

A Philosopher's Guide to Multidimensional Equality
Kristi A. Olson
Forthcoming in

Abstract: Take a distribution involving two or more dimensions—say, apples and oranges. If Smith fares better on one dimension, while Jones fares better on another, how do we compare them? Rawlsians call this the indexing problem: how do we identify the worst-off group overall? Proponents of the capabilities approach call it the aggregation problem: how do we weight the various dimensions? This essay examines possible solutions, including interpersonal comparisons of utility, the envy test, the egalitarian equivalent approach, undominated diversity, and the solidarity test. Although this essay does not solve the problem, it does identify the crucial question: should we equalize each person's bundle of resources or capabilities, the extent to which each person's bundle fits her preferences, or standing in the community? Since the question is normative—what should we equalize?—philosophers have much to contribute to the debate, and the payoff could be considerable: both policy decisions and other egalitarian debates could hinge on the answer.

Key words: indexing problem, aggregation, multidimensional equality, capabilities approach, envy test, egalitarian equivalent approach,

A Philosopher's Guide to Multidimension

as the list includes more than one, capability egalitarians face the : how should the various dimensions be weighted? (See, e.g., Robeyns, 2017; Arneson, 1990a).

These questions are not of mere theoretical interest. Policy decisions could depend on the answers. Philosophers and policymakers, then, need some way to compare multidimensional bundles. And, although tempting, deferring to public deliberation is not the solution; rather, deferring to public deliberation merely shifts the problem. W

is that the envy test does not deliver the desired information. If a distribution is not envy-free, the envy test tells us that the distribution is unequal. Yet, except in very simple scenarios (such as when only one person has envy), the envy test cannot identify the worst-off individual. For example, if Smith prefers Jones's bundle while Jones prefers Smith's bundle, the envy test cannot tell us who is worse off overall. Similarly, if Smith prefers Jones's bundle and Patel prefers Murphy's bundle but neither Smith nor Patel prefers the other's bundle, the envy test cannot identify the worst-off person. Moreover, except again in very simple scenarios, the envy test cannot rank scenarios according to the extent of inequality.

In response to this problem, we could modify the envy test to consider the amount of envy (see, e.g., Chaudhuri, 1986). For example, we might say that if Smith prefers Jones's bundle more than Jones prefers Smith's bundle, then Smith has more envy than Jones. We could then identify the worst-off person as the one with the most envy, and we could rank scenarios according to the extent of envy. Yet how exactly are we to measure the strength of their preferences? At first glance, the task might appear relatively simple. We could ask each person with envy how much of a certain resource added to her bundle would eliminate her envy: the more of that resource she needs to achieve indifference, the greater her envy (see Fleurbaey, 2008). But which resource do we use? If Jones and Smith diverge in how many apples they would trade for one orange, the amount of envy—and hence who is worse off overall—might change depending on whether we use apples or oranges as our metric (see Olson, 2018). I say more about this problem in Part III. For our purpose here, it suffices to point out that we have reintroduced a version of the calibration problem: we are, after all, attempting to measure and compare envy interpersonally.

According to the second objection, to the extent the envy test is used to say that one person is worse off than another, the envy test delivers verdicts at odds with ordinary properties of relations. Consider, for example, the following uncontroversial claim: if A is better all-things-considered than B, then B is worse off than A. Consider now the following claim: if A is better all-things-considered than B, then B is better off than A. This claim is false. Yet the envy test delivers verdicts at odds with ordinary properties of relations. Consider, for example, the following uncontroversial claim: if A is better all-things-considered than B, then B is worse off than A. Consider now the following claim: if A is better all-things-considered than B, then B is better off than A. This claim is false. Yet the envy test delivers verdicts at odds with ordinary properties of relations.

mutual—Smith can prefer Jones's bundle while Jones prefers Smith's bundle—the envy test can lead us to conclude that Smith is better off than Jones. Jones is better off than Smith, not merely in one respect but overall. Or consider this uncontroversial claim: If A is better than B and B is better than C, then A is better than C. Yet, since envy is not transitive, the envy test also violates this claim. For example, if Smith prefers Jones's bundle, Patel prefers Smith's bundle, but Patel does not prefer Jones's bundle, the envy test would conclude that Jones is better off than Smith and Smith is better off than Patel, but Jones is not better off than Patel. Thus, we can use the envy test to compare multidimensional bundles only if we are willing to abandon ordinary properties of relations.

III. THE EGALITARIAN EQUIVALENT APPROACH

Consider next the egalitarian equivalent approach (Pazner and Schmeidler, 1978; Fleurbaey, 2008, 2012a, 2012b; Fleurbaey and Blanchet, 2013; cf. Van Parijs, 1993). The envy test, recall, asks each person to compare her bundle to every other person's bundle. The egalitarian equivalent approach, in contrast, asks each person to compare her bundle to a designated reference bundle. To illustrate, suppose Smith has {one apple, two oranges} while Jones has {two apples, one orange}. If each is indifferent between her bundle and the designated reference bundle—say {three oranges}—then there is no inequality between them. In contrast, if Smith prefers his own bundle while Jones prefers the reference bundle, then Jones is deemed worse off: Smith's bundle I P

reference bundle will depend on the metric used to measure their indifference point. Suppose we ask each: How many _____ added to your bundle would make you indifferent between your bundle and the reference bundle? In this scenario, there is no inequality: Smith and Jones both require one additional orange to be indifferent between their bundle and the reference bundle. On the other hand, suppose we ask: How many _____ added to your bundle would make you indifferent between your bundle and the reference bundle? In this scenario, Jones will be deemed worse off: Jones requires two apples, whereas Smith requires only one apple. As a result, with one metric, there is no inequality; with another metric, there is an inequality.

We might attempt to sidestep the problem above by using money as the metric. Money, after all, can be used to buy both apples and oranges, and it is neutral between them (even though more money might be required for one than for the other). Indeed, we might be tempted to assign each bundle a dollar value and forgo the egalitarian equivalent approach altogether. Yet, in many of the distributions of interest to philosophers and policymakers, some resources or capabilities cannot be bought or sold. As a result, the choice of money as a metric still requires justification. Consider, for example, bundles of money and health: If Smith is rich but in poor health, while Jones is poor but in good health, who is worse off overall? If Smith and Jones value money and health differently—i.e., if they would make different tradeoffs between them—then the choice of money as the metric to measure indifference can affect not only whether there is an inequality but also who is disadvantaged and by how much. Of course, we cannot (directly) redistribute health, whereas we can directly redistribute money. But keep in mind that we are here just asking each person a hypothetical question: how much money added to your bundle would make you indifferent between your bundle and the reference bundle? No actual redistribution takes place. As a result, proponents of the egalitarian equivalent approach cannot justify money as the metric by pointing to the fact that money can be more readily redistributed.

Moreover, even if proponents of the egalitarian equivalent approach could justify their choice of reference bundle and their choice of metric, there is another objection. Namely, the egalitarian equivalent approach will sometimes deliver verdicts at odds with our ordinary judgments. Consider first the ordinary judgment that if two people have identical bundles, there is no inequality between them. Yet, as we saw above, the egalitarian equivalent approach will sometimes report an inequality. Consider second a distribution in which one person's bundle is worse in at least one respect and better in no respect than another person's bundle. Suppose, for example, that Jones has {one apple, two oranges}, while Smith has {two apples, two oranges}. Comparisons of such bundles are commonly assumed to be easy: Smith is better off than Jones (see, e.g., Wolff and de-Shalit, 2007). Yet, the egalitarian equivalent approach will sometimes deliver the opposite verdict. To illustrate, if, in the scenario above, the reference bundle is {four apples}, Smith (who values apples and oranges equally) will be indifferent between his bundle of {two apples,

undominated diversity), Jones nonetheless has envy (failing the envy test). Second, even when the envy test and undominated diversity reach the same verdict, they take different paths.

Undominated diversity appeals to everyone's preferences; no special weight is given to the preferences of the person assigned the bundle. The envy test, in contrast, caters to the preferences of the person assigned the bundle.

Of the approaches considered thus far, undominated diversity has several advantages. First, unlike the utility approach, undominated diversity avoids interpersonal comparisons of utility. Second, unlike the envy test, undominated diversity respects ordinary properties of relations. Third, unlike the egalitarian equivalent approach, undominated diversity respects our ordinary judgments about inferior bundles.

Of course, whether the last point counts as an advantage depends on the normative importance of fit. Indeed, if fit is taken into consideration, the plausibility of the undominated diversity approach falters. To illustrate, even if everyone prefers bundle X to bundle Y, it might be the case that Jones ranks bundles X and Y first and second (out of, say, 100 bundles), while Smith ranks bundles X and Y second to last and last. As a result, the unanimously dispreferred bundle might be a better fit for Jones than the unanimously preferred bundle is for Smith: Jones, after all, receives the bundle she ranks second, while Smith receives the bundle he ranks 99th.

Here I want to consider two (other) objections to undominated diversity. First, except in very simple cases, undominated diversity—like the envy test—cannot identify the worst-off person. For example, if everyone prefers bundle A to bundle B and bundle C to bundle D, then we can say that the persons with bundle B and bundle D are worse off than the persons with bundle A and bundle C, respectively. Nonetheless, if there are no unanimous preference rankings of bundle B and bundle D, we cannot say who—the person with bundle B or the person with bundle D—is

unanimously dispreferred bundles, the greater the inequality. After all, a distribution in which one person has \$100 dollars while everyone else has \$99 will generate the same number of unanimously dispreferred bundles as a distribution in which one person has \$100 dollars and everyone else has \$1.

Of course, we could supplement undominated diversity by asking how much compensation added to a bundle would change the marginal person's preferences such that the bundle is no longer unanimously dispreferred: the greater the compensation required, the greater the inequality. Any such modification, however, requires a metric: we must decide which resource—e.g., apples, oranges, money, or health—to use to measure the marginal person's indifference point. Yet, as we saw in Part III, if people make different tradeoffs between these resources, the choice of metric can affect the extent of inequality. As a result, the choice of metric would again require justification.

According to the second objection, undominated diversity is too weak (or too incomplete). Consider, for example, a bundle that is dispreferred by all but one of 100 people. Intuitively, such a bundle is inferior. Yet, proponents of undominated diversity cannot say so. In response to this objection, a proponent of undominated diversity could simply bite the bullet and insist that there is no inequality (or, at least, insist that the bundles are incommensurable). Second, she could attempt to rule out outlier preferences (see, e.g., Van Parijs, 1995; Arneson, 1990a). Of course, any such culling of preferences would need to be justified.

Third, she could relax the unanimity requirement. To be sure, we cannot say that bundle X is better than bundle Y whenever a simple majority prefers bundle X, on pain of cycles (Van Parijs, 1995). To illustrate, if Smith prefers X to Y to Z, Jones prefers Y to Z to X, and Patel prefers Z to X to Y, then a majority prefers X to Y, a different majority prefers Y to Z, and yet another majority prefers Z to X. As a result, the person with bundle X is better off than the person with bundle Y and, by transitivity, the person with bundle Y is better off than the person with bundle X, in violation of

ordinary properties of relations. Nonetheless, we could replace the unanimity requirement with, say, a near unanimity requirement. Iturbe-Ormaetxe and Nieto (1996) and Fleurbaey (2008) consider related proposals. These proposals, however, forfeit the appeal of unanimity.

In Part V, I introduce another way to make undominated diversity more demanding, while still appealing to everyone's preferences.

V. THE SOLIDARITY TEST

Instead of asking whether any bundle is unanimously dispreferred, the solidarity test asks: Is an envy-free distribution of these bundles compatible with everyone's preferences? (Olson, 2018, 2020a, 2020b). The solidarity test bears an obvious relationship to the envy test: whenever the envy test is satisfied, the solidarity test is also satisfied. Yet, unlike the envy test, the solidarity test does not ask whether an envy-free distribution is \geq everyone's preferences. To illustrate, if Smith prefers {two apples, one orange} while Jones prefers {one apple, two oranges}, then an envy-free distribution of these bundles is compatible with everyone's preferences. The solidarity test is thus satisfied—and, moreover, it is satisfied even if Smith and Jones receive their dispreferred bundles. Thus, the solidarity test can be satisfied even when the envy test is not satisfied. If, on the other hand, Smith and Jones both strictly prefer {two apples, one orange}, an envy-free distribution of these bundles is incompatible with everyone's preferences and thus neither the envy test nor the solidarity test is satisfied.

The solidarity test has several advantages over the earlier approaches. First, unlike the utility approach, the solidarity test does not require interpersonal comparisons of utility. Second, as I explain below, unlike the envy test, the solidarity test will not deliver verdicts at odds with ordinary properties of relations. Third, unlike the egalitarian equivalent approach, the solidarity test will not deliver verdicts at odds with our judgments of inferior bundles. (Of course, whether this counts in favor of the solidarity test depends on the normative importance of fit.) And fourth, like undominated diversity, the solidarity test appeals to everyone's preferences. Yet, since the

solidarity test is more demanding than undominated diversity, it is less vulnerable to the weakness objection.¹

Nonetheless, the solidarity test is also vulnerable to objections. Here I consider two. The first is a familiar one. Except in very simple scenarios, the solidarity test cannot identify the worst-off person or rank scenarios according to the extent of inequality. And any attempt to modify the solidarity test such that it could deliver these results would require a metric and hence would be

tracks, consider two scenarios. In the first scenario, Jones receives a bundle that she would not choose but someone else in the distributive community would choose. In the second scenario, Jones receives a bundle that she (including herself) would choose. The second scenario strikes some as normatively objectionable in a way that the first scenario is not. But what exactly is the difference? The difference cannot be about Jones's utility, her bundle, or her bundle's fit with her preferences. After all, we can hold those constant between the two scenarios. According to proponents of the solidarity test, the difference instead concerns Jones's standing in her community. When the bundles are incompatible with an envy-free distribution—and when this could have been avoided by a different allocation of resources or capabilities to bundles—the recipients of the bundles do not stand as equals. In response to the question above—why does it matter if Smith prefers Jones's bundle if Jones does not?—

some combination of these (see, e.g., Sen, 1979; Dworkin, 1981a, 1981b; Cohen, 1989). Nor am I referring to the point of equality debate, which asks whether the fundamental egalitarian concern should be about the distribution of something, as distributive egalitarians claim, or about interpersonal relations, as relational egalitarians claim (see, e.g., Anderson 1999, 2010; Scheffler, 2003, 2005). Rather, my claim here is that, even if we decide that something—say, resources or capabilities—should be distributed equally (thereby taking a position with respect to the two debates above), there is still a further question that arises when we consider multidimensional equality: namely, should we equalize the bundles, the extent to which each person's bundle fits her preferences, each person's standing in the community, or something else altogether?

Although I cannot argue for it here, my suspicion is that the three debates are, in fact, intertwined. To illustrate the rough idea

question Question advances alone. Rather, the q

REFERENCES

Anderson, E. (1999). What is the point of equality? *Ethics*, 109, 287-337.

<https://doi.org/10.4324/9781315236322-9>

Anderson, E. (2010). The fundamental disagreement between luck egalitarians and relational egalitarians. *Canadian Journal of Philosophy* 36, 1-23.

<https://doi.org/10.1080/00455091.2010.10717652>

Arneson, R. (1990a). Liberalism, distributive s

treatment. *Theory and Decision* 11, 267-288. <https://doi.org/10.1007/bf00126381>

Scheffler, S. (2003). What is egalitarianism? *Philosophy & Public Affairs* 31, 5-39.
<https://doi.org/10.1111/j.1088-4963.2003.00005.x>

Scheffler, S. (2005). Choice, circumstance, and the value of equality. *Politics, Philosophy, and*