

ON THE POSSIBILITY OF NONAGGREGATIVE PRIORITY FOR THE WORST OFF*

BY MARC FLEURBAEY, BERTIL TUNGODDEN,
AND PETER VALLENTYNE

I. INTRODUCTION

We shall focus on moral theories that are solely concerned with promoting the benefits (e.g., well-being) of individuals and shall explore the possibility of such theories' ascribing some priority to benefits to those who are worse off—without this priority being absolute. Utilitarianism (which evaluates alternatives on the basis of total or average benefits) ascribes no priority to the worse off, and leximin (which evaluates alternatives by giving lexical priority to the worst off, and then the second worst off, and so on) ascribes absolute priority to the worse off (i.e., favors even a very small benefit to a worse-off person over very large benefits to large numbers of better-off people). Neither extreme view, we assume, is plausible.

One intermediate view is *finitely weighted total prioritarianism*. It is like utilitarianism except that it is concerned with total weighted benefits. It differs, however, in that there is a prioritarian weighting for benefits: benefits to individuals have finitely decreasing marginal weights (i.e., the moral importance of a one-unit increase in benefits is smaller the more units the individual already has). Thus, like leximin, it gives priority to the worse off. Unlike leximin, however, this priority is not absolute or infinite. Large enough benefits to the better off can take priority over the benefits to a worse-off individual. Finitely weighted total prioritarianism is commonly considered an attractive compromise between utilitarianism and leximin. It shares with utilitarianism, however, a questionable property: namely, that it justifies that a significant benefit to the worst-off person can be outweighed by trivial benefits to enough best-off people. In this sense, it can be argued that it does not give enough priority to the worse off.

Is it possible to avoid the problem of finitely weighted total prioritarianism without assigning absolute priority to the worse off? We show that it is not easy to do so. We do this by considering the possibility of combining nonaggregative priority for the worst off over the best off while also allowing some aggregation (e.g., when the worst off are not affected).

* For most helpful comments, we thank Nils Holtug, Mike Otsuka, Ellen Paul, Alex Voorhoeve, the other contributors to this volume, and its editors.

II. GENERAL BACKGROUND

Throughout, we leave open the relevant conception of benefit (resources, primary goods, brute luck well-being, etc.). References to a person's being worse off than another should be understood in terms of whatever benefits are relevant.

We shall assume, for the sake of argument, that benefits are fully measurable and interpersonally comparable. This may seem like a strong assumption, but in the present context it avoids burdening our analysis with informational constraints. The assumption that benefits are so measurable and comparable does not entail that such information is *relevant* for the moral assessment of options. The assumption is simply that such information is available. This ensures that no principle of moral goodness is ruled out merely on the grounds that it presupposes that benefits are measurable or comparable in ways that they are not. Our main argument is that no moral theory satisfies all of certain seemingly plausible conditions—even if benefits are fully measurable and interpersonally comparable. In this context, the measurability and comparability assumptions should make it easier to find a satisfactory theory and thus strengthen the significance of our conclusions.

We shall be concerned with the assessment of the moral goodness of alternatives, where alternatives are possible objects of choice (e.g., actions or social policies). Alternatives may have all kinds of features: they generate a certain distribution of benefits, satisfy or violate various rights, involve various intentions, and so on. In what follows, we shall assume that the only relevant information for the assessment of moral goodness is the benefit distribution that an alternative generates. More formally, we shall assume:

Benefitism: Alternatives can be identified with, and their moral goodness assessed solely on the basis of, their benefit distributions.

Benefitism is a generalization of welfarism (the view that alternatives can be assessed solely on the basis of the resulting distributions of welfare). Although it does not assume that welfare is all that matters, it does assume that moral goodness supervenes on individual benefits. If two alternatives generate the same distribution of benefits, then they have the same status with respect to moral goodness.

Benefitism is a very strong assumption. It rules out the relevance, for example, of respecting the wills of agents (e.g., as reflected in their consent), if this is not part of the evaluation of alternatives. It holds that there is no difference with respect to moral goodness between forcing a person to go jogging against her will and merely successfully encouraging her to do so, if the benefits to her and everyone else are the same. Although we would reject this condition because of this implication (or doubt that

benefits can be the same in this example, on a reasonable measure of benefits), we shall here grant it, since it is not relevant for the issue at hand. The issue of whether something like leximin is the only plausible egalitarian theory of moral goodness arises even if equality-promotion is limited, for example, to those alternatives that respect the wills of agents (e.g., violate no rights). Granting Benefitism will simplify the presentation. The results could, however, be recast to apply only where the relevant differences between alternatives concern the distributions of benefits that they generate.

Given Benefitism, we can identify an alternative with the benefit distribution that it generates, and in what follows we shall do so for simplicity. We shall further assume that all distributions of benefits display non-negative benefits (zero is the lowest conceivable level). Finally, we assume that the set of benefit distributions generated by the set of possible alternatives is *rich* in the following sense:

Domain Richness: For any logically possible benefit distribution with a finite number of individuals and non-negative benefits, there is an alternative that generates that distribution.

This condition rules out, for example, the possibility that, where there are three people, the distribution $\langle 3, 7, 9 \rangle$ (3 units of benefit to the first person, 7 to the second, 9 to the third) is not one of the alternatives. It also rules out the possibility that distributions involving certain (e.g., very large) finite population sizes are not among the alternatives. The condition requires that all logically possible benefit distributions are among the alternatives. This is not to say that all are part of any given *feasible* set (the alternatives that are open to an agent on a given occasion). Of course, there are lots of logically possible benefit distributions that are not feasible on a given occasion. The claim here is about the range of benefit distributions that can be assessed in terms of moral goodness. The condition holds that such judgments can be made for all logically possible distributions.

We believe that this is a highly plausible condition. Benefit distributions here play the role of test cases for a theory of moral goodness. All logically possible test cases—assuming, as we shall, a finite population—are admissible.

We shall focus on various conditions governing moral goodness. More exactly, we shall first focus on the relation, between alternatives, of being morally at least as good as another. Following the standard definitions, we shall say that (1) an alternative is *morally better* than another if and only if it is at least as good and the other is not at least as good as it; and (2) an alternative is *equally morally as good as* another if and only if it is at least as good and the other is also at least as good as it. We do not assume that the moral betterness relation is complete (i.e., that for any two alter-

natives, at least one is at least as good as the other). Some alternatives may be morally incomparable with others.

III. MORAL BETTERNESS

We start by reporting a slight modification of a result from Marc Fleurbaey and Bertil Tungodden showing that, given some standard conditions, it is impossible to satisfy a very weak aggregation condition and a very weak prioritarian condition.¹ Throughout, we are concerned with theories that give priority to those who are *worse off* than others. Our conditions, however, shall focus on the special case of the priority given to the *worst-off* individuals in society.

Benefits to the worse off should, we assume, have some priority over benefits to the better off, and this should ensure that a given benefit to a *worst-off* individual has priority over a sufficiently smaller benefit to all of the *best-off* individuals—no matter how many best-off people there are. More exactly, we propose to capture this idea by the following condition:

Minimal Nonaggregative Priority: For any pair of benefit levels, for any alternative in which the benefits of the worst off are at or below the lower benefit level and the benefits of the best off are at or above the higher benefit level, and for any benefit gain to a worst-off person, there is a sufficiently small benefit loss to every best-off person such that, if those who are best off remain best off and the worst-off individual who gains does not become a best-off person, then the net result is morally better—no matter how many best-off persons there are.

Compare, for example, $\langle 3,8,8 \rangle$ and $\langle 1,9,9 \rangle$. Compared with the second alternative, the first alternative gives the worst-off person 3 units rather than 1 unit of benefit. It also gives the two best-off individuals 8 units rather than 9 units of benefit. Minimal Nonaggregative Priority does not require that the first alternative be judged better than the second—since it does not tell us whether the loss of 1 to the best off is small enough relative to a gain of 2 for the worst off. If, however, a one-unit loss is small enough to satisfy the condition when the worst-off person starts with 1 unit and gains 2 units and the best-off persons start with 9 units, then the

¹ Marc Fleurbaey and Bertil Tungodden, “The Tyranny of Non-Aggregation Versus the Tyranny of Aggregation in Social Choices: A Real Dilemma” (unpublished, 2007). In order to streamline the presentation of our later results, we introduced the following modifications of one of Fleurbaey and Tungodden’s results. We strengthened both the priority condition and the aggregation condition by requiring the modified benefit distribution to be better than the original—and not merely at least as good. This strengthening permits us to invoke Acyclicity rather than the stronger Transitivity and to drop Weak Pareto for certain results.

condition does require that the first alternative be judged better—not only is $\langle 3,8,8 \rangle$ better than $\langle 1,9,9 \rangle$, but also $\langle 3,8, \dots ,8 \rangle$ is better than $\langle 1,9, \dots ,9 \rangle$, no matter how many best-off persons there are.

Minimal Nonaggregative Priority merely affirms that the worst-off person has a certain weak kind of *nonaggregative* priority over others. It holds that, for a given gain to a worst-off person, there is a sufficiently small loss to the best off such that the gain to the worst off is favored—no matter how many best off have the small loss (and thus no matter how great the total loss for the best off). The condition rules out the tyranny of aggregation (e.g., only considering total benefits) by placing limits on admissible forms of aggregation.

This is a weak condition in several ways. First, it only gives priority to a *worst-off* person in comparison with the *best-off persons*. It does not require that the badly off in general have priority over those who are better off. Second, it does not give *absolute priority* to the worst off. It is compatible with favoring a larger benefit to a best-off person over a smaller benefit to the worst-off person. Finally, the size of the loss for the best off that makes it sufficiently small to be outweighed by a given gain for the worst off may depend both on the level of benefits (prior to any gain or loss) of the best off and the level of benefits of the worst off. For example, one might think that what counts as a sufficiently small loss to the best off is larger the lower the benefit level of the worst-off individual and also larger the greater the benefit level of the best off.

The Minimal Nonaggregative Priority condition combines priority for the worst off with a limitation on aggregation that has been much discussed in the literature. T. M. Scanlon asks, for example, whether, in order to eliminate an hour of extremely painful electrical shocks to one person, we should impose much smaller costs (e.g., the loss of fifteen minutes of World Cup television coverage) on several million people. His view is that we should and that this is true independently of whether the one person who is helped is, or would be (without help), the worst-off person.² The condition we are discussing imposes Scanlon's aggregation-

² T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA.: Belknap Press of Harvard University Press, 1998), 235. For further discussion of the relevance of aggregation of benefits (e.g., the small losses of many adding up to a large total loss), see David Brink, "The Separateness of Persons, Distributive Norms, and Moral Theory," in R. G. Frey and C. W. Morris, eds., *Value, Welfare, and Morality* (Cambridge: Cambridge University Press, 1993), 252–89; Frances Kamm, *Morality, Mortality I* (New York: Oxford University Press, 1993), chaps. 5–6; Rahul Kumar, "Contractualism on Saving the Many," *Analysis* 61 (2001): 165–70; Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135–67; Alastair Norcross, "Two Dogmas of Deontology: Aggregation, Rights, and the Separateness of Persons," elsewhere in this volume; Michael Otsuka, "Scanlon and the Claims of the Many Versus the One," *Analysis* 60 (2000): 288–93; Michael Otsuka, "Skepticism about Saving the Greater Number," *Philosophy and Public Affairs* 32 (2004): 413–26; Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy and Public Affairs* 34 (2006): 109–35; Derek Parfit, "Innumerate Ethics," *Philosophy and Public Affairs* 7 (1978): 285–301; Derek Parfit, "Justifiability to Each Person," *Ratio* 16

limitation idea, but applies it only to the case where the loss is imposed on best-off individuals and is sufficiently small. It is thus much weaker than the general aggregation-limitation idea.

This condition is violated by utilitarianism (which gives no priority to the worst off) and also by finitely weighted prioritarianism (which gives some finite priority to the worse off, but which can be outweighed by gains to best-off people that are as small as you like, as long as there are enough such people). Leximin, however, satisfies this condition (since it always gives absolute priority to the worse off).

The next condition requires that there should be at least some limited form of aggregation:

Minimal Aggregation: For any benefit distribution, imposing a loss on one person and gains on everyone else makes things morally better if the loss is small enough and the gains are large enough.

This condition is very weak. It rules out leximin, but both utilitarianism and finitely weighted prioritarianism satisfy it. It only requires that a loss for one person is outweighed by the gains to others, when the gain is sufficiently large and the loss is sufficiently small. It says, for example, that things are made morally better by increasing the benefits for a million people by 100 units and decreasing the benefits of the worst-off person by some sufficiently small amount (e.g., some completely trivial amount). This condition does not require that the alternative with the greater total benefits always be judged better. It is silent, for example, about whether things are made morally better by increasing the benefits for a million people by 100 units and decreasing the benefits of the worst-off person by 100 units (since the gain of 100 for others may not be large enough to offset the loss of 100 for the one person).

Finally, Fleurbaey and Tungodden endorse three standard conditions: Transitivity, Weak Pareto, and Replication Invariance. We discuss each of these in turn.

Consider first:

Transitivity: If, for alternatives, x , y , and z , x is morally better than y , and y is morally better than z , then x is morally better than z .

This, of course, is a very standard assumption. In what follows, we only need a much weaker assumption:

(2003): 368–90; T. M. Scanlon, “Replies,” *Ratio* 16 (2003): 424–39; John Taurek, “Should the Numbers Count?” *Philosophy and Public Affairs* 6 (1977): 293–316; and Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), chap. 6.

Acyclicity: If, for alternatives x_1, \dots, x_n , x_1 is morally better than x_2 , x_2 is morally better than x_3, \dots and x_{n-1} is morally better than x_n , then x_n is not morally better than x_1 .

Acyclicity is about as close to being uncontroversial as one can get when it comes to consistency requirements on the moral-betterness relation. If x_1 is better than x_2 , and x_2 is better than x_3 , then Acyclicity, like Transitivity, rules out the possibility that x_3 is better than x_1 . Unlike Transitivity, however, Acyclicity allows that x_1 and x_3 may be equally good or incomparable.

The second standard condition introduced in Fleurbaey and Tungodden's analysis is:

Weak Pareto: For any two alternatives x and y , if each person has greater benefits in x than in y , then x is morally better than y .

Weak Pareto is a weak efficiency condition on the promotion of benefits (much weaker than the utilitarian sum-total conception of efficiency). It requires, for example, that $\langle 2,4,6 \rangle$ be judged morally better than $\langle 1,3,5 \rangle$ and also morally better than $\langle 1,1,5 \rangle$. It is silent about whether $\langle 2,4,6 \rangle$ is morally better than $\langle 99,1,6 \rangle$. This is, we believe, an extremely plausible condition. It is satisfied by utilitarianism, leximin, and finitely weighted total prioritarianism.

The third standard condition introduced in Fleurbaey and Tungodden's analysis is the condition that moral goodness is not affected when the number of people in a distribution is increased by "replication." Let an n -replication of a given distribution of benefits be the result of adding, for each person with benefits b , $n-1$ new people with benefits b . For example, a 3-replication of $\langle 2,6 \rangle$ is $\langle 2,6,2,6,2,6 \rangle$. Our definition of the condition, then, is:

Replication Invariance: For any two alternatives, x and y , and any positive integer n , x is morally at least as good as y if and only if the n -replication of x is morally at least as good as the n -replication of y .

This requires, for example, that if $\langle 1,5 \rangle$ is at least as good as $\langle 2,3 \rangle$, then $\langle 1,5,1,5,1,5 \rangle$ is at least as good as $\langle 2,3,2,3,2,3 \rangle$. We believe that this condition is highly plausible. It is satisfied by leximin, utilitarianism, and finitely weighted total prioritarianism. Replication does not affect average benefits, and thus average utilitarianism satisfies the condition. N -replication does affect the total of a given alternative, but it does so by multiplying the total by n . Thus, if the total of one alternative is at least as great as that of a second, the total of an n -replication of the first will be at least as great as that of an n -replication of the second. Thus, total utilitarianism satisfies the Replication-Invariance condition. Finitely weighted total prioritarianism satisfies the condition for similar reasons (given that the weights are based

on the absolute level of the benefit). Leximin also satisfies the condition, because a replication will not change the level of the worst off (or second worst off, and so on) in any of the two alternatives in a comparison.

A slight modification of the proof in Fleurbaey and Tungodden's earlier analysis yields the following result:

Result 1: No moral-betterness relation satisfies Weak Pareto, Replication Invariance, Acyclicity, Minimal Nonaggregative Priority, and Minimal Aggregation.³

The proof takes roughly the following form: The first four conditions entail that an alternative that gives greater benefits to the worst-off person is not morally worse. This violates the last condition, which requires that, at least in some case (where the gain for the worst off is small enough in relation to the gain of others), the alternative in question is morally worse (or equal). Hence, we have an impossibility.

Let us note how some standard theories violate at least one of the foregoing five conditions. Leximin satisfies all the conditions except Minimal Aggregation. Utilitarianism (average and total) and finitely weighted total prioritarianism satisfy all the conditions except for Minimal Nonaggregative Priority. Thus, none of these standard theories satisfies both Minimal Nonaggregative Priority and Minimal Aggregation. Fleurbaey and Tungodden prove more generally that no theory satisfies all the conditions.⁴

This result suggests that it is impossible to combine a plausible kind of aggregation with a plausible kind of priority for the worse off. Indeed, because the priority condition is merely a special case of a more general limitation on the role of aggregation (roughly: not allowing a large number of small enough gains to take priority over a single large loss), the results show that it is not possible to limit aggregation in this way. One might, of course, wonder whether some of the conditions discussed above are implausibly strong. We shall now show that the impossibility result remains even with significant weakenings of these conditions, and even if we do not invoke Weak Pareto.

First, however, let us note that the impossibility is not an inherent incompatibility between the aggregation condition and the priority condition. These two conditions are compatible—even in conjunction with Weak Pareto and Transitivity. These conditions are satisfied, for example,

³ Fleurbaey and Tungodden ("The Tyranny of Non-Aggregation") also establish an impossibility result which relies on a version of Minimal Aggregation that focuses on the cases where there is a uniquely-worst-off individual who remains the worst off even after he receives the benefit gain, and where Replication Invariance is replaced with a condition called Reinforcement. Reinforcement states that if x is better than y , then x' is better than y' if the only difference between x and x' and y and y' is that a person who is better off in y than in x is removed in y' and in x' .

⁴ See Fleurbaey and Tungodden, "The Tyranny of Non-Aggregation Versus the Tyranny of Aggregation."

by the *geometric Gini* betterness relation (a modification of the standard Gini betterness relation, which weights benefits by the rank order: weight of 1 for the highest benefits, 3 for the second-highest benefits, 5 for the third-highest benefits, . . . and $2n - 1$ for the worst-off level of benefits, where there are n people). This relation ranks alternatives on the basis of their total *geometric Gini-weighted benefits*, where the geometric Gini-weights are, for instance, 1 for the benefits of the best-off position, 2 for the benefits of the second-best-off position, 4 for the third-best-off position, and 2^{n-1} for the benefits of the worst-off position (where there are n people). (If there is more than one person at a given benefit level, the relevant positions are allocated arbitrarily among them.) For example, the geometric Gini-weighted benefits for $\langle 2,4,6 \rangle$ add up to a total of 22 ($2 \times 2^2 + 4 \times 2^1 + 6 \times 2^0$). This betterness relation satisfies Minimal Aggregation, since, for any benefit to others, there is a small enough loss to the worst-off position such that the total of the Gini-weighted benefits is made greater by the combination of benefits to others and loss to the uniquely worst off. Geometric Gini also satisfies Minimal Nonaggregative Priority, since the weight for the worst-off position (2^{n-1}) is always greater than the total weight given to others ($1 + 2 + 4 + \dots + 2^{n-2}$, which equals $2^{n-1} - 1$). For example, for three people, the weight for the worst off is 4 and the total weight for the others is 3. Thus, for any given benefit to the worst off, if everyone else loses no more than that benefit amount, the total weighted benefits will be increased.

Geometric Gini, however, does not satisfy one of the standard conditions discussed above, namely, Replication Invariance. To see this, consider $\langle 5,5 \rangle$, $\langle 3,10 \rangle$, and their 2-replications, $\langle 5,5,5,5 \rangle$ and $\langle 3,3,10,10 \rangle$. Geometric Gini judges $\langle 5,5 \rangle$ (weighted benefits of $1 \times 5 + 2 \times 5$, or 15) as strictly worse than $\langle 3,10 \rangle$ (weighted benefits of $1 \times 10 + 2 \times 3$, or 16), but it judges $\langle 5,5,5,5 \rangle$ (weighted benefits of $1 \times 5 + 2 \times 5 + 4 \times 5 + 8 \times 5$, or 75) as strictly better than $\langle 3,3,10,10 \rangle$ (weighted benefits of $1 \times 10 + 2 \times 10 + 4 \times 3 + 8 \times 3$, or 66).

One way of avoiding the impossibility result, then, is to reject Replication Invariance. We believe, however, that the condition is highly plausible, and our focus will be on the other conditions.

Let us, then, reconsider Minimal Nonaggregative Priority. It requires that priority be given to a benefit to a single worst-off person over a sufficiently smaller benefit to each of the best-off individuals—no matter how many there are. There are at least two ways that the priority may be too strong.⁵ First, the benefit given to the worst-off individual may raise her above the average in terms of benefits. One might accept the strong

⁵ A third way that one might think that the priority condition is too strong is that it requires priority for just one worst-off person. One might agree that priority should be given to benefits to the worst-off individuals *when there are enough of them* but reject the idea that a single worst-off person should be given such priority. However, given Replication Invariance, this is equivalent to the condition with only one person gaining. Thus, we do not introduce this weakening of the priority condition.

priority where the beneficiary remains below average, but reject it in some cases where the beneficiary acquires significantly above-average benefits.

Second, the priority condition may be too strong when the worst-off individual is already superaffluent and the best off are only slightly better off. Minimal Nonaggregative Priority says that giving the worst-off individual a trivial benefit that keeps her below the others, makes things better when coupled with a sufficiently smaller loss to each of the (very slightly better off) best-off individuals—*no matter how many best-off individuals there are*. The total loss in benefits will be enormous, if there are enough best-off people. It may seem implausible to think that imposing such a loss makes things morally better—where the only benefit is a trivial benefit to one superaffluent individual. Similarly, one may object to Minimal Nonaggregative Priority when both the best off and the worst off are poor in absolute terms. One might, that is, favor priority where the worst-off person is sufficiently poor in absolute terms and the best off are sufficiently affluent in absolute terms without favoring it in other cases.

Consider, then, the following weakening of the priority condition:

Ultra Minimal Nonaggregative Priority: There is a pair of benefit levels such that for any alternative in which the benefits of the worst off are below the lower benefit level and the benefits of the best off are above the higher benefit level, and for any benefit gain to a worst-off person, there is a sufficiently small benefit loss to every best-off person such that, if (1) those who are best off remain best off and above the higher benefit level, and (2) the worst-off individual remains below average and below the lower benefit level, then the net result is morally better—no matter how many best-off persons there are.

There are two differences between this condition and Minimal Nonaggregative Priority. First, Minimal Nonaggregative Priority only requires that the worst-off individual to whom a benefit is given must remain worse off than the *best-off* individuals—whereas the Ultra Minimal version above requires that the worst-off individual must remain below *average* with the benefit. The second difference is that Ultra Minimal Nonaggregative Priority requires that the worst-off individuals be below some benefit level (e.g., abject poverty) and that the best-off individuals be above some higher benefit level (e.g., superaffluence). We believe that Ultra Minimal Nonaggregative Priority is extremely plausible, and we shall show that the impossibility result remains even with this weaker priority condition.

Let us also reconsider Minimal Aggregation. Its most problematic implication concerns cases where there are just *two* people. Minimal Aggregation holds that giving a sufficiently large benefit to one person and imposing a sufficiently small loss on the other makes things morally better. One

might agree that giving *enough people* a sufficiently large benefit and imposing a sufficiently small loss on one person makes things morally better without holding that this is so when there is only one person who gets the large benefit. For example, where the person suffering the loss is the worst-off person, one might hold, in the two-person case, that the changes make things morally worse. Consider, then, the following weakening of the aggregation condition:

Ultra Minimal Aggregation: There is some positive population size such that for any benefit distribution with a larger population, imposing a small enough loss on one person and a large enough gain on everyone else makes things morally better.

This is the same as Minimal Aggregation, except that it holds only for sufficiently large populations (and not necessarily two-person populations). We believe that the weakened condition is extremely plausible. We now note:

Result 2: No moral-betterness relation satisfies Replication Invariance, Acyclicity, Ultra Minimal Nonaggregative Priority, and Ultra Minimal Aggregation.

This result establishes that the impossibility result holds even when (1) Minimal Nonaggregative Priority is replaced with the much weaker Ultra Minimal Nonaggregative Priority, (2) Minimal Aggregation is replaced with the much weaker Ultra Minimal Aggregation, (3) Transitivity is replaced with the much weaker Acyclicity, and (4) Weak Pareto is removed.⁶ Here we can note that geometric Gini violates Replication Invariance, utilitarianism and finitely weighted prioritarianism violate Ultra Minimal Nonaggregative Priority, and leximin violates Ultra Minimal Aggregation. The proof generalizes this to show that *no moral-betterness relation* satisfies all the conditions.

Let us provide an illustration of the structure of the proof:

(1) Suppose that Ultra Minimal Nonaggregative Priority applies when the worst-off individual is below 10 units of benefit, the best-off individuals are above 80 units of benefit, and, for a gain to the worst off of 2 or more, the loss to the best off needs only to be 1 or less. Consider now the following sequence of alternatives for a population with five hundred individuals, where initially there are one hundred individuals with benefits of 0 and four hundred individuals with benefits of 200. In this sequence, each successive alternative increases the benefits of one indi-

⁶ The proofs of this and all other results are posted online at <http://mora.rente.nhh.no/projects/EqualityExchange/>.

vidual from 0 to 3, and decreases the benefits of four hundred individuals by one unit.

Z0: $\langle 0, 0, \dots, 0, 200, \dots, 200 \rangle$
 Z1: $\langle 3, 0, 0, \dots, 0, 199, \dots, 199 \rangle$
 Z2: $\langle 3, 3, 0, 0, \dots, 0, 198, \dots, 198 \rangle$
 ...
 Z99: $\langle 3, 3, \dots, 3, 0, 101, \dots, 101 \rangle$
 Z100: $\langle 3, 3, \dots, 3, 3, 100, \dots, 100 \rangle$

In this sequence, Ultra Minimal Nonaggregative Priority judges each alternative from Z1 to Z100 to be morally better than the previous one. Hence, by Acyclicity, Z0 is not morally better than Z100.

(2) Suppose that Ultra Minimal Aggregation applies when the population contains at least five people ($n = 5$), the gainers gain at least 50, and the one loser loses less than 5. Consider the following alternatives:

X: $\langle 3, 100, 100, 100, 100 \rangle$
 Y: $\langle 0, 200, 200, 200, 200 \rangle$

By Ultra Minimal Aggregation, Y is morally better than X.

(3) Consider the 100-replications of X and Y, 100-X and 100-Y. By (2), Y is morally better than X. Hence, by Replication Invariance, 100-Y is morally better than 100-X. But this violates the conclusion in (1), given that 100-Y is Z0 and 100-X is Z100.

Thus, in this example, no moral-goodness relation satisfies the four conditions. The proof generalizes this illustration for any values of the parameters in Ultra Minimal Nonaggregative Priority and Ultra Minimal Aggregation.

We can weaken the priority condition even further and still obtain an impossibility result. Ultra Minimal Nonaggregative Priority has force no matter how small the gain is to a worst-off individual. One might think that the priority should hold only when the gain is large enough (e.g., nontrivial). Consider, then:

Super Ultra Minimal Nonaggregative Priority: There is a pair of benefit levels and a minimum gain size (smaller than the lower benefit level) such that for any alternative in which the benefits of the worst off are below the lower benefit level and the benefits of the best off are above the higher benefit level, and for any benefit gain to a worst-off person that is no less than the minimum gain size, there is a sufficiently small benefit loss to every best-off person such that, if (1) those who are best off remain best off and above the higher benefit level, and (2) the worst-off individual remains below average and

below the lower benefit level, then the net result is morally better—no matter how many best-off persons there are.

Even this weakening, however, does not avoid the impossibility result, if we also impose the following very plausible condition:

*Pigou-Dalton:*⁷ A decrease in the benefits of one person combined with an increase of the same amount in the benefits of a worse-off person makes things morally better, if the originally worse-off person does not become better off than the originally better-off person and no other changes take place in the distribution.

This requires that a transfer from a better-off to a worse-off person must be judged to make the situation morally better. For example, the condition requires that $\langle 7,5,2 \rangle$ be judged morally better than $\langle 9,5,0 \rangle$. It is an extremely weak prioritarian condition. It merely requires that priority be given for the worse off when the total benefits are not affected. It is completely uncontroversial within a prioritarian approach. We now note:

Result 3: No moral-betterness relation satisfies Pigou-Dalton, Replication Invariance, Acyclicity, Super Ultra Minimal Nonaggregative Priority, and Ultra Minimal Aggregation.⁸

Both Ultra Minimal Aggregation and Super Ultra Minimal Nonaggregative Priority seem very plausible. If a sufficiently large number of people gain sufficiently, and just one person loses a sufficiently small amount relative to these gains, then the result is morally better. To give absolute priority to this one person in such cases seems wrong, even if the person is the worst off. Thus, Ultra Minimal Aggregation is in line with our intuition that moral goodness requires at least some aggregation. Super Ultra Minimal Nonaggregative Priority also seems to capture a basic moral intuition. It is always better to give a large enough gain to a worst-off person and impose trivial losses on the best off (relative to the gain for the worst off)—no matter how many best-off individuals there are. Aggre-

⁷ This is a standard prioritarian condition from the social choice literature and is named after Arthur Pigou and Hugh Dalton.

⁸ The following relation satisfies all the conditions except for Pigou-Dalton. Consider a sequence of benefit thresholds, $g, 2g, 3g, \dots$, where g is the minimum gain of the worst off appealed to in Super Ultra Minimal Nonaggregative Priority. The relation judges that x is morally at least as good as y if and only if (1) for some benefit threshold in the sequence, x has fewer people below that threshold than y does but has the same number of people below each of the lower thresholds, or (2) for every benefit threshold, the number of people below that threshold is the same in x as in y , and the weighted total prioritarian benefits are at least as great in x as in y . This violates Pigou-Dalton because a transfer from a person above a threshold to a person below the threshold may push the better-off person below the threshold, which makes things worse according to this relation.

gation of such trivial losses is not relevant. Result 3 thus seems to establish that no plausible theory of moral goodness is sufficiently sensitive to prioritarian and aggregative concerns. We shall now show that a similar result holds for theories of moral permissibility.

IV. MORAL PERMISSIBILITY

Above, we focused on the axiological issue of moral betterness. We shall now focus on the deontic issue of moral permissibility, which is not a binary relation like moral betterness but a function selecting a subset of permissible alternatives relative to a set of feasible alternatives. The alternatives selected from a feasible set are morally permissible (relative to this set), while the other feasible options are morally impermissible. Just as for moral betterness we consider all logically conceivable distributions of benefits, for moral permissibility we will assume that all logically conceivable sets of distributions are potential feasible sets over which moral permissibility may be defined. No subsets of alternatives are excluded *a priori* from the analysis.⁹

We shall reformulate the conditions of Result 2 in terms of a theory of moral permissibility, that is, a theory that, for any feasible set of options, specifies the subset that is permissible. Recall that throughout, we assume Domain Richness, so that every logically possible distribution of benefits is generated by some alternative. We will, however, restrict our attention, in the rest of the essay, to feasible sets with a well-behaved boundary, which implies that there is a maximum benefit for each individual. Consider, then:

Deontic Replication Invariance: For any two alternatives, x and y , and any positive integer n , if x is never (for any feasible set) permissible when y is feasible, then the n -replication of x is never permissible if the n -replication of y is feasible.

Consider, for example, the alternatives $\langle 1,5 \rangle$ and $\langle 2,3 \rangle$. Deontic Replication Invariance requires that if $\langle 1,5 \rangle$ is never permissible when $\langle 2,3 \rangle$ is feasible, then $\langle 1,5,1,5,1,5 \rangle$ is never permissible when $\langle 2,3,2,3,2,3 \rangle$ is feasible.

We believe that this condition is highly plausible. It is satisfied by the deontic version of leximin, utilitarianism, and finitely weighted total prioritarianism—for essentially the same reasons that the corresponding condition on moral betterness, Replication Invariance, is satisfied by these theories (as explained in the previous section).

⁹ The axiological (moral betterness) and deontic (moral permissibility) approaches have been exhaustively studied in the social choice literature, where they are known as the social ordering and social choice function approaches. See, for example, Amartya Sen, "Social Choice Theory: A Re-Examination," *Econometrica* 45 (1977): 53–89, reprinted in Amartya Sen, *Choice, Welfare, and Measurement* (Oxford: Basil Blackwell, 1982), 158–200.

Consider next the following two conditions:

Deontic Ultra Minimal Nonaggregative Priority: There is a pair of benefit levels such that for any alternative in which the benefits of the worst off are below the lower benefit level and the benefits of the best off are above the higher level, and for any benefit gain, there is a sufficiently small benefit loss such that an option is not permissible if, compared to it, some feasible alternative (1) gives the benefit gain to the worst-off individual and leaves her *below average and below the lower benefit level*, and (2) gives the benefit loss to all of the best off and leaves them best off *and above the higher benefit level*.

Deontic Ultra Minimal Aggregation: There is some positive population size such that, for any benefit distribution with a larger population, an option is not permissible if some feasible alternative (1) gives one person less benefits but the difference is small enough, and (2) gives everyone else more benefits and the difference is large enough.

These two conditions are essentially the same as the original conditions (Ultra Minimal Nonaggregative Priority and Ultra Minimal Aggregation). Instead of requiring that one alternative be judged *better* than another, they require that, when both are feasible, one be judged *impermissible*.

In our analysis of moral permissibility, we will not impose any requirement of internal consistency. We will, however, adopt the perspective of *practical moral permissibility*, according to which moral permissibility is purely a matter of doing favorably relative to a feasible set. Hence, we invoke the following condition:

No Prohibition Dilemmas: For any feasible set of alternatives, at least one feasible alternative is permissible.

Of course, moral permissibility can be understood in a more ideal way that does not guarantee that at least one feasible alternative is permissible, but we will focus here on practical moral permissibility.¹⁰

The following is the counterpart of Result 2:

¹⁰ It is important to note that the choice function approach of social choice theory has No Prohibition Dilemmas implicitly built into it. This is because the concept of a choice function—a function that, for any given feasible set, selects the subset consisting of all and only the permissible alternatives relative to that feasible set—is defined to always select a nonempty set. We here posit No Prohibition Dilemmas for practical moral permissibility, but for an argument that prohibition dilemmas are conceptually possible (e.g., not ruled out by deontic logic as such), see Peter Vallentyne, “Prohibition Dilemmas and Deontic Logic,” *Logique et Analyse* 18 (1987): 113–22; and Peter Vallentyne, “Two Types of Moral Dilemmas,” *Erkenntnis* 30 (1989): 301–18.

Result 4: No theory of moral permissibility satisfies No Prohibition Dilemmas, Deontic Replication Invariance, Deontic Ultra Minimal Nonaggregative Priority, and Deontic Ultra Minimal Aggregation.

Let us illustrate the proof with the same example and illustrative assumptions as in the illustration of the proof of Result 2 in Section III.

- (1) Consider the feasible set consisting of Z_0, \dots, Z_{100} , where Z_0, \dots, Z_{100} are defined as in the illustration of the proof of Result 2. By Deontic Ultra Minimal Nonaggregative Priority, for all $t = 0, \dots, 99$, Z_t is not permissible because Z_{t+1} is feasible.
- (2) Let X and Y be defined as in the illustration of the proof of Result 2. By Deontic Ultra Minimal Aggregation, X is not permissible when Y is feasible. By Deontic Replication Invariance, the 100-replication of X is not permissible when the 100-replication of Y is feasible. Given that the 100-replication of X is Z_{100} and the 100-replication of Y is Z_0 , Z_{100} is not permissible, given the feasibility of Z_0 .
- (3) By (1) and (2), no feasible alternative is permissible, which violates No Prohibition Dilemmas.

In the axiological case, we also had an impossibility where Pigou-Dalton was imposed and the priority condition was further weakened to have force only when the gains to the worst off are large enough (e.g., nontrivial). This impossibility also remains in the deontic case.

Deontic Pigou-Dalton: An option is not permissible if, compared to it, some feasible alternative decreases the benefits of one person, increases the benefits of a worse-off person by the same amount, leaves the originally worse-off person no better off than the originally better-off person, and involves no other changes in the distribution.

Deontic Super Ultra Minimal Nonaggregative Priority: There is a pair of benefit levels and a minimum gain size (lower than the lower benefit level) such that for any alternative in which the benefits of the worst-off level are below the lower benefit level and the benefits of the best-off level are above the higher benefit level, and for any benefit gain to a worst-off person that is no less than the minimum gain size, there is a sufficiently small benefit loss to all the best-off persons such that an option is not permissible if, compared to it, some feasible alternative (1) gives the benefit gain to the worst-off individual and leaves her below average and below the lower benefit level, and (2) gives the benefit loss to all of the best off and leaves them best off and above the higher benefit level.

Result 5: No theory of moral permissibility satisfies Deontic Pigou-Dalton, Deontic Replication Invariance, Deontic Super Ultra Minimal Nonaggregative Priority, and Deontic Ultra Minimal Aggregation.

These results establish that, if No Prohibition Dilemmas is imposed, the impossibility result cannot be avoided simply by moving from moral betterness to moral permissibility.

V. AVOIDING MORAL NIHILISM

No theory of moral goodness, or of moral permissibility, satisfies the foregoing seemingly very plausible conditions. This suggests that a plausible account of moral goodness or of moral permissibility is not possible. This is true, however, only if the conditions are plausible. Although we think that each condition seems highly plausible, this plausibility must be reassessed in light of the impossibilities. Because we believe that a plausible theory of morality is possible, we believe that at least one of the conditions is too strong. Here we shall examine some possible ways of weakening the conditions. We should emphasize that we will not endorse any particular way out. Our task is merely to give the reader a sense of possible ways out.

A. Dropping Acyclicity and No Prohibition Dilemmas

We strongly believe that Acyclicity and No Prohibition Dilemmas are basic requirements of moral goodness and practical moral permissibility, so we do not see the weakening or the removal of these requirements as an appealing way out of the impossibility. Still, it may be instructive for the understanding of our result, to briefly study what kind of conceptions of moral goodness or moral permissibility become available if we do not invoke these requirements.

Interestingly, it turns out that a rather prominent model of moral goodness and of moral permissibility, the complaint model proposed by Thomas Nagel and by T. M. Scanlon, satisfies all the other conditions (including the Pareto condition).¹¹ The complaint model relies on the idea that, in a comparison of two alternatives, we should do a pairwise consideration of individual complaints. Here we shall focus on the *binary* complaint model, which will be contrasted in Section V.D with the *global* complaint model. On the binary complaint model, a person's complaint for a given alternative is always relative to another alternative, and is equal to the amount by which her finitely weighted prioritarian benefits (for some specified set of weights) falls short from the largest amount she could get in either of the two alternatives. For example, in comparing $\langle 100, 50 \rangle$ with $\langle 10, 60 \rangle$,

¹¹ Thomas Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), chap. 8; Scanlon, *What We Owe to Each Other*, chap. 5.

the first person has no complaint in $\langle 100, 50 \rangle$, and the second person has a complaint equal to the prioritarian value of 60 minus the prioritarian value of 50. The *axiological binary complaint model* holds that one alternative is morally at least as good as another if and only if the maximum (prioritarian) complaint in the one alternative is no greater than that in the other—and, in case of ties, if and only if the second maximum complaint is no greater, etc.

Nagel and Scanlon proposed this framework as a prioritarian view that avoids the extreme implication of *leximin* that a worse-off person always has priority, no matter how small her benefit. In their view, the better off have priority when their gain is sufficiently greater than the gain of the worse off (the complaints of the better off are greater if the raw benefit shortfall is large enough). So the model satisfies Ultra Minimal Aggregation. However, if the worse-off person has the stronger claim, then the worse off is given priority independently of the number of better-off people that have a smaller claim, which implies that the model satisfies Ultra Minimal Nonaggregative Priority. Finally, it follows trivially that the model satisfies Replication Invariance, since the size of the greatest claim will be invariant to any replication of an alternative.

Similarly, we can show that the *deontic binary complaint model* satisfies the deontic versions of these conditions, where such a model judges an alternative permissible if and only if no feasible alternative has a smaller maximum complaint, or, in cases of ties, if no feasible alternative has a smaller second largest complaint, etc.

To illustrate the axiological binary complaint model, let us assume that the priority weighted benefits are measured by the square root of the benefits.¹² In comparing two alternatives, then, the magnitude of each person's complaint in a given alternative is equal to the difference between the square root of her maximum benefit in the two alternatives and the square root of her benefits in the given alternative. So, for example, if we compare $\langle 1, 25 \rangle$ and $\langle 4, 9 \rangle$, the respective complaints are $\langle 1, 0 \rangle$ and $\langle 0, 2 \rangle$. Thus, the greatest complaint is that of the second person in $\langle 4, 9 \rangle$. Her shortfall of 16 carries greater weight than the first person's shortfall of 3 in $\langle 1, 25 \rangle$ —even though the second person is better off than the first person. Thus, the axiological binary complaint model judges $\langle 1, 25 \rangle$ as morally better than $\langle 4, 9 \rangle$ (and favors the better-off person). As a second example, compare $\langle 4, 100 \rangle$ with $\langle 16, 81 \rangle$. Here, their respective complaints (i.e., shortfalls in the square roots of the benefits) are respectively $\langle 2, 0 \rangle$ and $\langle 0, 1 \rangle$. The greatest complaint is thus that of the first person in $\langle 4, 100 \rangle$. Thus, this complaint model judges $\langle 16, 81 \rangle$ as morally better (and favors the worse-off person even though her benefit gain of 12 in $\langle 16, 81 \rangle$ is less than the gain of 19 to the second person in $\langle 4, 100 \rangle$).

¹² See, for example, Alex Voorhoeve, "Should Losses Count? A Critique of the Complaint Model," *LSE Choice Group Working Papers*, vol. 2 (2006).

The axiological binary complaint model, unfortunately, violates Acyclicity. The square-root version, for example, judges $\langle 1,1,15 \rangle$ as better than $\langle 1,0,17 \rangle$ —since in the former the largest complaint is much less than 1 (that of the third person) and in the latter it is 1 (that of the second person). The model also judges $\langle 1,0,17 \rangle$ as better than $\langle 0,0,25 \rangle$ —since in the former the largest complaint is less than 1 (that of the third person) and in the latter it is 1 (that of the first person). Finally, this complaint model also judges $\langle 0,0,25 \rangle$ as better than $\langle 1,1,15 \rangle$ (the first option above)—since in the former the largest complaint is 1 (that of the first and second persons) and in the latter it is slightly more than 1 (that of the third person). Thus, the square-root version of the axiological binary complaint model violates Acyclicity.

Similarly, if we assume that this sequence constitutes the feasible set of alternatives, it follows straightforwardly that the deontic version of the binary complaint model will not satisfy No Prohibition Dilemmas.

Thus, if Acyclicity and No Prohibition Dilemmas are dropped, the binary complaint model satisfies all the remaining conditions. As indicated above, however, these two conditions are basic conditions of moral goodness and practical moral permissibility, and thus this is not a promising way of avoiding the impossibility result.

B. Dropping Replication Invariance

We have already shown that if we drop Replication Invariance, there are conceptions of moral goodness that satisfy the remaining principles. The geometric Gini moral-betterness relation satisfies both Ultra Minimal Aggregation and Ultra Minimal Nonaggregative Priority, even in conjunction with Weak Pareto and Transitivity. We find Replication Invariance quite compelling, and therefore we doubt that the removal or weakening of this condition could be a promising way out of the dilemma.

Note, however, how the geometric Gini manages to satisfy Ultra Minimal Nonaggregative Priority in spite of being clearly of the aggregative kind. This is because with exponential weights geometric Gini's aversion to inequality (or degree of priority to the worse off) increases with the size of the population. This might make sense, even if it is not a popular idea in the literature. A certain degree of inequality might appear more alarming in a large population than in a small population, because it means that more people are suffering from a disadvantaged position.

C. Weakening the aggregation condition

Ultra Minimal Aggregation requires that, for a sufficiently large population, imposing a sufficiently small loss on *anyone* and giving everyone else a sufficiently large gain makes things morally better. It requires this, even if the one loser is desperately poor and the gainers are all super-

affluent. One might well question whether things are made morally better in such a case. Clearly, some kind of aggregation is appropriate (we shouldn't focus solely on the worst off), but this kind of aggregation, given the other conditions, prevents giving sufficient priority to the worse off.

Consider, then, the following weakening of the aggregation condition:

Super Ultra Minimal Aggregation: There is some benefit level and some positive population size such that, for any benefit distribution with a larger population, imposing a loss on one person and gains on everyone else makes things morally better, if the loss is small enough, the gains are large enough, and those who gain remain below the specified benefit level.

This is exactly like the original Minimal Aggregation condition, except that it does not apply if the gainers are above some fixed benefit level (e.g., above a minimally adequate benefit level). It requires some kind of aggregation of benefits for those below the fixed level, but does not require the aggregation of benefits of those below that level with those above. This is in the same general spirit as John Rawls's idea of maximizing the aggregate benefits of the worst-off group (as opposed to the benefits of the worst-off individual).¹³ The impossibility result disappears if we so weaken the aggregation condition. Indeed, it disappears even if we impose Transitivity (rather than the weaker Acyclicity), Minimal Nonaggregative Priority (rather than the weaker Ultra Minimal version), Pigou-Dalton, and the following strong version of the Pareto condition:

Strong Pareto: For any two alternatives x and y , if each person has at least as great benefits in x as in y and some person has greater benefits in x than in y , then x is morally better than y .

Strong Pareto has force when only some individuals have strictly greater benefits in x than in y (whereas Weak Pareto has force only when everyone has strictly greater benefits). Strong Pareto is still extremely plausible.

We now note the following possibility result:

Result 6: There is a moral-betterness relation that satisfies Strong Pareto, Pigou-Dalton, Replication Invariance, Transitivity, Minimal Nonaggregative Priority, and Super Ultra Minimal Aggregation.

Any member of the family of *modified total poverty shortfall relations* satisfies all these conditions. Such prioritarian betterness relations are defined as follows: They specify some fixed benefit level that defines the poverty line (or minimum benefit level for a decent life). For each indi-

¹³ John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), 98.

vidual, her modified shortfall is defined as the shortfall in finitely weighted prioritarian benefits (for some specified set of weights) from the poverty line (and zero if not below that level). (For example, if the poverty line is 4 and the priority weighted benefits are the square root of the benefits, then an individual who gets 1 unit of benefit has a modified shortfall of 1 [square root of 4 minus square root of 1].) One alternative is judged morally at least as good as another just in case (1) its total shortfall is less than that of the other, or (2) its total shortfall is equal to that of the other, and its finitely weighted prioritarian total benefits (for some specified set of weights) is at least as great. (The second clause is needed to ensure that Strong Pareto is satisfied: that giving a benefit to someone above the poverty line makes things better. The appeal to finitely weighted prioritarian benefits ensures that the relation satisfies Pigou-Dalton.)

Modified total poverty shortfall relations satisfy all the conditions listed in Result 6. Due to space limitations, we shall here address only the aggregation condition. These relations violate Ultra Minimal Aggregation but satisfy Super Ultra Minimal Aggregation. They violate the former condition, because they give absolute priority to someone below the poverty line over any number of people getting any sized benefit above the poverty line (since benefits to those below the line reduce the total poverty shortfall but benefits to those above it do not). Modified total poverty shortfall relations, however, satisfy Super Ultra Minimal Aggregation. As long as those receiving the benefits are below the poverty line, these relations hold that giving enough people a large enough benefit and only one person a smaller loss makes things morally better (since the total poverty shortfall will be reduced).¹⁴ Result 6 establishes that it is possible to combine nonaggregative priority with the very weak kind of aggregation required by Super Ultra Minimal Aggregation. This possibility also applies in the deontic case (with a deontic version of the weaker aggregation condition), but we shall limit our attention to the axiological case (of moral betterness).

One possible way out of the impossibility result, then, is to weaken the aggregation condition so as not to require that benefits to those above some threshold can outweigh (under the specified conditions) benefits to those below. This allows that any loss to people above the threshold can

¹⁴ Super Ultra Minimal Aggregation avoids the impossibility result because the condition is applicable only when all the gainers are below some threshold (e.g., none are superaffluent). Another (similar) way of weakening the aggregation condition to avoid an impossibility result is to restrict its application to cases in which the one loser is above some threshold (e.g., a poverty threshold). The idea is that aggregation is required among well-off people, but, below the threshold, absolute priority to the worst-off may prevail. When such a weakening of the aggregation condition is introduced, a possibility result obtains with a *modified total affluence relation*, which deems x to be at least as good as y if and only if (1) the leximin criterion restricted to positions below the threshold prefers x , or (2) x and y are leximin indifferent with respect to positions below the threshold and x has finitely weighted total prioritarian benefits (for some specified set of weights) that are at least as great as those of y .

be outweighed by an insignificant gain to one person below the threshold, as long as this loss does not bring the losers below the threshold. For example, a modified total poverty shortfall relation holds that things are made morally better by (1) giving one person who is slightly below the threshold a very small benefit that leaves her below the threshold, combined with (2) reducing everyone else's benefits from significantly above the threshold to just at the threshold. Such a view is controversial, since many people believe that some aggregation should be accepted between the benefits of those below and those above the threshold. As we have shown, however, it is impossible to combine a requirement of minimal aggregation with a condition assigning minimal nonaggregative priority to the worst off. So if one insists on allowing for some aggregation at all levels, then the only option is to consider a weakening of the priority condition.

D. Weakening the priority condition

Super Ultra Minimal Nonaggregative Priority requires that a large enough benefit to a *worst-off* individual, who is below a poverty threshold and remains below the average, has priority over a sufficiently smaller loss to all of the best-off individuals above an affluence threshold, no matter how many best-off people there are. It gives nonaggregative priority to the worst off in the sense that the number of best-off gainers is irrelevant.

It is, of course, possible to weaken the priority condition so that it does not require nonaggregative priority. Finitely weighted prioritarianism, for example, gives priority to the worse off but in an aggregative way (since enough benefits to enough best-off people can outweigh any benefit to the worst off). We believe, however, that some kind of nonaggregative priority for the worst off is plausible. Therefore, in what follows, we shall focus on weakening the priority condition while still requiring some kind of nonaggregative priority.

For the axiological case (moral betterness), we have not been able to identify any promising weakening of the priority condition that both preserves a significant nonaggregative priority for the worst off and avoids the impossibility result. We have, however, identified such a weakening for the deontic case (moral permissibility).

Deontic Super Ultra Minimal Nonaggregative Priority holds that the presence of an option that favors the worst off in a certain way over a second option is sufficient for the second option to be impermissible. One might think, however, that the force of the priority for the worst off should be sensitive to what other options are feasible. For example, one might think that, in a given choice situation, the moral importance of giving a benefit to a given individual should depend in part on how far short it leaves her from the best that she could achieve. Thus, one might think (1) that $\langle 1,0,4 \rangle$ is not permissible when the only feasible alternative

is $\langle 1,2,3 \rangle$ (since the worst-off individual in the former has a shortfall of 2 and the best-off individual in the latter has a shortfall of only 1), but (2) that $\langle 1,0,4 \rangle$ may be permissible when $\langle 0,0,9 \rangle$ is also feasible (since now the best off in $\langle 1,0,4 \rangle$ has a shortfall of 5 units from her maximum, and this may offset the shortfall of 2 [relative to $\langle 1,2,3 \rangle$] of the worst off).

In light of this, it might be argued that the priority condition should be weakened so as to apply only when there are two feasible options. The following condition is identical to Deontic Ultra Minimal Nonaggregative Priority, except that it applies only when there are only two feasible alternatives:

Deontic Two-Option Ultra Minimal Nonaggregative Priority: There is a pair of benefit levels such that, for any alternative in which the benefits of the worst off are below the lower benefit level and the benefits of the best off are above the higher level, and for any benefit gain, there is a sufficiently small benefit loss such that an option is not permissible *if the only feasible alternative* is one that, compared to it, (1) gives the benefit gain to the worst-off individual and leaves her below average and below the lower benefit level, and (2) gives the benefit loss to all of the best-off individuals and leaves them best off and above the higher benefit level.

This condition leaves open what kind of priority should be given to the worst off when there are more than two feasible options. It merely insists that the requisite nonaggregative priority is given in the simple case where there are only two feasible options. This is enough, however, to rule out finitely weighted prioritarianism. If there are enough best-off individuals, this theory will judge an option permissible when the only feasible alternative gives a benefit gain to the worst-off individual in the case defined by the above condition. The condition is weak enough, however, to avoid the impossibility result. Indeed, the impossibility result is avoided even if we take the original axiological Minimal Nonaggregative Priority condition and convert it to a deontic condition but restrict it to the two-option case. Consider then:

Deontic Two-Option Minimal Nonaggregative Priority: For any pair of benefit levels, for any alternative in which the benefits of the worst off are below the lower benefit level and the benefits of the best off are above the higher level, and for any benefit gain, there is a sufficiently small benefit loss such that an option is not permissible *if the only feasible alternative* is one that, compared to it, (1) gives the benefit gain to the worst-off individual and leaves her below average and below the lower benefit level, and (2) gives the benefit loss to all of the best off and leaves them best off and above the higher benefit level.

This condition entails the previous condition. We now note that the impossibility result is avoided if we impose this “two-option” condition instead of Deontic Ultra Minimal Nonaggregative Priority—even if Strong Pareto and Pigou-Dalton are also imposed.

Result 7: There are theories of moral permissibility that satisfy No Prohibition Dilemmas, Strong Pareto, Pigou-Dalton, Deontic Replication Invariance, Deontic Two-Option Minimal (and Ultra Minimal) Nonaggregative Priority, and Deontic Ultra Minimal Aggregation.

We shall illustrate the possibility of such theories by appealing to the deontic complaint model, which judges a feasible alternative permissible if and only if no other feasible alternative has a smaller maximum complaint, or, in cases of ties, no smaller second-largest complaint, etc. Above, we considered the *binary* version of this model, for which a person’s complaint in a given option is always relative to some second specified option. We showed that that version of the complaint model violates No Prohibition Dilemmas. Here, we shall consider the *global* version of the model, for which a person’s complaint in a given option is relative to a specified *feasible set*. More exactly, relative to a given feasible set, a person’s complaint in a given option is equal to the shortfall in priority weighted benefits from the most that the individual can get *in that feasible set*. For example, if the feasible set consists of $\langle 5,3 \rangle$, $\langle 3,4 \rangle$, and $\langle 10,0 \rangle$, then the first person’s complaint in $\langle 5,3 \rangle$ is equal the shortfall in prioritarian benefits from 10. Unlike the binary model, the global model focuses solely on the complaint relative to the best option in the feasible set. For example, the fact that the first person has no complaint in $\langle 5,3 \rangle$ relative to $\langle 3,4 \rangle$ (since she gets 5 rather than 3) is deemed irrelevant.

Unlike the binary version, the global version of the deontic complaint model satisfies No Prohibition Dilemmas. There is always some option that minimizes the greatest complaints (etc.) relative to the feasible set.

The deontic global complaint model, however, does not satisfy Deontic Ultra Minimal Nonaggregative Priority. To see this, consider the following feasible set, consisting of twenty options, where each option has twenty poor people (to the left of the semicolon) and one rich person (to the right of the semicolon). In the sequence below, each successive alternative increases the benefits of two individuals from 0 to 1, and decreases the benefits of the best-off individual by half a unit.

$\langle 0,0, \dots .0;25 \rangle$
 $\langle 1,1,0,0, \dots .0;24.5 \rangle$
 $\langle 1,1,1,1,0,0, \dots .0;23 \rangle$
 \dots
 $\langle 1,1, \dots .1,1,0,0;16 \rangle$
 $\langle 1,1, \dots .1,1;15.5 \rangle$

Suppose, for illustration, that the priority condition applies when (1) the least well off is 2 or below and the best off is 10 or above, (2) there are at least two least-well-off people who gain, and (3) for a gain of 1 to the least well off, the best off lose 0.5 or less. Given these assumptions, Deontic Ultra Minimal Nonaggregative Priority requires that all options but the last be judged impermissible (on the basis of a comparison with any subsequent option). The deontic global complaint model, however, can judge one of these options (other than the last) permissible. If, for example, prioritarian benefits are equal to the square root of the benefit, then the model judges the second-to-last option permissible. This option minimizes maximum complaints (etc.). It has two poor people with complaints of 1 (square root of 1 minus square root of 0) and one rich person with a complaint of 1 (square root of 25 minus square root of 16), and everyone else has 0 complaints. The last option has greater maximum complaint, since the rich person has a complaint greater than 1 (the square root of 25 minus the square root of 15 is greater than 1). All other options have more than three people with complaints of 1. Thus, the second-to-last option is judged permissible, which violates Deontic Ultra Minimal Nonaggregative Priority.

The deontic global complaint model does, however, satisfy the weaker Two-Option Minimal (and Ultra Minimal) Nonaggregative Priority. For any given option, providing a benefit to a worst-off person and a small enough loss to the best-off individuals (no matter how many there are) produces a distribution with a smaller largest complaint. Thus, the original option will be judged impermissible, if these are the only two feasible options.

The global complaint model also satisfies all the remaining deontic conditions and thus establishes the foregoing possibility result (Result 7). Have we, then, finally found a significant possibility for moderate nonaggregative priority for the worse off? There are at least two ways of doubting this. One is the thought that the conditions imposed for Result 7 do not leave room for a significant form of nonaggregative priority. After all, the two-option priority conditions discussed above are completely silent about priority for the worse off when there are more than two options. The two-option priority conditions are indeed extremely weak. However, the fact that the global complaint model is nonaggregative much more generally indicates that it should be possible to strengthen the priority condition and still keep the possibility result. We have not yet established any interesting strengthening, however, and thus we leave this issue for future research.

A second way of doubting that we have found a significant possibility for moderate nonaggregative priority for the worse off is to claim that the impossibility result reappears when further plausible conditions are introduced. One such condition, deemed plausible by many, is the following contraction consistency condition (concerning how the permissibility of an option is affected when the feasibility set is contracted):

Alpha: If an alternative is permissible relative to a given feasible set, then it is also permissible relative to any subset containing the alternative.

This condition is violated, for example, if x is judged permissible relative to $\{x,y,z\}$, but is not judged permissible relative to $\{x,y\}$. If x is a winner against y and z , why wouldn't it also be a winner against y alone?

Alas, the possibility result disappears if Alpha is imposed:

Result 8: No theory of moral permissibility satisfies No Prohibition Dilemmas, Alpha, Deontic Replication Invariance, Deontic Two-Option Ultra Minimal Nonaggregative Priority, and Deontic Ultra Minimal Aggregation.

To illustrate the proof, let us see how the deontic global complaint model violates Alpha (while satisfying all the remaining conditions). Consider again the example used to illustrate the proof of Result 7. We showed how a deontic global complaint model based on the square-root weighting of benefits judges the second-to-last option, $\langle 1,1, \dots .1,1,0,0;16 \rangle$, permissible. Consider now the case where only the last two options— $\langle 1,1, \dots .1,1,0,0;16 \rangle$ and $\langle 1,1, \dots .1,1;15.5 \rangle$ —are feasible. In this case, $\langle 1,1, \dots .1,1,0,0;16 \rangle$ will not be permissible, because now the complaints are established solely in a comparison with $\langle 1,1, \dots .1,1;15.5 \rangle$ (instead of in comparison with the greatest feasible benefit for an individual in the earlier larger feasible set). And in this case, the greatest complaint in $\langle 1,1, \dots .1,1,0,0;16 \rangle$ is 1 (for the two people with 0), whereas it is less than 1 in $\langle 1,1, \dots .1,1;15.5 \rangle$ (namely, 0.5 for the person with 15.5). Hence, only $\langle 1,1, \dots .1,1;15.5 \rangle$ is permissible. Thus, $\langle 1,1, \dots .1,1,0,0;16 \rangle$ is judged permissible from the larger feasible set, but not from the two-option subset, which violates Alpha.

Given that the imposition of Alpha reintroduces the impossibility result, a crucial question is whether Alpha is a plausible condition. Two of us (Fleurbaey and Tungodden) believe that it is and thus deny that weakening the priority condition to the two-option version avoids the impossibility result. However, one of us (Vallentyne) believes that Alpha should be rejected. The reasons have been given elsewhere, and we will therefore only illustrate the concern here.¹⁵

Alpha is clearly a desirable condition for permissibility. The question, however, is whether it is a *mandatory* condition—one that any minimally adequate conception of permissibility *must* satisfy. To see just how strong this condition is, it suffices to note that it is violated by several standard kinds of theory. Consider, for example, a satisficing theory that judges an

¹⁵ Bertil Tungodden and Peter Vallentyne, "On the Possibility of Paretian Egalitarianism," *Journal of Philosophy* 102 (2005): 126–54.

option to be permissible, relative to a feasible set, if and only if it is not below the average (expected) value of options in the feasible set. This violates Alpha, since, if enough below-average options are removed from the feasible set, a given option can fall below the average value and thus cease to be judged permissible. There are other minimally plausible deontic theories that satisfy Alpha—such as the maximin relative benefit bargaining theory of David Gauthier, the regret theory of Graham Loomes and Robert Sugden, and Paretian egalitarianism.¹⁶ More generally, Sugden, Amartya Sen, Edward F. McClennen, and Frances Kamm have argued against the *a priori* imposition of Alpha.¹⁷ Thus, it is not obvious that Alpha is a plausible *a priori* condition on permissibility.

Because we, the authors, are not in agreement about the status of Alpha, we must leave open whether there is a way of avoiding the impossibility result by weakening the priority condition while preserving its nonaggregative nature.

VI. CONCLUSION

We have been exploring nonabsolute forms of prioritarianism. Utilitarianism fails to give any priority to the worst off, and leximin gives absolute priority to the worst off. One possibility is finitely weighted prioritarianism, but we have suggested that it fails to give appropriate priority to the worst off. For it allows that any benefit, no matter how large, to a worst-off person in abject poverty can be outweighed by an arbitrarily small benefit to enough extremely affluent people. We have been seeking a form of prioritarianism that does not allow such aggregative trade-offs. We have shown, however, that any moderate nonaggregative prioritarian theory of moral betterness or of practical moral permissibility must violate some seemingly plausible conditions.

If one is to avoid moral nihilism, one or more of the seemingly plausible conditions must be relaxed. We have suggested that (1) the aggregation condition can be relaxed to open the possibility for moral views based on minimizing the total prioritarian shortfall from the poverty level, and (2) the nonaggregative priority condition can be relaxed

¹⁶ David Gauthier, *Morals by Agreement* (London: Oxford University Press, 1986); Graham Loomes and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal* 92 (1982): 805–24; Tungodden and Valletyne, "On the Possibility of Paretian Egalitarianism."

¹⁷ Robert Sugden, "Why Be Consistent? A Critical Analysis of Consistency Requirements in Choice Theory," *Economica*, new series, vol. 52 (1985): 167–83; Amartya Sen, "Internal Consistency of Choice," *Econometrica* 61 (1993): 495–521, reprinted in Amartya Sen, *Rationality and Freedom* (Cambridge, MA: Harvard University Press, 2002), 121–57; Edward F. McClennen, *Rationality and Dynamic Choice* (Cambridge: Cambridge University Press, 1990), 182–84, 240, 251; Frances Kamm, *Mortality, Mortality II* (New York: Oxford University Press, 1996), 343–44; Frances Kamm, *Intricate Ethics* (Oxford: Oxford University Press, 2007), 169–73, 222 n. 16, 298, 484–487.

to open the possibility, if no requirement of consistency is imposed, for the deontic global complaint model, which seeks to minimize maximum prioritarian complaints.

Obviously, many people reject the need for some aggregation or the need for nonaggregative priority. We have not attempted to defend these assumptions. Instead, we have sought to explore the range of possibilities that they leave open. Our main conclusion is that this range is not great and that more investigation is needed in order to determine whether there is a plausible form of moderate nonaggregative prioritarianism.

Ethics and Economics, Université Paris Descartes

Economics, Norwegian School of Economics and Business Administration

Philosophy, University of Missouri–Columbia