A Philosopher's Guide to Multidimensional Equality

Kristi A. Olson

Forthcoming in *Philosophy Compass*

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, undominated diversity, solidarity test, John Rawls

A Philosopher's Guide to Multidimensional Equality

Take a distribution involving two or more dimensions—say, apples and oranges. If Smith fares worse on one dimension (fewer apples), while Jones fares worse on another (fewer oranges), who, if either, is worse off overall? Philosophers and policymakers sometimes need an answer.

To illustrate how the problem arises for philosophers, consider two examples. First, according to John Rawls's (1971) theory of justice, we should design institutions to improve the lot of the worst-off group, where the worst-off group is identified by its share of things like income, wealth, the prerogatives of authority, and, more controversially, leisure time (see, e.g., Rawls, 1982, 1988). Since multiple dimensions are at stake, Rawlsians face the *indexing problem:* if one group fares worst in one dimension, while other groups fare worst in other dimensions, which group is worst off overall? Of course, if one group fares worst across all dimensions—as some Rawlsians assume—the worst-off group will be easy to identify (see Arneson, 1990b). Yet, even then, Rawlsians still face the challenge of comparing the worst-off group under one institutional design to the worst-off group under alternative institutional designs (Gibbard, 1979; Arneson, 1990b). Rawls, however, never quite explained how this was to be done. Fifty years after the publication of *A Theory of Justice,* the indexing problem remains unsolved (cf. Hockett and Risse, 2006).

Consider second the capabilities approach, pioneered by Amartya Sen (1979) and Martha Nussbaum (1988). According to the capabilities approach, what matters is not each person's share of resources (things like income and leisure time), but rather what she is able to be or do with those resources—her capabilities. For example, is the person able to love, own property, and live a normal length life? Is she able to play, enjoy bodily health, and participate in political life? Although Nussbaum (1997) includes each of the above on her list of central capabilities, there is no general consensus about which capabilities to measure. Rather, the answer depends on the context and, in actual applications, on the data available. Yet, regardless of which capabilities are measured, so long

as the list includes more than one, capability egalitarians face the *aggregation problem*: 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. What, after all, is the public supposed to be doing when it deliberates? We need a principle that guides deliberation.

At the most general level, there are only two options: Either we use subjective assessments or we do not use subjective assessments. Yet, this way of framing the problem conceals the rich variety of ways in which subjective assessments can be used. In this essay, I examine five approaches that use subjective assessments and consider their strengths and weaknesses. I do not here attempt to solve the problem. Instead, I aim to show that the problem is a philosophical one: the solution depends on how we answer certain normative questions. By identifying what I take to be the central question, I hope to take us one step closer to a solution.

I. UTILITY

The most straightforward approach to comparing multidimensional bundles is to convert multiple dimensions into one dimension, such as utility. To illustrate, suppose each person reports her subjective utility (roughly, her happiness) with her bundle of resources or capabilities. We then compare her self-reported utility to those of everyone else.

Neither Rawlsians nor capability egalitarians would endorse this approach; utility is the wrong thing to equalize, they would say. Yet, since I am here just attempting to delineate the options, I will set that objection aside. There is, however, a different problem that affects the viability of the approach: namely, the need to compare utility across persons (see, e.g., Dworkin, 1981a; Fleurbaey, 2012). If we simply ask each person to rate, on a scale of 1 to 10, how happy she

would be with a particular bundle of resources or capabilities, then we need to know whether one person's rating of 5 really means the same degree of utility as someone else's rating of 5. If what Smith subjectively reports as 5 is what Jones would report as 7, any comparison of Smith and Jones will be distorted.

To be clear, the claim here is not that interpersonal comparisons of utility are impossible. As Sen (1970) and others have pointed out, we can indeed compare individuals' utility, even if only roughly. And in some circumstances the simplifying assumptions necessary to make such comparisons might be warranted. Nonetheless, in the present context, any such simplifying assumption would be directly question-begging. We are, after all, attempting to measure equality. The need to calibrate utility across persons is thus at the heart of the problem. As a result, unless the calibration problem can be solved, utility comparisons are not a viable solution.

II. THE ENVY TEST

The second approach is the envy test (Tinbergen, 1930, 1946; Foley, 1967; Kolm, 1971; Varian, 1975; Dworkin, 1981b; Parr, 2018; Heilmann and Wintein, 2021). Suppose Smith has {one apple, two oranges} while Jones has {two apples, one orange}. If neither prefers the other person's bundle, the envy test is satisfied and there is no inequality. If, on the other hand, Jones prefers Smith's bundle, the envy test is not satisfied and Jones is deemed worse off. More formally, under the envy test, each person compares her well-being with her current bundle to her expected well-being with every other person's bundle. These intrapersonal assessments are then converted into interpersonal assessments: If Jones judges that she would be better off if she had Smith's bundle (an intrapersonal assessment), Jones is deemed worse off than Smith (an interpersonal assessment).

One advantage of the envy test is that it allows comparisons of multidimensional bundles without requiring interpersonal comparisons of utility. It thus avoids the calibration problem discussed above. Yet, the envy test faces its own objections. Here I consider two. The first objection

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 not better all-things-considered than A. Yet, since envy can be

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 *and* 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 is better than the reference bundle, while Jones's bundle is worse than the reference bundle. And finally, if both prefer the reference bundle (or if both prefer their own bundle to the reference bundle) the egalitarian equivalent approach asks how much compensation added to (subtracted from) their bundle would make them indifferent between their bundle and the reference bundle. Whoever needs more added (less subtracted) is deemed worse off. More formally, each person compares her well-being with her bundle to her expected well-being with the reference bundle (an intrapersonal

assessment); the person who judges her bundle worst relative to the reference bundle is deemed worst off (an interpersonal assessment).

The egalitarian equivalent approach has an advantage over the envy test: namely, unlike the envy test, the egalitarian equivalent approach can identify the worst-off individual and rank scenarios according to the extent of inequality. And moreover, the egalitarian equivalent approach does so without requiring interpersonal comparisons of utility. Yet, the egalitarian equivalent approach achieves this apparent advantage by shifting the controversy to (i) the choice of reference bundle; and (ii) the choice of metric (see Olson, 2020a).

To illustrate the first controversy—the choice of reference bundle—suppose Smith and Jones both have {two apples, one orange}. Yet, whereas Smith is indifferent between apples and oranges, Jones prefers oranges. Specifically, suppose she is indifferent between one orange and two apples. As a result, if the reference bundle is {three oranges}, Smith will be indifferent between his bundle and the reference bundle, whereas Jones will prefer the reference bundle to her bundle. In contrast, if the reference bundle is {three apples}, Smith will again be indifferent between his bundle and the reference bundle, but Jones will prefer her bundle to the reference bundle. Under either reference bundle, Smith and Jones are deemed unequal despite having identical bundles; I say more on that below. Even granting that oddity, what the example just given shows is that *which* of them fares worse depends on which reference bundle is chosen: if the reference bundle is {three oranges}, Jones is worse off; if the reference bundle is {three apples}, Smith is worse off. The upshot is that the egalitarian equivalent approach is viable only if the choice of reference bundle can be justified (see Fleurbaey, 2008).

To illustrate the second controversy—the choice of metric—suppose Smith and Jones each have {two apples, one orange} and, as before, Jones deems one orange equivalent to two apples, while Smith is indifferent between apples and oranges. If the reference bundle is {two apples, two oranges}, both will prefer the reference bundle. Nonetheless, the extent to which each prefers the

reference bundle will depend on the metric used to measure their indifference point. Suppose we ask each: How many *oranges* 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 *apples* 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, two oranges} and the reference bundle while Jones (who values one orange as equivalent to two apples) will prefer her bundle of {one apple, two oranges} to the reference bundle. As a result, under the egalitarian equivalent approach, Jones could be deemed better off, despite having an inferior bundle.

Proponents of the utility approach face a similar objection: Jones's utility with {one apple, two oranges} might be greater than Smith's utility with {two apples, two oranges}. Yet, this fact will not embarrass proponents of the utility approach. After all, the point of their approach is to compare utilities and not resources or capabilities. Proponents of the egalitarian equivalent approach, then, might attempt a similar response. They could, for example, point out that our ordinary judgments fail to consider how well the bundle fits the person's conception of the good (Fleurbaey, 2012a; Fleurbaey and Blanchet, 2013; see also Brun and Tungodden, 2004). Fleurbaey (2012a) illustrates as follows: imagine a banker who considers his job a miserable fit; he would prefer to be an artist. Compare him to a less well-paid banker who considers banking her dream

job; no other profession would suit her better. Once the difference in fit is taken into account, the less well-paid banker might indeed be better off than the better-paid banker.

Yet, it is hard to see how this explanation could salvage the egalitarian equivalent approach. To the contrary, if we reject our ordinary judgments (e.g., that the person with the inferior bundle is worse off) because these judgments fail to consider fit, it seems as though we must reject the egalitarian equivalent approach for the same reason. After all, if the reference bundle is a better fit for Smith than for Jones, Jones might be worse off than Smith even if they both receive the reference bundle (or a bundle they deem equivalent to the reference bundle). Proponents of the egalitarian equivalent approach, then, must either provide a different reason for rejecting the judgment that a person with an inferior bundle is worse off or else explain why the egalitarian equivalent approach is not vulnerable to a similar objection. Much, then, hinges on how they defend their chosen reference bundle and chosen metric.

IV. UNDOMINATED DIVERSITY

The fourth approach—undominated diversity—appeals to unanimous preferences (Van Parijs, 1995; cf Sen, 1992). If *everyone* in the distributive scheme prefers Smith's bundle to Jones's bundle, Jones is deemed worse off than Smith. On the other hand, if at least one person prefers Jones's bundle or is indifferent between the two bundles, then there is no inequality (or, more modestly, the bundles are incommensurable).

In one respect, undominated diversity bears a close relationship to the envy test: any distribution that satisfies the envy test satisfies undominated diversity. Nonetheless, undominated diversity diverges in two important ways. First, some distributions satisfy undominated diversity but not the envy test; thus, undominated diversity is a weaker requirement. To illustrate, consider a three-person distributive scenario in which Smith and Jones both prefer Smith's bundle but Patel prefers Jones's bundle. Although Smith's bundle is not unanimously preferred (satisfying

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 worst off overall. Similarly, undominated diversity does not allow us to rank scenarios according to the extent of inequality. We cannot, for example, simply conclude that the more instances of

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 *realized.* 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 vulnerable to the problems discussed in Part III. My tentative conclusion in this essay is that egalitarians who take a subjective assessment approach to multidimensional comparisons might be out of luck: absent a solution to the calibration problem, subjective assessment approaches cannot identify the worst-off person or rank distributions according to the extent of inequality.

Even so, philosophers and policymakers presumably still want a test that tells us when (and why) a distribution is equal. That, however, takes us to the second objection. According to the second objection, the solidarity test does not in fact deliver any recognizable equality. To illustrate the problem, consider a two-person distribution in which Jones prefers Smith's bundle. According to the solidarity test, if Smith prefers Jones's bundle, there is no inequality. But—and here is the objection—why does it matter if Smith prefers Jones's bundle if Jones does not? (cf Segall, 2021). After all, what matters to Jones, presumably, is whether she receives a bundle she prefers. If she cannot trade bundles, it does her no good to receive a bundle Smith prefers. Why, then, does Smith's preference for Jones's bundle make an otherwise unequal distribution equal as the solidarity test insists? The viability of the solidarity test depends on an answer.

In response to this objection, a proponent of the solidarity test might explain that the various approaches track different equalities. To illustrate the equality that the solidarity test

[1] The solidarity test is also more demanding than the requirement that every bundle is such that someone ranks it first (see Van Parijs, 1995: 87). That requirement—insofar as it permits ties—would be satisfied by a scenario in which one person is indifferent among all bundles, but everyone else strictly prefers the same bundle. These bundles are incompatible with an envy-free distribution and thus would not satisfy the solidarity test.

tracks, consider two scenarios. In the first scenario, Jones receives a bundle *she* would not choose but someone else in the distributive community would choose. In the second scenario, Jones receives a bundle *no one* (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?—the answer is that, in virtue of Smith's preference, Smith and Jones stand as equals, and that is the equality the solidarity test tracks.[2]

We can now see why the solidarity test, unlike the envy test, does not deliver judgments at odds with ordinary properties of relations. The reason is straightforward: the solidarity test does tell us who, if either, is better off. The solidarity test instead tells us whether Smith and Jones stand as equals. And thus, the solidarity test will never (as the envy test sometimes might) say that Smith is better off than Jones, *and* Jones is better off than Smith.

The five approaches discussed above do not exhaust the realm of possibilities and, for each of the approaches, more would need to be said to make the approach plausible. Nonetheless, we can now identify the crucial normative question behind any attempt to compare multidimensional bundles: namely, what exactly should be equalized? To be clear, I am not here referring to the equality of what debate, which asks whether we should equalize resources, capabilities, welfare, or

[2] Of course, there might be other reasons why Smith and Jones do not stand as equals. If, for example, women are given their dispreferred bundle while men are given their preferred bundle, women and men do not stand as equals even if their bundles are compatible with an envy-free distribution. The solidarity test only addresses one of many considerations relevant to equal standing. Nonetheless, the solidarity test tracks one important consideration at issue in the comparison of multidimensional bundles.

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: if a distributive egalitarian faced with the question of how to compare multidimensional bundles decides to use the solidarity test, the gap between relational and distributive egalitarians might be rather narrow. Indeed, when faced with the task of comparing multidimensional bundles, an erstwhile distributive egalitarian might decide that she is ultimately concerned with equal standing, and thus with interpersonal relations. That is, when the distributive egalitarian considers multidimensional equality, she might discover that she is in fact a relational egalitarian. If this is right, then a solution to the problem of multidimensional equality might shed light on the point of equality debate. Similarly, as Arneson (1990a, 1990b) and others have pointed out, certain answers to the what is equality debate—e.g., resource egalitarianism—become less plausible when multidimensional equality is at stake.

CONCLUSION

Philosophers and policymakers sometimes need to compare multidimensional bundles. This essay does not solve that problem. Nonetheless, it does—I hope—take us one step closer to a solution. My main claim here is that the multidimensional problem cannot be solved by technical

quesionQuestion advances alone. Rather, the question that ultimately must be answered is normative: what exactly should we equalize? Philosophers, then, have much to contribute to the debate, and the payoff could be considerable: policy decisions and other egalitarian debates might depend on the answer.

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 subjectivism, and equal opportunity for welfare. Philosophy and Public Affairs, 19, 158-194.

Arneson, R. (1990b). Primary goods reconsidered. Noûs, 24, 429-254. https://doi.org/10.2307/2215774

Brun, B.C. and Tungodden, B. (2004). Non-welfaristic theories of justice: is "the intersection approach" a solution to the indexing impasse? Social Choice and Welfare, 22, 49-60. https://doi.org/10.1007/s00355-003-0276-0

Chaudhuri, A. (1986). Some implications of an intensity measure of envy. Social Choice and Welfare3, 255-270. https://doi.org/10.1007/bf00292731

Cohen, G.A. (1989). On the currency of egalitarian justice. Ethics, 99, 906-944.

Dworkin R. (1981a). What is equality? part 1. Equality of welfare. Philosophy & Public Affairs, 10, 185-246.

Dworkin R. (1981b). What is equality? part 2. Equality of resources. Philosophy & Public Affairs, 10, 283-345. https://doi.org/10.4324/9781315199795-7

Fleurbaey, M. (2008). Fairness, responsibility, and welfare. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199215911.001.0001

Fleurbaey, M. (2012a). The importance of what people care about. Politics, Philosophy, and Economics, 11, 415-447. https://doi.org/10.1177/1470594x12447775

Fleurbaey, M. (2012b). Economics is not what you think: a defense of the economic approach to taxation. Working Paper Series *Fondation Maison des sciences de l'homme*.

Fleurbaey, M. and Blanchet, D. (2013). Beyond GDP: measuring welfare and assessing sustainability. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199767199.001.0001

Foley, D. K. (1967). Resource allocation and the public sector. Yale Economic Essays 7, 45-98.

Gibbard, A. (1979). Disparate goods and Rawls' difference principle: a social choice theoretic

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

Heilmann, C. and Wintein, S. (2021). No envy: Jan Tinbergen on fairness. Erasmus journal for philosophy and economics, 14, 222-245.

Hockett, R.C. and Risse, M. (2006). Primary goods revisited: the "political problem" and its Rawlsian solution. Cornell Law Faculty Publications. Paper 55.
http://scholarship.law.cornell.edu/lsrp_papers/55

Iturbe-Ormaetxe I. and Nieto, J. (1996). On fair allocations and monetary compensations. Economic Theory, 7, 125-38. https://doi.org/10.1007/bf01212186

Kolm, S.C. (1971/1997). Justice and Equity. MIT Press.

Nussbaum, M. (1988). Nature, function, and capability: Aristotle on political distribution. Oxford Studies in Ancient Philosophy, 145-184.

Nussbaum, M. (1997). Capabilities and human rights. Fordham Law Review, 66, 273-300.

Olson, K.A. (2018). Impersonal envy and the fair division of resources. Philosophy and Public Affairs, 46, 269-292. https://doi.org/10.1111/papa.12122

Olson, K.A. (2020a). The solidarity solution: principles for a fair income distribution. New York: Oxford University Press. https://doi.org/10.1093/oso/9780190907457.001.0001

Olson, K.A. (2020b). The importance of what we care about: a solidarity approach to resource allocation. Political Studies. https://doi.org/10.1177/0032321720972872

Parr, T. (2018). How to identify disadvantage: taking the envy test seriously. Political Studies 66, 306-322. https://doi.org/10.1177/0032321717720377

Pazner, E.A. and Schmeidler, D. (1978). Egalitarian equivalent allocations: a new concept of economic equity. Quarterly Journal of Economics 92, 671-686. https://doi.org/10.2307/1883182

Rawls, J. (1971). A theory of justice. Cambridge, MA: Harvard University Press.

Rawls, J. (1982). Social unity and primary goods. In A. Sen and B. Williams (Eds.) Utilitarianism and beyond (pp. 159-185). Cambridge: Cambridge University Press.
https://doi.org/10.1017/cbo9780511611964.010

Rawls, J. (1988). The priority of right and ideas of the good. Philosophy and Public Affairs, 17, 251-276.

Robeyns, I. (2017). Wellbeing, freedom, and social justice: the capabilities approach re-examined. Cambridge: Open Book Publishers. https://doi.org/10.11647/obp.0130

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 Economics 4, 5-28. https://doi.org/10.1177/1470594x05049434

Segall, Shlomi (2021). The old man's bundle, still: Kristi Olson revisits the envy test. Analysis 81, 378-385.

Sen, A. (1970). Interpersonal aggregation and partial comparability. Econometrica, 38, 393-409. https://doi.org/10.2307/1909546

Sen, A. (1979). Equality of what? The Tanner Lectures on Human Values, Stanford University, Palo Alto, CA.

Sen, A (1985). Commodities and capabilities. Amsterdam: North Holland.

Sen, A. (1992). Inequality reexamined. Cambridge, Mass.: Harvard University Press. https://doi.org/10.1093/0198289286.001.0001

Tinbergen, J. (1930/2021). Mathematical psychology. Translated by Erwin Dekker, Conrad Heilmann, Ruth Hinz, and Stefan Wintein. Erasmus Journal for Philosophy and Economics.

Tinbergen, J. (1946). Redelijke Inkomensverdeling. Haarlem, Netherlands: De Gulden Pers.

Van Parijs, P. (1990). Equal endowments as undominated diversity. Louvain Economic Review 56, 327-355.

Van Parijs, P. (1993). Rawlsians, Christians and patriots: maximin justice and individual Ethics. European Journal of Philosophy 1, 309-342. https://doi.org/10.1111/j.1468-0378.1993.tb00088.x

Van Parijs P (1995). Real freedom for all: what, if anything, can justify capitalism? Oxford: Oxford University Press. https://doi.org/10.1093/0198293577.001.0001

Varian, HR (1975). Distributive justice, welfare economics, and the theory of fairness. Philosophy and Public Affairs, 4, 223-247.

Wolff J and De-Shalit A (2007). Disadvantage. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199278268.001.0001