_w < y_{w2}$, $y_{w1} < y_{w2} \leq x_w$, or $x_w \leq y_{w1} < y_{w2}$. If the first obtains, then Maximin Failure already follows: by stipulation, x and y are Anonymous Pareto incomparable, the worst off individual in x is worse off than the worst off individual in y ($y_{w1} < x_w$), but y is more equal than x . If the second obtains, then there is some number s ($s > 0$) such that $y_{w1} + s < x_w < y_{w2} + s$, and we may add s to the benefit level of all the individuals in y to obtain y''' . If the third obtains, then there is some number t ($0 < t < 1$) such that $y_{w1}t < x_w < y_{w2}t$, and we may multiply t

by the benefit level of all the individuals in y to obtain y''' . Either way, it follows by Compromise Condition that y''' and y'''' are both more equal than y . By stipulation, y is more equal than x , so by Restricted Equality Transitivity both y''' and y'''' are more equal than x . But by design y''' and y'''' are again Anonymous Pareto incomparable with x , and the worst off member in each is worse off than the worst off member in x ($y_{w'} + s < x_w$ and $y_{w'} t < x_w$). Thus, all three cases entail Maximin Failure. This proves (iii), and the result follows.

7. References

- Barry B (1989) *Theories of Justice: Volume 1*. Berkeley: University of California Press.
- Bossert W and Pfingstein A (1990) Intermediate Equality: Concepts, Indices, and Welfare Implications. *Mathematical Social Sciences* 19(2): 117-134,
- Bossert W and Pfingstein A (2004) Utility in Social Choice. In Barbera S, Hammond PJ, Seidl C (eds) *Handbook of Utility Theory: Volume 2 Extensions*. New York: Springer, pp. 1099-1177.
- Broome J (2015) Equality Versus Priority: A Useful Distinction. *Economics and Philosophy* 31 219-228.
- Cohen, G. A. (2008) *Rescuing Justice & Equality*. Cambridge: Harvard University Press.
- Dalton H (1920) The Measurement of The Inequality of Incomes. *The Economic Journal* 30(119); 348-361.
- d'Aspremont C and Gevers L (2002) Social Choice Functionals and Interpersonal Comparability. Arrow KJ, Sen AK, and Suzumura K (eds) *Handbook of Social Choice and Welfare: Volume 1*. New York: North-Holland, pp. 459-541.
- Deschamps R and Gevers L (1978) Leximin and Utilitarian Rules: A Joint Characterization. *Journal of Economic Theory* 17(2): 143-168.

- Hammond PJ (1976) Equity, Arrow's Conditions, and Rawls' Difference Principle. *Econometrica* 44(4): 793-804.
- Hardin R (2003) *Indeterminacy and Society*. Princeton: Princeton University Press.
- Hicks J (1939) The Foundations of Welfare Economics. *The Economic Journal* 49(196): 696-712.
- Kaldor N (1939) Welfare Propositions in Economics and Interpersonal Comparisons of Utility. *The Economic Journal* 49(195): 549-552.
- Kolm S (1999) The Rational Foundations of Income Inequality Measurement. In Silber J (ed) *Handbook of Income Inequality Measurement*. New York: Springer, pp. 19-100.
- Kolm S (1976) Unequal Inequalities I. *Journal of Economics Theory* 12(3): 416-442.
- McKerlie D (1994) Equality and Priority. *Utilitas* 6(1): 25-42.
- Parfit D (1997) Equality and Priority. *Ratio* 10(3): 202-221.
- Pigou AC (1912) *Wealth and Welfare*. London: MacMillan.
- Rawls J (1999) *A Theory of Justice*. Cambridge: Harvard University Press.
- Sen A (1997) *On Economic Inequality*. Oxford: Clarendon Press.
- Sen A (1977) On Weights and Measures: Informational Constraints on Social Welfare. *Econometrica* 45(7): 1539-1572
- Suppes P (1966) Some Formal Models of Grading Principles. *Synthese* 16: 284-306.
- Temkin L (2003) Equality, Priority, or What?. *Economics and Philosophy* 19: 61-87,
- Temkin L (1993) *Inequality*. New York: Oxford University Press.
- Temkin L (2011) *Rethinking the Good*. New York: Oxford University Press.
- Tungodden B (2000) Egalitarianism: Is Leximin the Only Option?. *Economics and Philosophy* 16: 229-245.
- Tungodden B (2003) The Value of Equality. *Economics and Philosophy* 19: 1-44.

Tungodden B and Vallentyne P (2005) On The Possibility of Paretian Egalitarianism. *Journal of Philosophy* 102(3): 126-154.

Vallentyne P (2000) Equality, Efficiency, and the Priority of the Worst Off. *Economics and Philosophy* 16: 1-19.

Vallentyne P (2009) Sen on Sufficiency, Priority, and Equality. In Morris CW (ed) *Amartya Sen*. New York: Cambridge University Press, pp. 138-169.