

Utils and Shmutils^{*}

Jacob M. Nebel

University of Southern California

1 Introduction

How should governments decide between alternative taxation schemes, environmental protection regulations, infrastructure plans, climate change policies, healthcare systems, and other policies? One kind of consideration that should bear on such decisions is their effects on people’s well-being. The most rigorous methodology for evaluating such effects is the “social welfare function” (SWF) approach originating in the work of Abram Bergson and Paul Samuelson and further developed by Kenneth Arrow, Amartya Sen, and other economists.¹ Matthew Adler’s *Measuring Social Welfare* is an introduction to this methodology.

This essay focuses on what Adler identifies as two key tools of the SWF methodology. The first is a measure of well-being. This measure is a function that takes an outcome and returns a list of numbers—a *well-being vector*, as Adler calls it—that represents the distribution of well-being in that outcome. Each component in the list—each *well-being number*, as

^{*}A review essay on Matthew D. Adler, *Measuring Social Welfare: An Introduction* (Oxford University Press, 2019). Page references are to this book unless otherwise noted. Thanks to Matthew Adler, Kara Dreher, Orri Stefánsson, and Trevor Teitel for helpful comments.

¹Abram Burk (Bergson), “A Reformulation of Certain Aspects of Welfare Economics,” *The Quarterly Journal of Economics* 52 (1938): pp. 310–34; Paul A. Samuelson, *Foundations of economic analysis* (Harvard University Press, 1947); Kenneth Joseph Arrow, *Social choice and individual values* (New York: Wiley, 1951); Amartya Sen, *Collective choice and social welfare* (San Francisco: Holden-Day San Francisco., 1970).

Adler calls it—represents a person's lifetime well-being in that outcome. Outcome x is assigned well-being vector $w(x) = (w_1(x), w_2(x), \dots, w_n(x))$, where $w_i(x)$ is the well-being number that represents person i 's well-being in outcome x . The second tool is a rule for ranking these well-being vectors. This rule—which we can think of as the SWF itself—will tell us how to compare any two well-being vectors. Put together, these tools let us compare outcomes by comparing their well-being vectors. Given any two outcomes, we apply the well-being measure to obtain their well-being vectors. Then we use the rule for ranking those vectors to determine which is better.

So described, the SWF methodology may appear to be of interest only to welfarist consequentialists, who believe that the rightness of acts depends only on the goodness of their outcomes, and that the goodness of outcomes depends only on their distributions of welfare. Indeed, Adler explicitly characterizes the SWF approach as both welfarist and consequentialist (p. 7). But, I believe, the SWF methodology should be of much broader interest. Even if welfare is not all that matters, it does matter. We can think, without being welfarists, that if one outcome's distribution of well-being is better than another's, then the one outcome is better than the other in one important respect—namely, with respect to the value of well-being. And we can think, without being consequentialists, that the comparison of outcomes with respect to this value is sometimes relevant to what we ought to do. (In this spirit, when I say that one outcome is better than another, this should be understood to mean only that it is better with respect to the value of well-being.)

Adler's central aim is to explain the power, flexibility, and normative foundations of the SWF methodology in an accessible way. A secondary aim is to argue that the SWF methodology is superior to the method of cost-benefit analysis more widely used by policymakers. The book handily achieves these aims. Adler offers a helpful overview of the SWF framework, discusses its ethical presuppositions, highlights axiomatic choice points for deciding between different types of SWFs, applies the SWF framework to a practical case study,

proposes an institutional role for the SWF framework, and suggests extensions of the SWF framework to consider variable-population choices and individual responsibility. It will be an invaluable resource for advanced undergraduates, graduate students, and faculty working in philosophy, economics, law, political science, and public policy.

In this essay, I question some ideas at the core of the SWF framework that, if true, would have profound implications for moral and political philosophy, but which unfortunately have received little attention from philosophers. These ideas have to do with the relation between the two tools mentioned above: the well-being measure and the rule for ranking well-being vectors.

The facts about individual well-being do not single out a particular scale on which well-being must be measured. As with physical quantities, there are multiple scales that can be used to represent the same information about well-being; no one scale is special. Like physical laws, the SWF and its ranking of distributions cannot depend on exactly which of these scales we use. Adler and other theorists in the SWF tradition have used this idea to derive highly restrictive constraints on the shape of the SWF. These constraints rule out seemingly plausible views about distributive justice and population ethics.

I argue, however, that these constraints stem from a simple but instructive mistake. The SWF should not be applied to vectors of *numbers* such as 1 and 2, but rather to vectors of *dimensioned quantities* such as 1 util and 2 utils. This seemingly pedantic suggestion turns out to have far-reaching consequences. Unlike the orthodox SWF approach, treating welfare levels as dimensioned quantities lets us distinguish between real changes in well-being and mere changes in the unit of measurement. It does this without making the SWF depend on the scale on which welfare is measured, and in a way that avoids the restrictive constraints on the shape of the SWF. We'll see how it does this, and why that is important, in section 5. Until then, we'll build up to my proposal by raising a problem for Adler's own preferred SWF (section 2), tracing the problem to these issues about measurement (section 3), and drawing

out those issues' implications for other views about distributive justice and population ethics (section 4).

2 The Atkinson SWF

The aim of *Measuring Social Welfare* is not to defend any particular SWF. Adler does, however, have a favorite kind of SWF, defended at length in his previous book *Well-Being and Fair Distribution*,² and it receives special attention in *Measuring Social Welfare*.

Adler favors a kind of what he calls *continuous-prioritarian* SWF. (I'll just say "prioritarian," taking the "continuous" part for granted.) This kind of SWF evaluates distributions by first applying a transformation to each person's well-being, which I will call the *priority weighting* function. A person's priority-weighted well-being is a strictly increasing and strictly concave function of her well-being. This means that a person's well-being has positive but diminishing marginal priority-weighted value: an increment of well-being always increases a person's priority-weighted well-being, but by less the better off she is. A prioritarian SWF adds up everyone's priority-weighted well-being and judges one distribution to be better than another just in case the one has a greater sum of priority-weighted well-being.

Adler's case for a prioritarian SWF is axiomatic. First, prioritarian SWFs satisfy the *strong Pareto* principle: if each person is at least as well off in one distribution as they are in another, and at least one person is better off, then the one distribution must be better than the other (p. 97). Second, prioritarian SWFs satisfy the principle of *anonymity*, which requires any two distributions that are permutations of each other to be equally good (p. 97). Two distributions are permutations of each other if they contain the same number of lives at each welfare level, with those levels (possibly) rearranged among people. These two axioms—

²Matthew Adler, *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* (Oxford University Press, 2011).

strong Pareto and anonymity—are satisfied by all of the SWFs Adler considers.

Third, prioritarian SWFs satisfy the *Pigou-Dalton* principle, which recommends any *pure, gap-diminishing* transfer of well-being from a better-off to a worse-off person in which everyone else remains unaffected (p. 89). A pure, gap diminishing transfer is one in which one person's loss equals the other's gain and the difference between their welfare levels is reduced. The Pigou-Dalton principle rules out the utilitarian SWF, which compares distributions by their sums of well-being. It also rules out what Adler calls *sufficientist* (often called "sufficientarian") SWFs, which give priority to the worse off only when they are below some sufficient threshold of well-being.³ Unlike the utilitarian SWF, sufficientist SWFs satisfy what Adler calls the *minimal* Pigou-Dalton principle: a pure, gap-diminishing transfer from the better-off to the worse-off sometimes makes a distribution better, and a pure transfer from the worse-off to the better-off never makes a distribution better (p. 95).

Fourth, prioritarian SWFs satisfy the axiom of *separability*. Separability says that the ranking of distributions cannot depend on the welfare of unaffected individuals—i.e., people whose welfare remains the same in the distributions compared (p. 89). Separability distinguishes prioritarian SWFs from paradigmatically egalitarian SWFs, on which the importance of benefiting a person depends not only on her own well-being and the size of the benefit but also on how she fares relative to others. For example, *rank-weighted* SWFs, such as the Gini SWF, evaluate distributions by multiplying each person's welfare by a weight that depends on her position in the distribution's rank-ordering of welfare levels, giving greater weight to the worse off, and summing these weighted welfare levels.⁴ Rank-weighted SWFs violate separability because the weight assigned to a person's welfare depends on the welfare of others, including unaffected individuals (p. 91).

³See Roger Crisp, "Equality, Priority, and Compassion," *Ethics* 113 (2003): pp. 745–63.

⁴For a recent defense of rank-weighted SWFs, see Lara Buchak, "Taking Risks behind the Veil of Ignorance," *Ethics* 127 (2017): pp. 610–44. I offer a critique of Buchak's argument in Jacob M. Nebel, "Rank-Weighted Utilitarianism and the Veil of Ignorance," *Ethics* 131 (2020).

Fifth, prioritarian SWFs satisfy the axiom of *continuity*. Continuity says that if one distribution is better than another, then a distribution resulting from a sufficiently small change to the one must also be better than the other (p. 102). For example, if distribution (1, 3) is better than (1, 2), then for some sufficiently small ε , $(1 - \varepsilon, 3)$ must also be better than (1, 2). This implication is violated by the *leximin* SWF, which gives absolute priority to the worst-off affected individual. Leximin judges (1, 3) to be better than (1, 2), because the former is better for the only affected individual. But it judges $(1 - \varepsilon, 3)$ to be worse than (1, 2), because the latter is better for the worst-off affected individual. Continuity also rules out the sufficientist SWFs that Adler considers, because they give absolute priority to people below the threshold when their interests conflict with those above it.

Table 1 shows which axioms are satisfied by each of the SWFs Adler considers. The SWFs are listed in rows, axioms in columns. A bullet in a cell indicates that the corresponding SWF satisfies the corresponding axiom. Prioritarian SWFs are the only SWFs that satisfy all five of Adler’s axioms: strong Pareto, anonymity, Pigou-Dalton, separability, and continuity (p. 105).


Table 1: Properties of Adler’s SWFs

	Strong Pareto	Anonymity	Pigou-Dalton	Minimal Pigou-Dalton	Separability	Continuity
Prioritarian SWF	•	•	•	•	•	•
Utilitarian SWF	•	•			•	•
Rank-Weighted SWF	•	•	•	•		•
Leximin SWF	•	•	•	•	•	
Sufficientist SWF	•	•		•	•	

There are infinitely many prioritarian SWFs, one for each possible priority weighting function. For reasons we’ll see in section 3, Adler favors a particular subfamily of prioritarian SWFs known as *Atkinson* SWFs (pp. 154–55).⁵ Atkinson SWFs use a priority weighting

⁵Named after Anthony B Atkinson, “On the measurement of inequality,” *Journal of Economic Theory* 2 (1970): pp. 244–63.

Figure 1: Atkinson priority weighting functions



function $g(\cdot)$ with a particular shape, determined by a positive *priority parameter* γ . When $\gamma = 1$, the priority weighting function is the natural logarithm: $g(w) = \log w$. Otherwise, it has the form:

$$g(w) = \frac{w^{1-\gamma}}{1-\gamma}$$

For example, when $\gamma = 1/2$, $g(w) = 2\sqrt{w}$: priority-weighted well-being is twice the square root of well-being. (Though there are infinitely many Atkinson SWFs, I sometimes use “the Atkinson SWF” to refer to this entire family.)

Figure 1 shows some examples of Atkinson priority weighting functions. The horizontal axis represents well-being, the vertical axis priority-weighted well-being. The blue curve plots $g(\cdot)$ when $\gamma = 1/2$. The red curve plots $g(\cdot)$ when $\gamma = 1$. The green curve plots $g(\cdot)$ when $\gamma = 2$. All three curves have a positive but decreasing slope: increments of well-being do less to increase priority-weighted well-being from higher levels.

Figure 1 shows Atkinson weighting functions only over nonnegative welfare levels. This is because, as Adler mentions in a footnote,

Atkinson SWFs do have a significant downside: they require well-being to be

non-negative or, in the case of $\gamma \geq 1$, strictly positive. The Atkinson $g(\cdot)$ function is either undefined or, if defined, not both strictly increasing and strictly concave with non-negative well-being numbers as inputs (or with 0 as an input for $\gamma \geq 1$).
(p. 155, fn. 33)

This is a serious problem for the Atkinson family, though the exact shape of the problem depends on what it means for welfare to be negative.

Adler does not commit to a particular view about the zero level in this book (p. 156). But he mentions a standard approach in population ethics, which is to define the zero level as the value of a *neutral life*: a life that is neither worth living nor worth *not* living.⁶ It seems possible for a life to be sufficiently bad—for example, because it is filled only with unmitigated suffering—that it is worth not living. Such a life would, on this account of the zero level, have negative well-being.

Suppose that some people actually have lives that are worth not living. Then, if the priority-weighted value of negative welfare is undefined, the Atkinson SWF cannot be used to evaluate any choice whatsoever, because the sum of priority-weighted well-being will remain undefined no matter what. And, even if everyone actually has a life worth living, Atkinson weighting functions that are undefined over negative welfare would lead the SWF to imply that it wouldn't be worse to reduce people's welfare to the point where their lives would be worth not living. (Remember that the prioritarian SWF is biconditional: one distribution is better than another if *and only if* the one contains a greater sum of priority-weighted well-being.) If we instead use an Atkinson weighting function that is defined over negative welfare, and if some people have lives worth not living, then the SWF either will recommend that we make these people even *worse* off (because the priority weighting function is strictly decreasing over negative welfare—e.g., when $\gamma = 3$), or will recommend any pure transfer from the worse off to the *better off* among these people (because the priority

⁶Adler takes this approach in *Well-Being and Fair Distribution*, 219.

weighting function is strictly convex over negative welfare—e.g., when $\gamma = 2/3$). All of these results are inconsistent with prioritarian convictions.

Adler might respond that Atkinson SWFs should just be restricted to populations in which all lives are worth living and to choices in which all lives would remain worth living no matter what.⁷ But that restriction would not be enough. Suppose it *might* be true, for all we know, that some lives are worth not living or would be if some policy were chosen. Adler's preferred method of applying a SWF under uncertainty is to maximize expected social welfare. On this method, Atkinson SWFs will have absurd implications similar to those above. So the Atkinson family would need to be restricted to choices in which we are certain that all members of society have lives worth living and would have such lives no matter what. But there are no such choices: we cannot reasonably be certain that every member of any contemporary society has and would invariably have a life worth living no matter what we do. And, even if there were such choices, a proponent of the SWF methodology would need some other SWF to evaluate our other choices. That SWF would have to violate one of the axioms that characterize the Atkinson family (one of the five prioritarian axioms, or a sixth we'll see in section 3). We might then wonder why we shouldn't just use that SWF to evaluate all choices, and why we should accept Adler's axioms when and only when all welfare levels are certainly positive.

These problems could be avoided by reinterpreting the zero level in such a way that it is impossible for a person's well-being to be negative. For example, suppose there is a worst possible life. Define the zero level to be the value of such a life. Then it will be impossible for people to have negative well-being.

But, first, there does not appear to be any lower bound on the value of lives. For any life, no matter how bad, it seems possible for there to be a life that is much worse—for example, by containing much more suffering. Second, Atkinson SWFs would have absurd implications

⁷Adler seems to suggest this in *Well-Being and Fair Distribution*, p. 392.

regarding variable-population choices if all possible lives were assigned positive welfare, no matter how bad. The most natural extension of prioritarian SWFs to variable-population choices is to compare distributions by their sums of priority-weighted well-being.⁸ If every life were assigned positive welfare, then Atkinson SWFs (so extended) would imply a particularly implausible version of Derek Parfit’s *repugnant conclusion*.⁹ When $\gamma < 1$, the Atkinson SWF would imply that, for any number of excellent lives, there must be some number of lives whose existence would be better, even though all of them would be just barely better than the worst life people could possibly live. This is because, when $\gamma < 1$ and all lives have positive welfare, all lives have positive priority-weighted value, so enough lives of arbitrarily low well-being could contain a greater sum of priority-weighted well-being. When $\gamma > 1$, it would imply that, for any number of horrible lives, there must be some number of lives whose existence would be worse, even though all of them would be wonderful. This is because, when $\gamma > 1$ and all lives have positive welfare, all lives have *negative* priority-weighted value, so enough lives of arbitrarily high well-being could contain a lower—more negative—sum of priority-weighted well-being. These conclusions are unacceptable. (The $\gamma = 1$ case avoids both implications. But, in variable-population cases, the logarithm function does not satisfy the key axiom that distinguishes the Atkinson family from other prioritarian SWFs. See footnote 18.)

Another approach would be to claim that the zero level is arbitrary, and to insist that we can always renormalize our scale so that everyone in a distribution is assigned nonnegative welfare. To do this, the zero level would have to be renormalized, in each choice, to the

⁸By the “sum of priority-weighted well-being” in outcome x , I mean $\sum_{i=1}^n g(w_i(x))$, where n is the number of people in x . This extension of prioritarianism is endorsed by Matthew D. Adler, “Future Generations: A Prioritarian View,” *Geo. Wash. L. Rev.* 77 (2009): p. 1478. Adler now prefers a different view, which adds up each person’s priority-weighted well-being after first subtracting the priority-weighted well-being of a neutral life, which is set greater than the zero level (Matthew D. Adler and Nicolas Treich, “Prioritarianism and Climate Change,” *Environmental and Resource Economics* 62 (2015): pp. 279–308). This avoids the problem I’m about to raise, but at the cost of violating the key axiom that distinguishes Atkinson SWFs in variable-population cases. See footnote 24.

⁹Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984).

worst level of well-being of any person in any available option.¹⁰ This procedure, however, would have bizarre consequences for Atkinson SWFs, which require a meaningful zero level (more in this in section 3).¹¹ For example, let $\gamma = 1/2$. Suppose that Ann and Bob have equally miserable lives. We can either benefit both Ann and Bob by 1 unit or benefit Bob by 3 units. If the zero level is normalized to the worst-off welfare level in any available outcome, the available distributions are represented by the vectors $(1, 1)$ and $(0, 3)$. The Atkinson SWF, with $\gamma = 1/2$, judges $(1, 1)$ to be better, so benefiting both people maximizes social welfare. Now suppose we have a third option. We can make Ann's and Bob's lives much worse, by 100 units each. This third option would result in the worst outcome for both, and their lives would be equally bad, so it gets represented as $(0, 0)$; benefiting both then gets represented as $(101, 101)$; benefiting Bob alone gets represented as $(100, 103)$. Our SWF judges $(100, 103)$ to be best: in the three-option choice, benefiting Bob, rather than both, maximizes social welfare. As this example shows, renormalizing the zero level to make all welfare levels nonnegative would mean the Atkinson SWF's comparison between two options can change with the introduction of a third option—one that is, in this case, Pareto-inferior to the others. Some people think that adding or subtracting an option can sometimes affect the permissibility of choosing other options. But it seems absurd to think that the addition of *this* kind of option—making everyone gratuitously worse off—should have such an effect. If it could, then social decisionmakers could be manipulated to change their decisions, at some cost, by the addition or removal of options that have no attraction whatsoever. That is an undesirable result.¹²

I conclude that Atkinson SWFs must be rejected. This conclusion leaves us with some work to do, because there appear to be good reasons for prioritarrians such as Adler to favor

¹⁰ Adler raises this possibility in Adler and Treich, 'Prioritarianism and Climate Change' [p. 294; p. 305, fn. 305].

¹¹ Consequences along these lines were pointed out to me by Orri Stefánsson.

¹² On this kind of manipulability in decision theory, see John Quiggin, "Regret theory with general choice sets," *Journal of Risk and Uncertainty* 8 (1994): pp. 153–65.

the Atkinson SWF.

3 Invariance

Why, among prioritarian SWFs, does Adler favor the Atkinson family? The answer has to do with the measure of well-being and brings us to the main topic of this essay.

The first key tool of the SWF methodology, recall, is the well-being measure $w(\cdot)$. This is a function that takes an outcome x and returns a well-being vector $w(x) = (w_1(x), w_2(x), \dots, w_n(x))$, where $w_i(x)$ represents person i 's well-being in x . For example, suppose that outcomes x and y contain two people, Ann and Bob. Suppose that $w(x) = (2, 4)$ and $w(y) = (3, 3)$, with Ann's and Bob's welfare levels listed in alphabetical order. How exactly do these numbers *represent* these people's welfare? And what, if anything, is special about *these* numbers in particular?

The SWF framework has a standard answer. The well-being numbers represent people's well-being in the sense that relations between the welfare levels of different people in different outcomes are mirrored by the relations between the numbers assigned to those people in those outcomes. The precise relations that are so represented depend on the measurability and interpersonal comparability of well-being. For example, suppose that well-being is only ordinally and intrapersonally measurable, with no interpersonal comparability. Then the vectors assigned by $w(\cdot)$ to x and y above— $w(x) = (2, 4)$ and $w(y) = (3, 3)$ —represent the facts that y is better for Ann than x , and that x is better for Bob than y . But we could not extract from these vectors that, say, x is better for Bob than it is for Ann, that Ann and Bob are equally well off in y , or that y is better for Ann than x by as much as x is better for Bob than y . The fact that $w_{\text{Ann}}(y) = w_{\text{Bob}}(y)$ is an arbitrary feature of the numbers assigned by $w_{\text{Ann}}(\cdot)$ and $w_{\text{Bob}}(\cdot)$, if there is no interpersonal comparability of welfare levels.

The distinction between relations that are and relations that are not represented by

the well-being measure can be made precise in the following way. There is not just one well-being measure $w(\cdot)$. There are many, and no particular one of them is special. Some of these measures conflict, or disagree, in the sense that they contain different information about people's well-being. For example, consider a well-being measure $w'(\cdot)$ that assigns $w'(x) = (4, 12)$ and $w'(y) = (3, 7)$. Even with a merely ordinal scale of well-being with no interpersonal comparability, $w(\cdot)$ and $w'(\cdot)$ conflict. This is because $w(\cdot)$ assigns greater well-being to Ann in y than to her in x , whereas $w'(\cdot)$ does the opposite. But if we instead assign $w'(x) = (3, 12)$ and $w'(y) = (4, 7)$, then $w(\cdot)$ and $w'(\cdot)$ would contain the same ordinal, intrapersonal information about these people's well-being, at least in these two outcomes.

Suppose two well-being measures, $w(\cdot)$ and $w'(\cdot)$ represent the same information about people's well-being—remaining open, for now, about what information there is to represent. Then it should not matter which measure we use. This idea is captured by what Adler calls “the Fundamental Principle of Invariance,” which says that a SWF should be “invariant to admissible rescalings of the well-being measure” (p. 43). An admissible rescaling of a well-being measure is just another well-being measure that contains the same information as the original. A SWF is invariant to such rescalings just in case, for any outcomes x and y , however the SWF ranks vectors $w(x)$ and $w(y)$, it ranks $w'(x)$ and $w'(y)$ the same way, where $w'(\cdot)$ is an admissible rescaling of $w(\cdot)$. To see why this principle is important, suppose that a SWF is not invariant to admissible rescalings of the well-being measure. Then there must be some outcomes such that the SWF's comparison between those outcomes depends on which well-being measure is used, even among well-being measures that represent the exact same information about people's well-being. This seems indefensible, because the choice between $w(\cdot)$ and $w'(\cdot)$ is arbitrary. It is like being willing to run some distance when measured in miles but not in kilometers, or changing preferences between job offers depending on whether one thinks of their salaries in dollars or in cents.

The Fundamental Principle of Invariance is, as Adler emphasizes, one of the most

important ideas in the SWF literature.¹³ Its precise implications depend on what kinds of information can meaningfully be represented by a scale of well-being. The less information there is to represent, the stricter its requirements. This is because the set of admissible rescalings is larger if there are fewer relations that must be preserved for two measures to agree. And if that set is larger, then it is harder for a SWF to be invariant with respect to all measures in that set.

We can think of the different possible kinds of meaningful information along a spectrum. On one extreme is the ordinal, purely intrapersonal scale mentioned above. This is an impoverished informational framework for purposes of social welfare evaluation. The SWF methodology cannot get very far with a purely ordinal, intrapersonal scale of well-being. Such a scale would make it admissible to rescale each person's well-being measure by increasing the well-being numbers assigned to some people, decreasing the numbers assigned to others, and by shrinking or stretching different people's scales of well-being, so long as each person's ordering of alternatives is preserved. If such individual-specific rescalings are admissible, then the Fundamental Principle of Invariance implies that they cannot affect the SWF's rankings. But the only kind of SWF is invariant to such rescalings while satisfying the strong Pareto principle is a serial dictatorship, in which the comparison between any two distributions is determined by the well-being of some particular person (the dictator), unless she is equally well off in both distributions, in which case it's determined by the well-being of some second person (the deputy dictator)—and so on.¹⁴

¹³For helpful overviews of invariance conditions and their implications for social welfare evaluation, see Claude d'Aspremont and Louis Gevers, "Social welfare functionals and interpersonal comparability," in Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, eds., *Handbook of Social Choice and Welfare*, Handbook of Social Choice and Welfare, 1 (Elsevier, 2002), pp. 459–541; Walter Bossert and John A. Weymark, "Utility in Social Choice," in Salvador Barberà, Peter J. Hammond, and Christian Seidl, eds., *Handbook of Utility Theory: Volume 2 Extensions* (Boston, MA: Springer US, 2004), pp. 1099–177; John Weymark, "Social Welfare Functions," in Matthew D. Adler and Marc Fleurbaey, eds., *The Oxford Handbook of Well-Being and Public Policy* (Oxford University Press, 2016).

¹⁴As Adler points out (p. 45), this result holds even if each person has a cardinal scale of well-being, so that the SWF must only be invariant to individual-specific *cardinal* rescalings (i.e., positive affine transformations, as opposed to any strictly increasing transformation). See R. Duncan Luce and Howard Raiffa, *Games and*

On the other extreme is an *absolute* scale of well-being, on which there is just a single welfare measure that accurately represents the available information about well-being. Measurement theorists sometimes characterize counting and probability as using absolute scales.¹⁵ There is a single number that represents how many pages there are in Adler’s book, or the probability of a fair coin landing heads conditional on its being flipped. No features of these numbers are arbitrary.¹⁶ This maximally rich informational framework would make the Fundamental Principle of Invariance trivial: every SWF would satisfy it. But it seems absurd to suppose that, for each person, there is a single number that represents how well off she is. There is no privileged, natural unit of well-being.

In between ordinal measurability without interpersonal comparability and absolute measurability with full interpersonal comparability are many possibilities. Adler’s preferred informational framework is closer to the absolute end of the spectrum. He takes there to be meaningful *ratios* of different people’s welfare levels. If $w(\cdot)$ is an accurate representation of people’s well-being on a ratio scale, then $w(x) = (1, 4)$ and $w(y) = (2, 2)$ imply that Bob is four times as well off as Ann in outcome x , that y is twice as good as x for Ann, and that x is twice as good as y for Bob. These relations are preserved by any *common ratio rescaling* of $w(\cdot)$: any $w'(\cdot)$ that differs from $w(\cdot)$ by multiplying everyone’s welfare by the same positive number. For example, let $w'(\cdot) = 3 w(\cdot)$. Then $w'(x) = (3, 12)$ and $w'(y) = (6, 6)$, which still imply that Bob is four times as well off as Ann in outcome x , that y is twice as good as x for Ann, and that x is twice as good as y for Bob. A ratio scale also lets us compare ratios of

decisions: Introduction and critical survey (Games and decisions: introduction and critical survey; Oxford, England: Wiley, 1957), pp. xix, 509, at 344; Louis Gevers, “On Interpersonal Comparability and Social Welfare Orderings,” [Wiley, Econometric Society] *Econometrica* 47 (1979): pp. 75–89.

¹⁵On counting, see Fred S Roberts, *Measurement theory with applications to decisionmaking, utility, and the social sciences* (1984), 64. On probability, see R. Duncan Luce and Patrick Suppes, “Representational Measurement Theory,” in Hal Pashler, ed., *Stevens’ Handbook of Experimental Psychology* (Hoboken, NJ, USA: John Wiley & Sons, Inc., 2002), 8. I suspect that many philosophers would disagree with Luce and Suppes and maintain that the $[0, 1]$ scale is an arbitrary but especially convenient normalization.

¹⁶At least, on the standard definition of conditional probability as the ratio between two unconditional probabilities. For critique, see Alan Hájek, “What Conditional Probability Could Not Be,” *Synthese* 137 (2003): pp. 273–323.

differences between welfare levels. $w(\cdot)$ and $w'(\cdot)$'s assignments both imply that Bob's gain from x coming about rather than y would be twice as great as Ann's loss. But, on a ratio scale, absolute differences or levels are not themselves significant. There is nothing special about Bob being 3 units better off than Ann in x , or about them both having welfare level 2. These values are not special in the sense that a welfare measure can, as $w'(\cdot)$ does, vary these numerical levels and differences while preserving all of $w(\cdot)$'s information.

If any common ratio rescaling of our welfare measure is admissible, then the Fundamental Principle of Invariance appears to imply what I will call

Multiplicative Invariance: For any well-being vectors $w(x)$ and $w(y)$, and any $k > 0$: $w(x)$ is at least as good as $w(y)$ if and only if $kw(x)$ is at least as good as $kw(y)$.¹⁷

Multiplicative invariance implies, for example, that if our SWF ranks vector $(2, 2)$ better than $(1, 4)$, then it must also rank $(6, 6)$ better than $(3, 12)$, and more generally $(2k, 2k)$ better than $(1k, 4k)$ for any positive k . The ranking of well-being vectors cannot change if everyone's welfare is multiplied by k .

Multiplicative invariance is widely accepted in the SWF literature. To see its appeal, suppose we have some welfare measure $w(\cdot)$ —call it the *util* measure—and our SWF ranks $w(x) = (w_1(x), \dots, w_n(x))$ better than $w(y) = (w_1(y), \dots, w_n(y))$. A util is just the value of some arbitrary life that is mapped by $w_i(\cdot)$ to level 1, or the difference in value between lives that are mapped by $w_i(\cdot)$ to levels that differ by 1. Now consider another welfare measure $w'(\cdot)$ —call it the *shmutil* measure, where 1 util is equal to k shmutils. Suppose that, in violation of multiplicative invariance, our SWF ranks $w'(y) = kw(y) = (kw_1(y), \dots, kw_n(y))$ better than $w'(x) = kw(x) = (kw_1(x), \dots, kw_n(x))$. Then it appears to matter whether welfare is measured in utils or in shmutils. For $w'(x)$ and $w'(y)$ are just the lists of x 's and y 's

¹⁷I take the name from Thierry Marchand, "Scale invariance and similar invariance conditions for bankruptcy problems," *Soc Choice Welfare* 31 (2008): pp. 693–707. Adler calls this principle "Invariance to Common Ratio Rescalings" (p. 276). I have changed the name both for brevity and because, I will argue, the condition goes beyond invariance to what are intuitively mere rescalings of the welfare measure.

welfare levels in numbers of shmutils, whereas $w(x)$ and $w(y)$ are their lists in numbers of utils. So, to figure out which of x or y is really better, we must first figure out whether well-being should be measured in utils or in shmutils. That is ridiculous. It cannot plausibly matter whether welfare is measure in utils or shmutils, just as it cannot matter whether mass is measured in grams or kilograms. If we reject multiplicative invariance (while taking ratios of well-being levels to be meaningful), then we seem forced to conclude that welfare is measurable on an absolute scale, unlike even the most natural physical quantities with which we are familiar, and that there is a single number, waiting to be discovered, that represents your well-being. That seems implausible.

It turns out that Atkinson SWFs are the only members of the prioritarian family that are multiplicatively invariant.¹⁸ That is why Adler favors the Atkinson family of SWFs: he accepts multiplicative invariance as well as prioritarianism. We have seen, however, that Atkinson SWFs have unacceptable implications regarding miserable lives. There is some irony in this. Atkinson SWFs assume a ratio scale of well-being, which requires there to be a meaningful zero level. But Atkinson SWFs are unworkable or implausible precisely when we try to take that level into account, by considering lives at or below it.¹⁹

Atkinson SWFs' problems with negative well-being are instances of a more general problem: there simply is no prioritarian SWF that is multiplicatively invariant and can handle both positive and negative welfare levels.²⁰ To see this, consider well-being vector $(0, 0, 0, 0)$. Prioritarianism implies (by continuity) that, for some negative welfare level $-w$, $(0, 0, 0, 0)$ and $(1, -w, 0, 0)$ are equally good. Now double everyone's welfare in both distributions. Multiplicative invariance implies that $(0, 0, 0, 0)$ and $(2, -2w, 0, 0)$ must also be

¹⁸Except in variable-population cases when $\gamma = 1$. The sum of priority-weighted well-being, using the logarithm function, is as great in $(1, 1)$ as it is in $(1, 1, 1)$, but not as great in $(2, 2)$ as it is in $(2, 2, 2)$.

¹⁹Compare Charles Blackorby and David Donaldson, "Ratio-Scale and Translation-Scale Full Interpersonal Comparability without Domain Restrictions: Admissible Social-Evaluation Functions," *International Economic Review* 23 (1982): pp. 249–68, at 253.

²⁰See Campbell Brown, "Prioritarianism for Variable Populations," *Philosophical Studies* 134 (2007): pp. 325–61, for another proof and extensive discussion of this issue.

equally good. The Pigou-Dalton condition then implies that $(2, -2w, 0, 0)$ can be improved by transferring a unit of well-being from the first person to the third: $(1, -2w, 1, 0)$ is better than $(2, -2w, 0, 0)$. $(1, -2w, 1, 0)$ can be similarly improved by transferring w units of well-being from the fourth person to the second: $(1, -w, 1, -w)$ is better than $(1, -2w, 1, 0)$. So, by the transitivity of *better than*, $(1, -w, 1, -w)$ must be better than $(0, 0, 0, 0)$. But since $(0, 0, 0, 0)$ and $(1, -w, 0, 0)$ are supposed to be equally good, so must $(0, 0, 1, -w)$ and $(1, -w, 1, -w)$ by separability, since the last two people have the same welfare in each pair of distributions. By anonymity, $(1, -w, 0, 0)$ and $(0, 0, 1, -w)$ are equally good. So, by the transitivity of equal goodness, we have both that $(1, -w, 1, -w)$ and $(0, 0, 0, 0)$ are equally good and that the former is better—contradiction. This shows that, with a domain that includes both positive and negative welfare, prioritarian SWFs cannot be multiplicatively invariant.

We saw, in section 2, why a prioritarian SWF cannot plausibly be restricted to domains in which all welfare levels are positive. Those who are attracted to multiplicative invariance might see this as an easy argument against prioritarianism. But I am reluctant to reject prioritarianism on this basis. Even though some of the prioritarian axioms can be reasonably denied, they seem to me more ethically compelling than multiplicative invariance. And, as we will now see, multiplicative invariance rules out several other seemingly plausible SWFs. So even some opponents of prioritarianism may have reason to question multiplicative invariance.

4 Implications of Multiplicative Invariance

So far, it may seem as though multiplicative invariance is a niche issue in the SWF literature that should be of concern only to prioritarians such as Adler. But in fact it has wide-ranging implications for population ethics and distributive justice, and it is deeply rooted in the SWF framework. I draw out some of these implications in this section.

4.1 Levels

Multiplicative invariance implies that the only morally significant facts about welfare levels are the ratios between them and how they compare to the zero level—i.e., whether they are positive or negative. But we do not just categorize lives as good (worth living) or bad (worth not living). Some lives are very good, wonderful, or barely worth living. Others are very bad, utterly miserable, or nearly worth living. This information cannot be cashed out entirely in terms of ratios of well-being levels. For example, consider a wonderful life and a life that is barely worth living. It is not at all obvious that, if we decrease these well-being levels by a factor of n , the one life would remain wonderful and the other barely worth living. Intuitively, if n is very large, the wonderful life would be reduced to one that is barely worth living. Multiplicative invariance requires us to ignore this information.

According to some SWFs, this kind of information can be morally significant. The most obvious example is one we have already encountered: sufficientist SWFs. A sufficientist SWF gives priority—on some versions, absolute priority—to people who fall below a certain threshold level of well-being. Unless this threshold is just the zero level, sufficientist SWFs will violate multiplicative invariance.

To see this, suppose that we give some priority to the worse off only when they are below some sufficiently wonderful level of well-being, which is mapped by our individual well-being measure $w_i(\cdot)$ to 50. Suppose that we aggregate benefits above that threshold in a utilitarian fashion, giving no priority to the worse off who are sufficiently well off. Table 2 depicts the vectors assigned to four outcomes by a single well-being measure $w(\cdot)$.²¹

The distributions on the left, those of x and y , are represented by the vectors $w(x) = (45, 45)$ and $w(y) = (60, 30)$. Our sufficientist SWF judges x 's distribution to be better than

²¹A similar example is given by Michael Morreau and John A. Weymark, “Measurement scales and welfarist social choice,” *Journal of Mathematical Psychology* 75 (2016): pp. 127–36, though theirs involves absolute priority.

Table 2: Sufficientist SWF with nonzero threshold violates multiplicative invariance

	Ann	Bob		Ann	Bob
$w(x)$	45	45	$w(x')$	90	90
$w(y)$	60	30	$w(y')$	120	60

y 's, because they contain the same total well-being but Bob's well-being is below the threshold and is greater in x . The distributions on the right are those of x' and y' , in which both people's welfare levels are doubled. Using the same well-being measure $w(\cdot)$, these distributions are represented by the vectors $w(x') = (90, 90)$ and $w(y') = (120, 60)$. Our sufficientist SWF judges $w(x') = (90, 90)$ and $w(y') = (120, 60)$ to be equally good because both people are above the threshold and they contain the same total well-being. This violates multiplicative invariance: $w(y')$ is as good as $w(x')$, but $w(y')/2 = w(y)$ is not as good as $w(x')/2 = w(x)$.²²

Another violation of multiplicative invariance comes from views in population ethics. According to *critical-level utilitarianism*, there is a fixed, positive level of well-being—the *critical level*—below which a person's existence makes things worse, even if her life is worth living.²³ Critical-level utilitarians evaluate distributions by first subtracting the critical level from each person's well-being and then adding up these critical-level-adjusted values.

To see how critical-level utilitarianism violates multiplicative invariance, consider table 3. Suppose the critical level is strictly between 1 and 10, and let a blank cell represent a person's nonexistence. The left side of the table compares distributions $w(x) = (10)$ and $w(y) = (10, 1)$; the right side compares distributions $w(x') = (100)$ and $w(y') = (100, 10)$, in which everyone who exists is made ten times better off. Critical-level utilitarianism, with

²²Adler observes that sufficientists must reject multiplicative invariance in Adler, *Well-Being and Fair Distribution* [p. 347, n. 69; pp. 388–89]. But he rejects sufficientism for reasons independent of multiplicative invariance. His rejection of sufficientism appeals to the prioritarian axioms, in particular the Pigou-Dalton principle, which sufficientist SWF violate. But if multiplicative invariance is true, then one of the prioritarian axioms must be rejected anyway in order to handle negative welfare levels.

²³Adler discusses critical-level utilitarianism in section 7.1.3. See John Broome, *Weighing Lives* (Oxford; New York: Oxford University Press, 2004); Charles Blackorby, Walter Bossert, and David Donaldson, *Population Issues in Social Choice Theory, Welfare Economics, and Ethics* (Cambridge University Press, 2005).

a critical level between 1 and 10, judges $w(x)$ to be better than $w(y)$ because the only difference between them is the addition of a person in y below the critical level. But $w(x')$ will be worse than $w(y')$, because the additional person in $w(y')$ is above the critical level.²⁴

Table 3: Critical-level utilitarianism violates multiplicative invariance

	Ann	Bob		Ann	Bob
$w(x)$	10		$w(x')$	100	
$w(y)$	10	1	$w(y')$	100	10

We have seen two kinds of views that violate multiplicative invariance by assigning moral significance to some welfare level other than zero: sufficientist and critical-level views. Of course, these views face problems of their own. My point is not that multiplicative invariance must be false because one of these views is true. The point, rather, is that these views are live ethical options. It is not obviously mistaken or confused to care about whether people fall below some sufficient threshold of well-being or critical level of existence. If, as Adler believes, there is one meaningful and morally significant level of well-being—the zero level—then why could there not be more? Why insist that the level at which a life becomes worth living is morally significant, but that no other level could possibly be similarly significant? Maybe these views should ultimately be rejected. But I do not think they should be rejected on the grounds that, because they violate multiplicative invariance, they presuppose the existence of an absolute scale or natural unit of well-being.

²⁴Multiplicative invariance would also rule out critical-level *prioritarianism*, which first subtracts the priority-weighted value of some nonzero critical level from each person’s priority-weighted well-being, and compares distributions by the sums of those differences. Adler and Treich, ‘Prioritarianism and Climate Change’, endorse critical-level prioritarianism with the critical level set to the value of a neutral life, which is greater than zero (see footnote 8). The same paper defends the Atkinson weighting function by appealing to multiplicative invariance. Since critical-level prioritarianism violates multiplicative invariance in variable-population cases, and the motivation for multiplicative invariance appeals to claims about the measurability of well-being that do not seem sensitive to the fixed- vs. variable-population distinction, this situation strikes me as dialectically unstable.

4.2 Differences

We do not just regard some people as better off than others, or as being made n times better off (if we make ratio comparisons between welfare levels at all). We regard some as being *much* or *only slightly* better off than others, or as being made *much* or *only slightly* better off. Like the intuitive properties of lives at various welfare levels, this kind of information cannot be cashed out entirely in terms of ratios. For example, suppose that Ann is vastly better off than Bob. We do not think that, for any n , if their welfare were decreased by a factor of n , Ann would remain vastly better off than Bob. If n were sufficiently large that both welfare levels were made extremely small, we would judge Ann to be only slightly better off than Bob. And we might assign moral significance to these differences, in a way that violates multiplicative invariance.

This kind of information may be especially relevant for an egalitarian SWF. Larry Temkin suggests that the badness of inequality may not be linear with respect to its size.²⁵ For example, consider the distributions in table 4. Suppose we judge $w(x) = (0, 4)$ and $w(y) = (1, 1)$ to be equally good: though $w(x) = (0, 4)$ has a greater sum of well-being, $w(y) = (1, 1)$ has it more equally distributed. Multiplicative invariance would require that $w(x') = (0, 100)$ and $w(y') = (25, 25)$ then be equally good as well. But we might judge the inequality in $w(x') = (0, 100)$ to be more than 25 times as bad as the inequality in $w(x) = (0, 4)$. So we might judge $w(y')$ to be better than $w(x')$, in violation of multiplicative invariance.

Table 4: Nonlinear weighting of egalitarian complaints

	Ann	Bob		Ann	Bob
$w(x)$	0	4	$w(x')$	0	100
$w(y)$	1	1	$w(y')$	25	25

²⁵Larry S. Temkin, *Inequality* (New York: Oxford University Press, 1993).

For another egalitarian example, Temkin also suggests that inequality matters more at lower levels of well-being: other things being equal, a person’s complaint against being worse off than others is stronger the worse off she is. Now consider the distributions in table 5. We might judge $w(x) = (-10, -10)$ to be better than $w(y) = (-50, 100)$ on the grounds that Ann is much worse off than Bob. But we might judge the inequality in $w(y') = (-5, 10)$ to be much less bad the inequality in $w(y)$ because Ann is so much less badly off in $w(y')$. So, in the choice between $w(x') = (-1, -1)$ and $w(y') = (-5, 10)$, we might think it more important to benefit Bob so that he has a life worth living, despite the inequality in $w(y')$.

Table 5: Stronger complaints at lower levels

	Ann	Bob		Ann	Bob
$w(x)$	-10	-10	$w(x')$	-1	-1
$w(y)$	-50	100	$w(y')$	-5	10

More generally, egalitarian violations of multiplicative invariance are commonly associated with “absolute” measures of inequality, on which the magnitude of inequality depends only on the differences between people’s well-being levels. On an absolute measure, adding some quantity of well-being to each person can’t increase or decrease inequality, but multiplying everyone’s well-being by a common factor can. These are in contrast to “relative” measures of inequality, which are unaffected by rescalings by a common factor but affected by common translations (i.e., addition or subtraction). Serge-Christophe Kolm influentially called the absolute and relative measures “leftist” and “rightist,” respectively. He motivated the use of absolute measures with the example of the May 1968 student protests and workers’ strike in France. The strike ended with an agreement to increase everyone’s pay by the same percentage. “Thus, laborers earning 80 pounds a month received 10 pounds more, whereas executives who already earned 800 pounds a month received 100 pounds more. The Radicals felt bitter and cheated; in their view, this widely increased incomes inequality [*sic*].”²⁶ One

²⁶Serge-Christophe Kolm, “Unequal inequalities. II,” *Journal of Economic Theory* 13 (1976): pp. 82–111, at

might hold a similar view about proportional increases in well-being rather than income.

Egalitarians with an absolute measure of inequality, or even with an intermediate measure that combines absolute and relative elements, will reject multiplicative invariance, because multiplying everyone's welfare by a common factor can affect such measures. For example, in table 4, they will judge the inequality in $w(x') = (0, 100)$ to be greater than the inequality in $w(x) = (0, 4)$. In the choice between $w(x) = (0, 4)$ and $w(y) = (1, 1)$ they may prefer to double the sum of well-being in exchange for the slight inequality in $w(x)$. In the choice between $w(x') = (0, 100)$ and $w(y') = (25, 25)$, they may be unwilling to double the sum of well-being in exchange for the (on their view) greater inequality in $w(x')$. For similar reasons, in table 5, they might prefer the equal distribution $w(x) = (-10, -10)$ to the vastly unequal $w(y) = (-50, 100)$, while preferring the slightly unequal $w(y') = (-5, 10)$, which ensures that Bob at least has a life worth living, to $w(x') = (-1, -1)$, in which both lives are worth not living.

Again, my point is not that multiplicative invariance must be false because some (partially) absolute measure of inequality is correct, or because the badness of inequality is not linear with respect to its size or is worse at lower levels. The point, rather, is that these are live ethical options. I think it would be a mistake to rule them out on the grounds that, because they violate multiplicative invariance, they require an absolute scale or natural unit of well-being. For, as we're about to see, they require no such thing.

5 Welfare as a Dimensioned Quantity

We have seen that multiplicative invariance is quite restrictive: it is inconsistent with critical-level utilitarian, sufficientist, and some egalitarian SWFs, as well as prioritarian SWFs if welfare levels can be negative. But as we saw in section 3, multiplicative invariance appears to

419.

follow from the widely accepted view that welfare is, at most, ratio-scale measurable. There is no absolute scale or privileged unit of well-being. So we should be able to shrink or stretch our scale of well-being without changing the ranking of distributions, as multiplicative invariance implies. If our SWF violates multiplicative invariance, then figuring out the correct ranking of distributions requires us first to know whether well-being should be measured in utils or shmutils or something else, as if one of these units is uniquely correct. That is absurd.

There is something fishy about this argument from mere ratio-scale measurability to multiplicative invariance.²⁷ It is one thing to require that, say, halving the unit of welfare measurement should not change a SWF's ranking of distributions. It is, intuitively, quite another thing to require that doubling everyone's welfare should not change the SWF's ranking, as in the examples we considered in section 4. The SWF methodology appears unable to distinguish between these two cases. This inability to distinguish between real changes in welfare and mere changes of scale is observed by Amartya Sen in his groundbreaking work on the informational basis of social welfare evaluation. Sen notes that "in all cases of measurability-comparability frameworks discussed here (and in other works), the invariance requirement covers both interpretations [halving units and doubling welfare] since there is no natural 'unit' of measurement of personal welfare."²⁸

My aim in the rest of this essay is to sketch an alternative picture of the quantitative representation of well-being on which these two cases can be distinguished, and on which

²⁷This is carefully spelled out by Morreau and Weymark, 'Measurement scales and welfarist social choice'.

²⁸Amartya Sen, "On Weights and Measures: Informational Constraints in Social Welfare Analysis," *Econometrica* 45 (1977): pp. 1539–72, at 1542. In response to this problem, Morreau and Weymark, 'Measurement scales and welfarist social choice', suggest a framework of "scale-dependent" social welfare functionals (functions from profiles of individual well-being measures to social orderings). On their approach, given a scale of well-being, there is a unique social welfare functional. But if we use a different scale, we need to use a different social welfare functional as well. This allows each social welfare functional to violate conditions like multiplicative invariance, without implying that there is some single correct scale of welfare. But it seems to abandon the project of formulating a single function that ranks distributions no matter which of some informationally equivalent scales we use to represent those distributions. My proposal avoids this problem.

violations of multiplicative invariance do not require there to be an absolute scale or natural unit of well-being. We can have a SWF that violates multiplicative invariance without insisting that welfare is measurable on anything stronger than a ratio scale. The appearance to the contrary, I think, stems from another assumption that is deeply rooted in the SWF literature, which has to do with the values of the individual well-being measure.

The well-being measure $w(\cdot)$ takes an outcome x and returns a well-being vector $w(x) = (w_1(x), \dots, w_n(x))$. What are the components of this vector—i.e., the values of individual i 's well-being measure $w_i(\cdot)$? They are generally said to be *real numbers* (Adler calls them “well-being numbers”), and I have represented them accordingly in this paper so far. On this orthodox view, $w(x)$ and $w(y)$ are vectors of real numbers, corresponding to x 's and y 's distributions of welfare measured in (say) utils. Let $w'(x) = k w(x)$ and $w'(y) = k w(y)$ represent these same distributions, but measured in shmutils. Violations of multiplicative invariance mean that our SWF can rank x 's and y 's distributions in one way—say, x 's better than y 's—when fed $w(\cdot)$'s vectors of welfare-in-utils, but another when fed $w'(\cdot)$'s vectors of welfare-in-shmutils.

Why, though, should we treat the values of individual well-being measures as real numbers? In statements of physical laws, such as “Force = mass \times acceleration,” we do not treat “mass” as standing in for a number. It stands in for a *dimensioned quantity*, such as 5 kilograms. And dimensioned quantities are not real numbers.

Why not? For one thing, the real numbers are closed under multiplication and addition.²⁹ If dimensioned quantities such as 5 kg, 2 kg, and 3 kg were real numbers, then the product $5 \text{ kg} \times 2 \text{ kg} = 10 \text{ kg}^2$ would also be a real number. And so would the sum $10 \text{ kg}^2 + 3 \text{ kg}$. But the operation of adding mass to mass-squared is not even defined. I am not denying, of course, that we can convert these dimensioned quantities to real numbers—from quantities

²⁹This argument is given by George W Hart, *Multidimensional Analysis Algebras and Systems for Science and Engineering* (New York, NY: Springer New York, 1995).

of mass to *numbers-of-kilograms-of-mass*—and then multiply and add the resulting quantities: $5 \times 2 + 3 = 13$. But the resulting sum would not represent any dimensioned quantity. (One way to see this is that, if we instead multiplied and added *numbers-of-grams-of-mass*, we would get $5000 \times 2000 + 3000 = 10,003,000$, which relates to the earlier sum in no interesting way: $10,003,000/13 \approx 769,461.54$.)

Another reason to distinguish dimensioned quantities from numbers is to respect requirements of dimensional homogeneity: we can only compare quantities of the same dimension.³⁰ If dimensioned quantities such as 5 kilograms and 4 meters were real numbers, then one would be greater than the other or they would be the same. But no quantity of mass is greater than, less than, or equal to any quantity of length. Again, we can of course compare an object's number-of-kilograms-of-mass with its number-of-meters-of-height, but the former isn't a mass and the latter isn't a length.

The same considerations, I suggest, apply to quantities of welfare. There is no quantity of welfare that is the sum of Ann's welfare and the product of Bob's and Cat's welfare, because the operation of adding welfare-squared to welfare is undefined. And no one's welfare is greater than, less than, or equal to their height, or to π . Just as physical laws relate dimensioned quantities to one another—e.g., force to mass and acceleration, none of these being numbers—we should expect the SWF to relate dimensioned quantities to one another: namely, social welfare to the welfare of individuals.

What are these dimensioned quantities, if not real numbers? Take a particular example, such as one kilogram. One simple answer is that a kilogram is a particular equivalence class of objects under the relation of having equal mass: namely, the set of objects that have the same mass as the standard kilogram.³¹ This answer won't quite do. For suppose that nothing weighs exactly π kg, and that nothing is exactly π^2 m long. Then both equivalence

³⁰This is emphasized by Henry E. Kyburg, "Quantities, Magnitudes, and Numbers," *Philosophy of Science* 64 (1997): pp. 377–410.

³¹See Kyburg, 'Quantities, Magnitudes, and Numbers'.

classes are the empty set, so π kg would be identical to π^2 m. This can be avoided by taking the equivalence classes to be sets of *possible* objects (or object-world pairs), and by assuming that every quantity along each dimension can be instantiated. There is no assumption that these quantities are observable. I wish to remain neutral about other metaphysical questions about dimensions and dimensioned quantities. For example, are there fundamentally such quantities as 5 kg, or are there only fundamentally such relations as having greater mass or betweenness of mass?³² I suspect that any answer to these questions has to be able to do the work that we need from dimensioned quantities for present purposes.

I will make some familiar assumptions about the mathematical operations that can be performed on dimensioned quantities. Only quantities of the same dimension can be added, subtracted, or compared. We can multiply quantities of arbitrary dimensions to form quantities of yet other dimensions. We can also divide such quantities. The ratio of two quantities of the same dimension is a dimensionless number. A number divided by a dimensioned quantity yields a “reciprocal” dimensioned quantity. For example, the number 1 divided by 1 meter yields 1 meter⁻¹. I will assume that, for any possible quantity x with dimension $[x]$, and any positive integer n , there is a possible quantity $x^{1/n}$ with dimension $[x]^{1/n}$. Such a quantity, if multiplied by itself n times, yields x .³³

My suggestion is twofold. First, the values of the well-being measure $w(\cdot)$ should be vectors of dimensioned quantities rather than of numbers—e.g., (1 util, 2 utils) rather than (1, 2). And, second, our SWF must be properly expressed to reflect the proposed dimensionality of welfare. Some SWFs may need no adjustment—e.g., the utilitarian SWF, which

³²See Shamik Dasgupta, “Absolutism Vs Comparativism About Quantity,” *Oxford Studies in Metaphysics* 8 (2013): pp. 105–50.

³³I follow David H. Krantz et al., *Foundations of measurement*, 3 vol, 1 (New York: Academic Press, 1971), 460. They note: “One just never hears physicists speak of dimensions such as $M^{1/113}$. But that objection has little to do with roots as such; the problem is that one does not hear mention of M^{113} either. As we pointed out, only very small integers arise in practice. To our knowledge, no one has yet provided an axiomatization that imposes any limitation on the multiplying and dividing of dimensions, nor do we know of any explanation in physics for these limitations.”


simply adds up quantities of welfare. Others may require very simple adjustment. For example, the critical-level utilitarian SWF subtracts the value of a critical level c from each person's well-being before adding up these critical-level-adjusted values. c is often identified as a number, but it makes no sense to subtract a dimensionless number from a dimensioned quantity. c must therefore be a quantity of well-being, rather than a number. The dimensionality of welfare, however, may appear to pose a problem for other SWFs.

After noting the Atkinson SWF's problems with negative well-being, Adler mentions a family of prioritarian SWFs that play nicely with both positive and negative welfare: the family of Kolm–Pollak SWFs (p. 155, fn. 33). Kolm–Pollak SWFs use a negative exponential priority weighting function:

$$g(w) = -e^{-\lambda w}$$

where e is the base of the natural logarithm and λ is positive. Figure 2 shows an example of a Kolm–Pollak priority weighting function. Kolm–Pollak SWFs satisfy the axioms of prioritarianism over both positive and negative welfare levels (they do not even require a fixed zero level of well-being). But Adler rejects them because they violate multiplicative invariance.

Figure 2: Kolm–Pollak priority weighting function



I have suggested that w is a dimensioned quantity. And λ is generally treated as a number. Then the exponent $-\lambda w$ must be dimensioned quantity. But what does it mean to raise a number (e) to the power of a dimensioned quantity? Exponents cannot have dimension. Or consider, for another example, the Atkinson SWF with $\gamma = 1$: $g(w) = \log w$. It makes no sense to take the logarithm of a dimensioned quantity, such as 5 grams: there is no exponent such that raising the base (a number) to that exponent will yield 5 grams.³⁴

There is a simple solution to this apparent problem: we should interpret λ not as a real number, but as a *dimensional constant*, like critical-level utilitarianism's c . Where $[w]$ is the dimension of well-being, λ must have dimension $[w]^{-1}$ (reciprocal well-being). Then the exponent $-\lambda w$ will be a number, because any quantity of utils^{-1} multiplied by any quantity of utils is a number. (Some may prefer λ to have dimension $[w]$ and to rewrite the priority weighting function as $g(w) = -e^{-w/\lambda}$, where the function gives greater priority to the worse off as λ approaches zero.) The logarithmic SWF must be similarly clarified with the introduction of a dimensional constant within the argument of the logarithm.³⁵

I am not suggesting that there *is* such a dimensional constant—only that Kolm–Pollak SWFs require one if welfare levels are dimensioned quantities rather than numbers. Perhaps it is a disadvantage of a SWF that it posits additional dimensional constants that we have no independent reason to believe exist.³⁶ We might disfavor such a SWF on grounds of simplicity. But simplicity is just one theoretical consideration among others. If a dimensional constant is required by our most ethically attractive SWF, that may provide sufficient reason to accept it. Compare the case of physical laws. We posit a gravitational constant of

³⁴This argument is made by Chérif F. Matta et al., “Can One Take the Logarithm or the Sine of a Dimensioned Quantity or a Unit? Dimensional Analysis Involving Transcendental Functions,” *J. Chem. Educ.* 88 (2011): pp. 67–70.

³⁵An additional dimensional constant is also needed, in both the Kolm–Pollak and logarithmic cases, to ensure that priority-weighted well-being is a dimensioned quantity rather than dimensionless number. I ignore this complication in what follows.

³⁶See Bradford Skow, “How to Adjust Utility for Desert,” *Australasian Journal of Philosophy* 90 (2012): pp. 235–57, at 247.

dimension $[\text{force}] \cdot [\text{length}]^2 \cdot [\text{mass}]^{-2}$ because we know that the gravitational force between two bodies is proportional to the product of their masses and inversely proportional to the square of the distance between them. The gravitational constant is needed for force to be a dimensionally consistent function of these variables. We cannot reasonably have as much confidence in any particular SWF as we have in a physical law. But it would be unreasonable, I think, to reject an ethically attractive SWF purely on the grounds that it requires a dimensional constant.

If a SWF is properly expressed as a function of dimensioned quantities of welfare, then its rankings will not depend on the unit in which welfare is measured. Take the Kolm–Pollak SWF. Suppose we initially measure welfare in utils. Let $w = 1$ util and $\lambda = 1 \text{ util}^{-1}$. Then $g(w) = -e^1$. If we instead measure w in shmutils, so that $w = k$ shmutils, this will not change the value of $g(w)$: $k \text{ shmutils} \times 1 \text{ util}^{-1} = k \text{ shmutil} \times 1/k \text{ shmutils}^{-1} = 1$, so we still have $g(w) = -e^1$. This is because λ is not a number, but a dimensional constant. It has the same value regardless of the unit in which it is measured: $1 \text{ util}^{-1} = 1/k \text{ shmutils}^{-1}$. This ensures the Kolm–Pollak SWF is dimensionally invariant—i.e., that its truth does not depend on the unit of measurement.

Let me give another example that does not involve well-being. Suppose Ann and Bob each plant two trees, which start at 1 meter tall and double in height each year until they are chopped down. Ann and Bob get to decide when to chop down their trees, and each aims to have the greatest average tree height. Ann grows one tree for a year and the other tree for two years. Bob grows both of his trees for a year and nine months. Bob's trees will end up, on average, taller than Ann's. Suppose next that Ann grows one tree for 12 years and the other for 24 years. Bob grows his for 21 years. Ann's trees will end up, on average, taller than Bob's. Since the longer growth periods are just twelve times the shorter ones, does this mean that the formula for tree growth with respect to time grown presupposes a stronger-than-ratio (i.e., absolute) scale of temporal duration? Of course not. A month is just as good a unit of

time as a year.

To unpack the analogy a bit more, let $h(t)$ be the height of a tree grown for time t . It may be tempting to write $h(t) = h(0)2^t$. But since t is a dimensioned quantity, that cannot strictly be correct. It should instead be $h(t) = h(0)2^{t/(1 \text{ year})}$. This introduction of a dimensional constant would be unnecessary if t were treated as a number—i.e., time-in-years rather than time—as the orthodox SWF methodology does with well-being. But then the formula would not be invariant to changes in the measure of time, as the example in the previous paragraph illustrates. It would be absurd to suggest that since the formula would not be invariant to changes in the measure of time, there must be a single correct unit of time if a tree really doubles in height every year. The impression of requiring times to have “correct” units only arises from treating times as having no unit as all—i.e., as dimensionless numbers—which effectively regards them as if they were measurable on an absolute scale. Seen this way, Adler’s rejection of Kolm–Pollak SWFs on the grounds that a ratio scale of welfare requires multiplicative invariance is like insisting from the armchair that nothing can grow or decay exponentially with respect to time because temporal duration is no more than ratio-scale measurable.

We can now see, more generally, why recognizing welfare levels as dimensioned quantities allows a SWF to violate multiplicative invariance even if welfare is only ratio-scale measurable. Consider how Adler motivates multiplicative invariance:

Assume that an SWF ranks outcomes differently, depending on whether well-being values are assigned using $w(\cdot)$ or $w^+(\cdot)$, even though $w^+(\cdot)$ is a common ratio rescaling of $w(\cdot)$. The two measures contain identical information regarding levels, differences, and ratios, but the SWF is “reading” the well-being numbers as if they embody additional information, and thereby differentiating between $w(\cdot)$ and $w^+(\cdot)$. This is, intuitively, problematic. What additional well-

being information can there be? (p. 156)

We can now answer this argument. Recognizing welfare levels as dimensioned quantities allows us to distinguish two kinds of operations that might be called “common ratio rescalings.” One is changing the numerical component used to represent a dimensioned quantity of welfare—e.g., converting utils to shmutils. This kind of rescaling involves “multiplying” welfare levels by the conversion ratio k shmutils/1 util = 1. This operation preserves, as Adler says, all of the information regarding levels, differences, and ratios of well-being. If the SWF is properly expressed as a function of dimensioned quantities of well-being, with the proper dimensional constants, this will not change the ranking of distributions, as we have seen with the Kolm–Pollak SWF. This is because, if $w^+(\cdot)$ is a “common ratio rescaling” of $w(\cdot)$ in this sense, then $w^+(\cdot) = w(\cdot)$: the two measures assign the same vectors of *dimensioned quantities* to the same outcomes, even if those quantities are represented in different units (much as 1 kg = 1000 g). So of course the SWF will rank the resulting vectors the same way.

The second kind of operation is multiplying welfare levels by a dimensionless number $k \neq 1$. This does not, contra Adler, preserve all of the information regarding levels and differences: people’s well-being levels and the differences between them have changed. Then a SWF that is not multiplicatively invariant, such as the Kolm–Pollak SWF, might rank distributions differently. But why shouldn’t it? This kind of “common ratio rescaling” is a real change in well-being, not a mere change in scale. The absence of a natural unit of well-being gives us no reason to doubt that such a change could affect the ranking of distributions—any more than the absence of a natural unit of time gives us reason to doubt that atoms could decay exponentially with time. It is only because the SWF tradition represents welfare levels as numbers rather than dimensioned quantities that there is any temptation to conflate these two cases of “common ratio rescalings.”

Adler might insist that all of the *relevant* information about well-being levels and dif-

ferences is preserved by this second kind of “common ratio rescaling,” because, contrary to the views considered in section 4, the only information that could be relevant is captured by ratios. That may be true. But it does not follow from the mere ratio-scale measurability of welfare. Consider an analogy to another dimensioned quantity, namely temperature. The Celsius scale assigns the boiling point of water to 100°. We might reasonably care whether some amount of water is above or below 100° Celsius—say, because we want the water to boil. Does caring about such information mean that we must be sensitive to the scale on which temperature is measured, or that we are bizarrely taking a degree Celsius to be the correct unit of temperature? Of course not. The Celsius scale is uncontroversially not an absolute (or even ratio) scale. But that does not mean that if every object’s temperature-in-degrees-Celsius were doubled and increased by ten, everything would stand in all of the same relations that depend in some way on temperature (e.g., being in the same state of matter).

Multiplicative invariance simply does not follow from Adler’s insight that, since welfare is (at most) ratio-scale measurable, a SWF’s rankings should be independent of the unit in which well-being is measured. It only appears to follow when we mistakenly treat the values of the individual well-being measure as numbers rather than dimensioned quantities and assume that the SWF contains no dimensional constants. Since well-being levels are not numbers, prioritarrians and others can reasonably reject multiplicative invariance without being committed to the existence of an absolute scale or natural unit of well-being. They can do this by treating welfare levels as dimensioned quantities and thereby distinguishing between real changes in well-being and mere changes in scale, without having to change the SWF when well-being is represented in different units. Of course, there may be other reasons to reject SWFs that violate multiplicative invariance—e.g., because they posit additional dimensional constants like the Kolm–Pollak SWF’s λ or the critical level c . But proponents of those SWFs were already committed to such additional parameters; they may

have simply ignored their dimension.

6 Conclusion

I am suggesting that Adler and others in the SWF tradition have made what may seem a simple mistake: they have failed to heed the pedantic advice of our science teachers, to mind our units. If I am right, then why has this mistake been so widely made?

One speculative hypothesis is that the treatment of welfare levels as numbers, rather than dimensioned quantities, is an artifact of the ordinalist framework from which much of the SWF literature originally stems. If well-being were only ordinally measurable, talk of “units” of well-being might seem misleading. What would be the point of representing well-being levels in utils, for instance, if 2 utils were not twice the value of 1 util? It would perhaps seem better to assign numbers without including any unit, to avoid the impression of a cardinal or even ratio scale. But mere ordinal measurability is no justification to treat well-being levels as if they were dimensionless numbers. The Mohs scale is an ordinal scale of mineral hardness, but Mohs levels are not numbers; they are levels of hardness. In any case, those who take welfare to be measurable on a ratio scale—such as Adler—should find it less tempting to omit units.

Another possibility is that theorists in the SWF tradition have been under the influence of something like Luce’s “principle of theory construction” in the philosophy of science.³⁷ Luce’s principle says that admissible transformations of independent variables (e.g., ratio rescalings of the individual welfare measure) should lead an admissible transformation of the dependent variable. But Luce himself qualifies this principle to functional relationships in which there are no dimensional constants. Luce takes his principle of theory construction, properly understood, to imply that “psychologists (as well as other scientists)

³⁷See R. Duncan Luce, “On the Possible Psychophysical Laws,” *Psychological Review* 66 (1959): pp. 81–95.

either are restricted to a very few possible types of laws ...*or* they cannot avoid including dimensional parameters in the statement of their laws.”³⁸ Multiplicative invariance and other restrictive conditions on the shape of the SWF can be avoided by including such dimensional constants, which we may have independent reason to recognize.³⁹

Here I have only drawn out a few consequences of treating well-being levels as dimensioned quantities. It lets us distinguish between real changes of well-being from mere changes in the unit of measurement, without requiring the SWF to change when welfare is represented on a different scale. It highlights the importance of dimensional constants for SWFs. It opens the possibility for welfare levels and differences to have moral significance beyond the ratios between them, without requiring a privileged, natural unit of well-being. And it would allow Adler to apply a prioritarian SWF to distributions that contain negative welfare levels, without going beyond a ratio scale. I do not know what other implications this approach may have for the SWF methodology.

³⁸R. Duncan Luce, “Comments on Rozeboom’s Criticisms of ‘On the Possible Psychophysical Laws,’” *Psychological Review* 69 (1962): pp. 548–51, at 550, emphasis mine.

³⁹Skow, ‘How to Adjust Utility for Desert’, suggests that the dimension of social welfare is distinct from that of individual welfare. This seems plausible because the goodness of an individual life does not seem comparable with a distribution of well-being. This means that even utilitarian SWFs may require a dimensional constant to relate individual welfare to social welfare.