

EQUITY BY THE NUMBERS: MEASURING POVERTY, INEQUALITY, AND INJUSTICE

*Matthew D. Adler**

INTRODUCTION	551
I. EQUITY METRICS: AN OVERVIEW	559
<i>A. Inequality Metrics</i>	560
<i>B. Social Welfare Functions</i>	566
<i>C. Poverty Metrics</i>	569
<i>D. Social-Gradient Metrics</i>	573
<i>E. A Summary</i>	576
II. WHY THE PIGOU-DALTON PRINCIPLE? A GENERIC JUSTIFICATION	579
III. WHAT IS THE BEST CURRENCY FOR THE PIGOU-DALTON PRINCIPLE?.....	583
IV. SHOULD THE PIGOU-DALTON PRINCIPLE BE “RESTRICTED”? HEREIN OF POVERTY AND SOCIAL-GRADIENT METRICS.....	592
<i>A. Poverty Metrics</i>	593
<i>B. Social-Gradient Metrics</i>	596
V. EQUALITY OF OPPORTUNITY METRICS.....	599
CONCLUSION.....	606

INTRODUCTION

Can we measure inequity? A naïve response denies the possibility of quantification. And yet scholarly and popular discourse abounds with numbers purporting to measure inequality, poverty, and other aspects of equity. A more sophisticated answer acknowledges the wide use of *equity metrics*, but stresses their plurality. “How to quantify unfairness is a value choice. Different numbers mirror different ‘takes’ on the meaning of equity. End of story.”

A yet more nuanced response sees structure in the plurality of equity metrics. There are deep, unifying axioms, but divergence with respect to the application and specification of those axioms. Seeing both the deep consensus, and the differences in specification, helps us think clearly about the normative position presupposed by the use of one or another metric. That, at least, is the story I will tell here.

Has the United States become a less equitable society over the last thirty-five years, since the 1970s? Answers to that highly salient question illustrate the plurality of tools for quantifying equity. A classic *inequality* metric (such as the Gini coefficient, or the variance-based “coefficient of variation”) quantifies the population-wide distribution of some attribute.¹ But which attribute? The inequality of periodic *income* (in particular, annual income) has increased since the 1970s. This is true for different definitions of “income.” It is also robust to the choice of inequality metric, although metrics that are especially sensitive to what happens at the top of the income distribution will record a more dramatic change over the last thirty-five years.²

Wage and *wealth* inequality have also increased.³ But if we move from income, wages, and wealth, to a different indicator of economic well-

* Richard A. Horvitz Professor of Law and Professor of Economics, Philosophy, and Public Policy, Duke University. This Article is based upon my Meador Lecture at the University of Alabama School of Law. Thanks to the dean and faculty of the School of Law for the generous invitation to give a lecture, and to the editors of the *Alabama Law Review* in helping to prepare this article for publication.

1. See *infra* Section I.A (describing structure of inequality metrics).

2. Two detailed studies of U.S. income inequality over the last several decades are REBECCA M. BLANK, CHANGING INEQUALITY (2011); Jonathan Heathcote et al., *Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006*, 13 REV. OF ECON. DYNAMICS 15 (2010). Heathcote et al. look at a variety of definitions of income (including wage earnings, pre-tax income, and disposable income) and a variety of inequality metrics (including the Gini coefficient, variance of logarithm, and percentile ratios). They generally find increasing inequality since the 1970s (indeed, since the beginning of their time series, in 1967) as illustrated by their figures 8, 9, 11, and 12. Blank examines inequality over the period 1979–2007, looking specifically at earnings and “total income” (the sum of earnings, government income, and unearned income from other sources), and using the Gini, coefficient of variation, and percentile ratios as her inequality metrics. When she looks at all adults 18–64 (including nonworking household members allocated a share of household income), she finds an increase in all inequality metrics applied to both total annual income and annual earnings. See BLANK, *supra*, at 64–65 tbl.5; *cf. id.* 30–31 tbl.2 (finding more mixed picture with respect to earnings of all workers). See also Jeffrey Thompson & Timothy M. Smeeding, *Country Case Study—USA*, in THE GREAT RECESSION AND THE DISTRIBUTION OF HOUSEHOLD INCOME 202 (Stephen P. Jenkins et al. eds., 2013) (measuring income inequality with Gini coefficient and percentile ratios, as well as top percentile shares, and finding increasing inequality from 1979 through first years of Great Recession).

These findings of increasing U.S. income inequality are confirmed in a recent detailed comparative analysis that includes U.S. time trends. See Andrea Brandolini & Timothy M. Smeeding, *Income Inequality in Richer and OECD Countries*, in THE OXFORD HANDBOOK OF ECONOMIC INEQUALITY 71, 83 fig.4.2, 88 fig.4.5 (Wiemer Salverda et al. eds., 2009). Atkinson, Piketty, and Saez, using measures of inequality focused on the disparity between top incomes and the rest of the population (in particular, top percentile shares), find a very dramatic increase in U.S. income inequality from the 1970s through 2007. See Anthony B. Atkinson, Thomas Piketty & Emmanuel Saez, *Top Incomes in the Long Run of History*, 49 J. ECON. LIT. 3 figs.1, 2 & 3 (2011). Piketty, in his influential recent book, extends the timeline through the Great Recession; a dramatic increase in top percentile income shares from the 1970s through 2010 can be observed in figures 8.5 through 8.8. See THOMAS PIKETTY, CAPITAL IN THE TWENTY-FIRST CENTURY 291–300 (Arthur Goldhammer trans., 2014). The figures in Thompson and Smeeding, *supra*, at 216–217, tell the same story.

3. On wage inequality, see BLANK, *supra* note 2, at 30–31 tbl.2; Heathcote, *supra* note 2, at 23 fig.4. On wealth inequality, see Fabian T. Pfeffer, Sheldon Danziger & Robert Schoeni, *Wealth*

being—“consumption,” i.e., individual expenditure on goods and services—the picture is more muddled. A number of studies find that inequality in the population-wide distribution of consumption in the United States has increased much less, since the 1970s, than inequality in the distribution of income⁴—although these findings have been contested.⁵ Consumption tends to be more stable than income, since individuals can borrow and save to “smooth out” transitory changes to income. Relatedly, an individual’s *lifetime* income is arguably better proxied by her periodic (e.g., annual) consumption than by her periodic income. Thus, it may be the case that *lifetime* income inequality increased substantially less during the period 1980–present, as compared to *annual* income inequality.⁶

What about the poor? Has the degree of poverty gone up or down? Poverty metrics, unlike inequality metrics, are “truncated” at the poverty line: they ignore the distribution of the relevant attribute above that line.⁷ Traditionally, scholarly discussion of poverty—and certainly official poverty measurement by the U.S. government—has focused on *income* as the relevant attribute, with the poverty line set either by identifying the income required to meet certain needs (for example, nutritional needs), or as some fraction of population mean or median income. The *headcount ratio* (the fraction of the population with incomes below the poverty line) is the simplest measure of the degree of income poverty.⁸ The headcount ratio, combined with the official U.S. definition of “income” and official poverty line, yields an oscillating pattern: from 13% (1980), to 13.5% (1990), to 11.3% (2000), to 15.1% (2010).⁹ Using some alternative definitions of “income,” Meyer and Sullivan find an oscillating pattern that

Disparities Before and After the Great Recession, 650 ANNALS AM. ACAD. POL. & SOC. SCI. 98 (2013); Heathcote, *supra* note 2, at 42 fig.19.

4. See ORAZIO P. ATTANASIO, ERICH BATTISTIN & MARIO PADULA, *INEQUALITY IN LIVING STANDARDS SINCE 1980* (2011); Richard Blundell et al., *Consumption Inequality and Partial Insurance*, 98 AM. ECON. REV. 1887 (2008); David S. Johnson et al., *Economic Inequality Through the Prisms of Income and Consumption*, 128 MONTHLY LAB. REV. 11 (2005); Dirk Krueger & Fabrizio Perri, *Does Income Inequality Lead to Consumption Inequality? Evidence and Theory*, 73 REV. ECON. STUD. 163 (2006); Heathcote, *supra* note 2; Bruce D. Meyer & James X. Sullivan, *Consumption and Income Inequality in the U.S. Since the 1960s* (Apr. 2, 2013) (working paper) (<http://www3.nd.edu/~jsulliv4/Inequality3.6.pdf>).

5. See Mark A. Aguiar & Mark Bils, *Has Consumption Inequality Mirrored Income Inequality?* (Dec. 23, 2013) (unpublished manuscript), *available at* http://scholar.princeton.edu/sites/default/files/mirror_0.pdf (last visited Nov. 21, 2014); Orazio Attanasio et al., *The Evolution of Income, Consumption, and Leisure Inequality in the U.S., 1980–2010* (Nat’l Bureau of Econ. Research, Working Paper No. 17982, 2012).

6. Cf. Jeremy Arkes, *Trends in Long-Run Versus Cross-Section Earnings Inequality in the 1970s and 1980s*, 44 REV. INCOME & WEALTH 199 (1998) (finding Gini coefficient of five-year earnings to be lower than average Gini coefficient of annual earnings during the 1970s and 1980s).

7. See *infra* Section I.C (describing structure of poverty metrics).

8. Buhong Zheng, *Aggregate Poverty Measures*, 11 J. ECON. SURV. 123, 124 (1997).

9. See Bruce D. Meyer & James X. Sullivan, *Winning the War: Poverty from the Great Society to the Great Recession*, in BROOKINGS PAPERS ON ECONOMIC ACTIVITY 133, 150–51 tbl.1 (2012).

yields a net decrease in the headcount ratio from 1980 to 2010; and, switching to *consumption* poverty, a larger decrease.¹⁰

To be sure, it is deeply troubling (as a moral matter) and puzzling (as a matter of social science) that poverty in the United States has not decreased very substantially over the last thirty-five years, despite dramatic gains in GDP per capita.¹¹ The point here is that poverty metrics are insensitive to distributional patterns *above* the poverty line. Thus, income inequality metrics, which *are* sensitive to such patterns, see a time trend of increasing inequity, while poverty metrics see a flat or perhaps somewhat downward trend.

The picture becomes yet more muddled if we move away from indicators of economic well-being: income, wealth, consumption, wages. One such indicator is *longevity*. Peltzman examined the population distribution of lifespan in the United States, starting in the 1740s through 2002. Using the Gini coefficient to measure the inequality of lifespan, he found a dramatic decrease—with continuing declines in lifespan inequality through the 1980s and 1990s.¹²

The cumulative reduction of mortality inequality is startling. In the first century of our data, mortality Ginis range around 40 to 50. That is, they exceed the upper range of household income Ginis that prevail today in the developed world. The decline in mortality inequality since the mid nineteenth century is hardly interrupted by either of the two world wars This decline has taken the mortality Ginis of today down to levels that are much lower than—on the order of half—the lowest contemporary income Ginis in these countries.¹³

Edwards and Tuljapurkar reach broadly similar conclusions, using a variance-based measure of lifespan inequality.¹⁴

A different indicator of individual well-being is *happiness*, as determined by surveys asking individuals, “How happy are you on a scale of 1 to 5?” or some similar question. Looking at the variance in the population-wide distribution of answers to these happiness surveys, Stevenson and Wolfers find that “happiness inequality fell sharply during

10. *Id.* For earlier studies—using alternative measures and finding oscillating poverty over the decades after 1980—see KEVIN LANG, POVERTY AND DISCRIMINATION 87 fig.4.5 (2007); Thesia I. Garner & Kathleen S. Short, *Identifying the Poor: Poverty Measurement for the U.S. from 1996 to 2005*, 56 REV. INCOME & WEALTH 237, 252 fig.2 (2010).

11. See LANG, *supra* note 10, at 89 fig.4.6.

12. Sam Peltzman, *Mortality Inequality*, 23 J. ECON. PERSP. 175 (2009).

13. *Id.* at 181.

14. See Ryan D. Edwards & Shripad Tuljapurkar, *Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries*, 31 POPULATION & DEV. REV. 645 (2005).

the 1970s and continued to fall in the 1980s, before rising slightly in the 1990s and 2000s.”¹⁵

While inequality metrics look at the population-wide distribution of some attribute (income, wealth, happiness, longevity, wages, etc.), and poverty metrics focus on the poverty-line-truncated distribution, yet a third—quite prevalent—approach examines *disparities* between socially salient groups.¹⁶ Call this the “social gradient” approach. Much work in this vein, particularly in the United States, looks at racial disparities. How serious are racial skews in income, wealth, longevity, educational attainment, etc.? A different strand in social-gradient work looks at socioeconomic differences, in particular at the correlation between socioeconomic status and some non-economic indicator (often, health).¹⁷

Thus, the question, “Has the United States become a less equitable society?” can be construed in social-gradient terms—more specifically, in racial-disparity terms. Have racial disparities in economic or non-economic indicators increased or decreased since the 1970s? The answer to that question is complex. A recent study of racial disparities in *health* looked at black–white differentials with respect to 15 health measures over the period 1990 to 2005, finding that disparities narrowed for 7 of the 15, and increased for 5.¹⁸ As for *income*, it is unclear whether the decline in racial disparities in income that occurred prior to 1980¹⁹ has continued since then. Black–white wage differentials seem to have increased over the period

15. Betsey Stevenson & Justin Wolfers, *Happiness Inequality in the United States*, 37(2) J. LEGAL STUD. S33, S51 (2008). See also Andrew E. Clark et al., *The Great Happiness Moderation* 37 tbl.4.D (IZA Discussion Paper No. 6761) (2012). Dutta and Foster, using ordinal inequality indices, reach the following conclusion: “In terms of broad trends, happiness inequality [in the U.S.] decreased from its highest level in the 1970s, through the 1980s and 1990s. Only in the 2000s did it start to rise again. However, in 2010 there has been a remarkable decline in inequality” Indranil Dutta & James Foster, *Inequality of Happiness in the U.S.: 1972–2010*, 59 REV. INCOME & WEALTH 393, 413 (2013).

16. See *infra* Section I.D.

17. See generally C.J.L. Murray et al., *Health Inequalities and Social Group Differences: What Should We Measure?*, 77 BULL. WORLD HEALTH ORG. 537, 538 (1999).

18. Jennifer M. Orsi et al., *Black–White Health Disparities in the United States and Chicago: A 15-Year Progress Analysis*, 100 AM. J. PUB. HEALTH 349, 351–352 (2010). More precisely, statistically significant improvements were observed with respect to 7 measures (all-cause mortality, cancer mortality, lung cancer mortality, motor vehicle crash mortality, percentage low birth weight babies, percentage no prenatal cancer in the first trimester, and primary and secondary syphilis case rate); and statistically significant widenings were observed with respect to 5 (heart disease mortality, female breast cancer mortality, diabetes mortality, suicide mortality, and tuberculosis case rate).

Using data on self-rated health from the Panel Study of Income Dynamics, Rohde and Guest find a reduction in black–white health inequality from 1990 through 2007. Nicholas Rohde & Ross Guest, *Multidimensional Racial Inequality in the United States*, 114 SOC. INDICATORS RES. 591, 597–599 (2013).

19. See T. Kirk White, *Initial Conditions at Emancipation: The Long-Run Effect on Black–White Wealth and Earnings Inequality*, 31 J. ECON. DYNAMICS & CONTROL 3370 (2007).

1980 to 2007.²⁰ Finally, Stevenson and Wolfers observe a large decrease in the black–white disparity in *happiness*.²¹

In short, different generic approaches to measuring equity, and different choices of the relevant attribute (income, health, happiness, etc.), can yield strikingly different stories about the direction and magnitude of the change in inequity in the United States over the last thirty-five years. And the same would be true if we shifted our focus to *global* inequity,²² or to a different time period.²³

This observation—I should stress—is *not* meant to downgrade the significance of the increase in U.S. income inequality. A high level of income inequality distorts the democratic process,²⁴ and may well cause other sorts of inequities.²⁵ Every citizen should be deeply concerned that the top 10% of the income distribution captures almost one half of the total, and that the top 1% captures one-fifth—returning us to disparities not seen in the United States since the 1920s.²⁶ Income inequality is not merely intrinsically relevant—as one way of capturing inequity—but has a wide range of causal impacts. The focus of this Article is the measurement of inequity, not the (much larger) question of assessing the causal impacts of income inequality or other kinds of inequities. The reader should certainly be alive to that question, even though it is not one that I will be addressing here.

In any event, what I have illustrated to this point—using a mini-case-study of time trends in the United States over the last thirty-five years—is that equity *can* be quantified (e.g., via inequality, poverty, or social-gradient metrics), but that the choice of metric can greatly affect the analyst’s conclusions (here, about the time trend). Of course, this is not

20. Jake Rosenfeld & Meredith Kleykamp, *Organized Labor and Racial Wage Inequality in the United States*, 117 AM. J. SOC. 1460, 1468 (2012).

21. See Stevenson & Wolfers, *supra* note 15, at S61 fig.6. Yang also finds a shrinking of the black–white happiness gap (albeit more modest than that observed by Stevenson and Wolfers) in *Social Inequalities in Happiness in the United States, 1972 to 2004: An Age–Period–Cohort Analysis*, 73 AM. SOC. REV. 204 (2008); and *Long and Happy Living: Trends and Patterns of Happy Life Expectancy in the U.S., 1970–2000*, 37 SOC. SCI. RES. 1235 (2008).

22. See Koen Decancq et al., *The Evolution of World Inequality in Well-Being*, 37 WORLD DEV. 11 (2009); Mark McGillivray & Nora Markova, *Global Inequality in Well-Being Dimensions*, 46 J. DEV. STUD. 371 (2010).

23. See Clayne Pope, *Measuring the Distribution of Material Well-Being: U.S. Trends*, 56 J. MONETARY ECON. 66 (2009) (finding divergent long-term trends with respect to income versus non-income inequality).

24. See, e.g., LARRY M. BARTELS, *UNEQUAL DEMOCRACY* (2008).

25. Richard Wilkinson and Kate Pickett have documented many correlations between income inequality and social ills. See RICHARD WILKINSON & KATE PICKETT, *THE SPIRIT LEVEL* (2009). Some of these correlations surely reflect a causal linkage *from* income inequality to the correlated item—although teasing apart mere correlation and causation is not easy.

26. PIKETTY, *supra* note 2, at 291–300 figs.8.5, 8.6, 8.7 & 8.8.

really surprising. What is less obvious is the deep connection among equity metrics that I asserted at the beginning of this Introduction.

The central claim of this Article is that the *Pigou-Dalton principle* constitutes that deep connection. The Pigou-Dalton principle both unifies equity metrics, and differentiates them insofar as different metrics adopt divergent specifications of the principle. The Pigou-Dalton principle is the linchpin of equity, just as the Pareto principle is the linchpin of efficiency.²⁷ So, first, the Article aims to popularize Pigou-Dalton. How many law professors have never heard of the Pareto principle? How many *are* familiar with Pigou-Dalton? And yet any law professor (or economist or scholar of public health or . . .) who cares about equity ought to know what the principle says.

So what it says is:

The Pigou-Dalton Principle

Assume that, in scenario x , one person has more of valuable attribute C than a second person. In scenario y , the first person's level of C has decreased by some amount, and the second person's level of C has increased by *the very same amount*. In other words, a “pure” (non-leaky) transfer of the attribute from the first person to the second has occurred. However, this transfer is not so large as to reverse their positions: in scenario y , the first person still has more C or an equal amount of C as the second person. Finally, these are the only two individuals affected (with respect to C) by the move from x to y . Everyone else's level of C in x is the same as her level of C in y .²⁸

Under these conditions, scenario y is *more equitable* than scenario x . We have reduced the gap between the first person's holdings of C and the second's, and have done so via a pure transfer—so that

27. The principle can be traced to the work of A.C. Pigou and Hugh Dalton, and in the various literatures on equity discussed below it is regularly given the name “Pigou-Dalton” (but not always, sometimes instead being labeled, e.g., as the “principle of transfers”). See A.C. PIGOU, *WEALTH AND WELFARE* 24 (1912); Hugh Dalton, *The Measurement of the Inequality of Incomes*, 30 *ECON. J.* 348, 351 (1920).

28. Sometimes, the principle is formulated more broadly—allowing for a transfer that does reverse the two individuals' positions, but diminishes the difference in their holdings of C . It is more straightforward, I believe, to provide a direct normative “story” for the narrower principle. See MATTHEW D. ADLER, *WELL-BEING AND FAIR DISTRIBUTION: BEYOND COST-BENEFIT ANALYSIS* 339–40 (2012); *infra* Part II. However, if the currency C is well-being or some good proxy for well-being, the two principles are equally normatively attractive, since the broader principle follows from the narrower one given an axiom of “anonymity”: requiring that distributions of C which are the same except for *who* receives various amounts be seen as equally equitable. If C is well-being (at least), the anonymity axiom seems very powerful. See ADLER, *supra*, at 52.

the first person cannot complain, as a matter of equity, about the disparity between his loss and the other individual's gain.

While its cousin, Pareto, comes "pre-specified" in terms of preference-satisfaction, the Pigou-Dalton principle is unspecified. For what *is* the valuable item *C*? Is it income? Utility? Happiness? Health?

Part I of the Article will show how the Pigou-Dalton principle, *in some form*, underlies a wide range of equity metrics. I focus here on four families of metrics, each the basis for much current scholarship as well as information gathering by official statistical bureaus. Those families are inequality metrics, poverty metrics, social-gradient metrics, and social welfare functions. My survey of these tools will describe both traditional approaches and newer, innovative techniques within the four families—namely, so-called "multidimensional" poverty and inequality metrics, inspired by Amartya Sen and Martha Nussbaum's work on "capabilities." Part I will explain how each of the four kinds of metrics satisfies the Pigou-Dalton principle in terms of *some* "currency": some specification of the valuable attribute *C*.

Part II will explore the normative defensibility of the Pigou-Dalton principle. Why *does* so much current thinking about equity presuppose that principle? I offer one possible justification for the principle: a justification which is "welfarist" in focusing on the connection between Pigou-Dalton transfers and well-being and presupposing the possibility of interpersonal well-being comparisons. Although alternative defenses of the principle are certainly possible, it is the welfarist approach that I believe most convincing, and present in Part II.

Part III and Part IV explore difficult questions about the specification of the principle. Two key questions emerge from the survey of the four families in Part I. First, what *is* the appropriate "currency" for the Pigou-Dalton principle? Second, should the principle be applied in a *restricted* form? As we shall see, both poverty metrics and social-gradient metrics—in their own ways—limit the scope of the Pigou-Dalton principle. Poverty metrics do not satisfy the principle with respect to transfers above the poverty line, and social-gradient metrics only satisfy it for transfers from higher to lower social status individuals.

Plausibly, the best defense of the social-class-restricted version of the Pigou-Dalton principle relates to considerations of individual responsibility. If Nadja is *responsible* for being worse off than Juan—if she has frittered away her resources, while Juan has been prudent—then equity no longer favors an improvement in her holdings at Juan's expense. Now imagine that Nadja is worse off than Juan *because of* her lower social status. Surely *that* is not her responsibility. By restricting the Pigou-Dalton principle to transfers from higher- to lower-social-status individuals, social-

gradient metrics work to ensure that considerations of individual responsibility do not vitiate the normative case for a transfer.

However, this defense of social-gradient metrics may be challenged. Social class is a crude indicator of individual responsibility. The differing statuses of Nadja and Juan may not, in fact, have been the cause of their differential holdings of *C*. Moreover, Robert may have the same social status as Juan, yet have less *C* without being responsible for this differential. Shouldn't the problem of responsibility be handled in a more systematic way?

This question sets the stage for Part V, which explores an emerging, fifth family of equity metrics: responsibility-sensitive metrics (sometimes termed "equality-of-opportunity" metrics). Such metrics seek to demarcate between an individual's "circumstances" (for which she is not accountable) and her "effort," and to wash out differential effort in the measurement of inequality. This is an exciting, new development in the measurement of equity—and here, too, the Pigou-Dalton principle can be seen as foundational. Responsibility-sensitive metrics satisfy the principle after some kind of normalization for individual "effort." For example, one such metric, proposed by John Roemer—a pioneering figure in this field—favors Pigou-Dalton transfers from someone who has more of the appropriate "currency" to someone who has less, *if* they fall in the same "effort" class.

This Article has three aims: descriptive, explanatory, and normative. First, I hope to describe a variety of methodologies, each widely utilized, for quantifying equity. Many readers will not be familiar with all or even any of these methodologies. Second, I aim to shine a light on the Pigou-Dalton principle, showing how that fundamental axiom of equity both unites the methodologies, and yet also differentiates them insofar as the methodologies choose to specify the principle in different ways. Finally, I hope to excavate the normative debates underlying such differences in specification. What *is* the best way to measure equity? This is a pretty tough question to answer, but at least we can be clear about why it is so hard.

I. EQUITY METRICS: AN OVERVIEW

This Part describes four families of equity metrics, stressing how each family—in its own way—is connected to the Pigou-Dalton principle. These families, between them, comprise the overwhelming bulk of contemporary scholarly work attempting to provide some kind of quantitative assessment of equity.

Throughout the Article, I will speak of a "Pigou-Dalton transfer" occurring with respect to some particular "currency"—income, utility,

happiness, etc.—and an equity metric satisfying the “Pigou-Dalton principle” with respect to some currency. A *Pigou-Dalton transfer*, with respect to a given currency, means a non-leaky transfer of the currency from someone who has more to someone who has less, still leaving the first person with a greater than or equal amount of the currency and changing no one else’s holdings of it. A given metric satisfies the *Pigou-Dalton principle* with respect to a given currency, if it necessarily counts a Pigou-Dalton transfer in that currency as reducing the degree of inequity.

A. Inequality Metrics

Traditional measures of inequality are *unidimensional*. They focus on the distribution of some single attribute in a population. The most popular unidimensional inequality measures are the Gini coefficient; variance-based measures, such as the coefficient of variation; the Theil index; and the Atkinson index.²⁹

For the most part, unidimensional inequality metrics have been used by scholars to study the distribution of income (or some other economic “currency,” such as wages, wealth, or consumption). But nothing in the formal structure of such metrics requires that income, or another economic attribute, be the input. All that is required is (1) information sufficient to estimate the distribution of the attribute in the population of interest (and the distribution’s intertemporal path, if time trends in inequality are being quantified), and (2) some cardinal scale for measuring the attribute.³⁰

Thus, a burgeoning body of work looks at the inequality of longevity or some cardinal measure of health.³¹ And *if* the researcher is comfortable

29. For overviews of the theory of unidimensional inequality metrics, see HILDE BOJER, DISTRIBUTIONAL JUSTICE: THEORY AND MEASUREMENT 63–133 (2003); SATYA R. CHAKRAVARTY, INEQUALITY, POLARIZATION, AND POVERTY ch. 1 (2009); FRANK A. COWELL, MEASURING INEQUALITY (3d ed., 2011); PETER J. LAMBERT, THE DISTRIBUTION AND REDISTRIBUTION OF INCOME 1–132 (3d ed., 2001); Charles Blackorby, Walter Bossert & David Donaldson, *Income Inequality Measurement: The Normative Approach*, in HANDBOOK ON INCOME INEQUALITY MEASUREMENT 133 (Jacques Silber ed., 1999); Frank Cowell, *Measurement of Inequality*, in 1 HANDBOOK OF INCOME DISTRIBUTION 87 (Anthony Atkinson & Francois Bourguignon eds., 2000); Frank Cowell, *Inequality and Poverty Measures*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); Bhaskar Dutta, *Inequality, Poverty and Welfare*, in HANDBOOK OF SOCIAL CHOICE AND WELFARE 597 (Kenneth J. Arrow et al. eds., 2002).

30. By “cardinal,” here, I mean that the attribute is measurable on a scale which is *finer* than an ordinal scale—a scale that captures *more* than the levels of the attributes, e.g., differences between the levels, or ratios. Standard income inequality metrics assume either that income is measurable up to ratio rescalings, or up to a common translation factor. See CHAKRAVARTY, *supra* note 29, at 15–22; Blackorby, Bossert & Donaldson, *supra* note 29, at 144–51.

Some recent scholarship attempts to devise inequality metrics for attributes measured on an ordinal scale. See Dutta & Foster, *supra* note 15; Frank A. Cowell & Emmanuel Flachaire, *Inequality with Ordinal Data* (London Sch. of Econ., 2014), available at <http://darp.lse.ac.uk/pdf/IneqOrdinal.pdf>.

31. See, e.g., ANGUS S. DEATON & CHRISTINA H. PAXSON, *Aging and Inequality in Income and Health*, 88 AM. ECON. REV. (Papers & Proc.) 248 (1998); Edwards & Tuljapurkar, *supra* note 14; Neal

converting answers to a happiness survey to a cardinal scale, then the “currency” for the Gini coefficient, etc., can be happiness.³²

The so-called “Lorenz curve” is the conceptual foundation for unidimensional inequality metrics. (See Figure 1.) Consider, to begin, two populations with the same population size and same average individual level of the attribute (for example, the same average income), but different distributions. The Lorenz curve of a given distribution arranges individuals in rank order, from those who have the least amount of the attribute, to those who have the most. For each group of individuals along this spectrum (the lowest-ranked individual, the lowest- and second-lowest ranked individuals, the three lowest-ranked individuals, etc.), the group’s percentage of the total *number* of individuals in the entire population is mapped onto their share of the total *amount* of the attribute in the population. For example, if the first 5% of the population earns 1% of total income, and the first 10% earns 3% of total income, the Lorenz curve would include the points (5%, 1%) and (10%, 3%).

The axiom of “Lorenz dominance” says that if the Lorenz curve for one distribution is sometimes inside the Lorenz curve for a second, and never outside, then an inequality metric should assign the first distribution a lower inequality number (degree of inequality).³³ Lorenz dominance is generally seen as *the* most fundamental axiom that an inequality metric should satisfy, and indeed the standard metrics *do* satisfy it.³⁴ The axiom of Lorenz dominance is very intuitive: if one distribution Lorenz-dominates a second, then—in a very clear and intuitive sense—it is “closer” to perfect equality than the second. Where standard inequality metrics differ is in how

Fann et. al., *Maximizing Health Benefits and Minimizing Inequality: Incorporating Local-Scale Data in the Design and Evaluation of Air Quality Policies*, 31 RISK ANAL. 908 (2011); Mark McGillivray et. al., *Health Inequality and Deprivation*, 18 HEALTH ECON. (Supplement) S1 (2009); Peltzman, *supra* note 12; Pope, *supra* note 23; David E. Sahn & Stephen D. Younger, *Measuring Intra-Household Health Inequality: Explorations Using the Body Mass Index*, 18 HEALTH ECON. (Supplement) S13 (2009).

32. See sources cited *supra* note 21.


33. Let $(\pi, L(\pi))$ be the Lorenz curve for one distribution, where π is the fraction of the total population in a given group of rank-ordered individuals, and $L(\pi)$ their fraction of the total population holdings of the attribute. Let $(\pi, L^*(\pi))$ be the Lorenz curve for a second distribution. Then this second curve Lorenz-dominates the first if $L^*(\pi) \geq L(\pi)$ for all π , with this inequality strict for some.

34. Note that a perfectly equal distribution of the attribute will Lorenz-dominate every other distribution.

Recall that we are focusing here on two distributions with the same population size and the same mean. In that case *all* standard inequality metrics satisfy the axiom of Lorenz dominance. In a more general case, we might draw Lorenz curves for two distributions with the same population sizes but different means. In this case, so-called *relative* inequality metrics (those that assign the same degree of inequality to a distribution and all ratio rescalings) will continue to satisfy Lorenz dominance, but other kinds of inequality metrics need not. The most widely-used inequality metrics—the coefficient of variation, Gini, Atkinson, and Theil—are relative metrics. See CHAKRAVARTY, *supra* note 29, ch. 1 (discussing relative and absolute metrics); COWELL, *supra* note 29, at 61–74 (discussing relative metrics under heading of “income scale independence”).

they rank two distributions with “crossing” Lorenz curves, such that neither Lorenz-dominates the second.

Figure 1: The Lorenz Curve



Explanation: This figure illustrates a Lorenz curve for four individuals. The first individual has 5% of the total population holdings of the attribute, the second 15%, the third 30%, and the last 50%. The point A corresponds to the group comprised of the single lowest-ranked individual; the point B to the group comprised of the two lowest-ranked individuals; the point C to the group comprised of the three lowest-ranked individuals; and the point D to the entire population. The value on the x axis is the percentage of the population in the group, and the value on the y axis (also indicated in parentheses) is the percentage of the total population holdings of the attribute held by the group.

The nexus between inequality metrics and the Pigou-Dalton principle, in turn, grows from the Lorenz curve. In a famous 1934 book, Hardy, Littlewood, and Polya proved the following (nonobvious) mathematical fact: one distribution Lorenz-dominates a second if and only if the first can

be reached from the second by a series of Pigou-Dalton transfers (plus perhaps a permutation).³⁵

For example, imagine that there are four individuals in the population, and that they begin with the unequal distribution (70, 95, 110, 125), where 70 is the amount of the attribute held by Individual 1 (e.g., \$70,000 in annual income if the attribute is income), 95 the amount of the attribute held by Individual 2, etc. The sum total of the attribute is 400. This distribution corresponds to a Lorenz curve whereby the first 25% of the population has 70/400 of the total attribute, the first 50% has 165/400 of the total, and the first 75% has 275/400 of the total.

That distribution is (of course) Lorenz-dominated by a perfectly equal distribution (100, 100, 100, 100). The starting point (70, 95, 110, 125) can be turned into perfect equality via the following series of Pigou-Dalton transfers: from the starting point to (70, 100, 105, 125), to (75, 100, 100, 125), to (100, 100, 100, 100).³⁶

The distribution is also Lorenz-dominated by the unequal distribution (75, 92, 113, 120)—a distribution “closer” to perfect equality. Why? Note that in this second distribution the first 25% of the population has 75/400 of the total (a greater percentage than 70/400); the first 50% of the population has 167/400 (which exceeds 165/400); and the first 75% has 280/400 (which exceeds 275/400). Moreover, the Lorenz-dominating distribution (75, 92, 113, 120) can be reached from the starting point (70, 95, 110, 125) via the following series of Pigou-Dalton transfers: from the starting point to (75, 90, 110, 125), to (75, 90, 115, 120), to (75, 92, 113, 120). The “magic” of the Pigou-Dalton transfer, as proved by Hardy, Littlewood and Polya, is that this will be true for *every* case in which one distribution (perfectly equal or not) Lorenz-dominates a second.

I have focused thus far on the inequality ranking of two distributions with the same population size and mean. Unidimensional inequality metrics

35. See GODFREY H. HARDY, JOHN E. LITTLEWOOD & GEORGE POLYA, *INEQUALITIES* (1934); ALBERT W. MARSHALL & INGRAM OLKIN, *INEQUALITIES: THEORY OF MAJORIZATION AND ITS APPLICATIONS* 21–22 (Richard Bellman, ed.) (1979) (reproducing the Hardy/Littlewood/Polya theorem). Although the theorem is formulated in terms of Pigou-Dalton transfers in the broader sense of gap-diminishing transfers, *see supra* note 28, the proof strategy also shows that: one distribution is Lorenz-dominated by a second if and only if there is a series of Pigou-Dalton transfers in the narrower sense discussed in this Article (non-rank-switching transfers), combined perhaps with a permutation of individuals' holdings, that leads from the first distribution to the second. To see why a permutation may be needed, consider the distribution among two individuals (30, 70), with 30 the holdings of the first individual, and the distribution (60, 40). The second distribution Lorenz-dominates the first, and can be generated from the first by a Pigou-Dalton transfer of 10, yielding (40, 60); and then a permutation to yield (60, 40). Clearly a Pigou-Dalton transfer in the narrower sense *alone* will never get us from (30, 70) to (60, 40).

36. The second distribution is reached from the starting point via a Pigou-Dalton transfer of 5 from the third individual to the second; the third distribution from the second via a Pigou-Dalton transfer of 5 from the third individual to the first; and the perfectly equal distribution from the third distribution via a Pigou-Dalton transfer of 25 from the fourth individual to the first.

can be (and often are) employed to compare distributions with different means, population sizes, or both. In such a case, the Pigou-Dalton principle does not directly constrain inequality measurement.³⁷ However, it shapes such assessment indirectly, if quantification is undertaken using a metric (Gini, coefficient of variation, Theil, etc.) that is sure to respect the Pigou-Dalton principle in every case where the principle does apply.

Multidimensional inequality metrics have recently become popular.³⁸ This development coincides with theoretical and empirical work on “capabilities” and “functionings.”³⁹ A key theme in the capabilities/functionings literature is the multiplicity of determinants of individual well-being. An individual’s welfare cannot be reduced to her income, her health, her social and professional life, or the quality of her leisure, etc., but depends upon all of these, and more.

One response to this (undeniably true) proposition about well-being is to construct some *inclusive measure of individual well-being* that integrates information about each individual’s multidimensional attainments, and then to use this all-things-considered number as the “currency” for a unidimensional inequality metric or social welfare function.⁴⁰

But multidimensional inequality metrics do not take this path. Instead, two other routes are pursued:

- a) Dimension-by-dimension assessment. The simplest approach is to apply a unidimensional inequality metric to each of the several dimensions about which the researcher has distributional data; and then take an average (or some similar function⁴¹) of these dimensional numbers. For example, if there is data about the distribution of income, health, and educational attainment (each

37. If there has been a pure transfer, without any other change, then necessarily neither the population nor the mean level of the attribute has changed.

38. See CHAKRAVARTY, *supra* note 29, ch.5 ; Asis Kumar Banerjee, *A Multidimensional Gini Index*, 60 MATHEMATICAL SOC. SCI. 87 (2010); Satya R. Chakravarty & Maria Ana Lugo, *Multidimensional Indicators*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); Jean-Yves Duclos et al., *Partial Multidimensional Inequality Orderings*, 95 J. PUB. ECON. 225 (2011); Esfandiar Maasoumi, *Multidimensioned Approaches to Welfare Analysis*, in HANDBOOK ON INCOME INEQUALITY MEASUREMENT 437 (Jacques Silber ed., 1999); Ernesto Savaglio, *Three Approaches to the Analysis of Multidimensional Inequality*, in INEQUALITY AND ECONOMIC INTEGRATION 269 (Francesco Farina & Ernesto Savaglio eds., 2006); John Weymark, *The Normative Approach to the Measurement of Multidimensional Inequality*, in INEQUALITY AND ECONOMIC INTEGRATION, *supra*, at 303.

39. See, e.g., SABINA ALKIRE, VALUING FREEDOMS (2002); WIEBKE KUKLYS, AMARTYA SEN’S CAPABILITY APPROACH (2005); THE CAPABILITY APPROACH: CONCEPTS, MEASURES AND APPLICATIONS (Flavio Comim et al. eds, 2008); Ingrid Robeyns, *The Capability Approach: A Theoretical Survey*, 6 J. HUM. DEVELOPMENT 93 (2005). The field is inspired by Martha Nussbaum’s and Amartya Sen’s scholarship. See, e.g., MARTHA NUSSBAUM, WOMEN AND HUMAN DEVELOPMENT: THE CAPABILITIES APPROACH (2000); AMARTYA SEN, INEQUALITY REEXAMINED (1992).

40. See ADLER, *supra* note 28, at 119–24.

41. By “similar,” here, I mean that the overall inequality assigned to a multidimensional distribution is monotonically increasing in inequality in each of the dimensions.

cardinally measured), we might estimate multidimensional inequality as the average of the Gini coefficient of income, health, and educational attainment.

- b) “Sophisticated” multidimensional inequality metrics. The dimension-by-dimension format has various deficits. In particular, it ignores interdimensional correlations. (Imagine that two populations have the very same distributions of income, health, and educational attainment—taken separately—but in the first case those with higher incomes tend to have better health and education, while in the second case there are no systematic connections between an individual’s health, income, and education. Then, surely, the second distribution is more equitable.) “Sophisticated”⁴² multidimensional metrics are (mathematically complicated) formulas for assigning an inequality number to a multidimensional distribution—formulas which are not merely an average or some other function of the inequality in each dimension taken separately.

How do multidimensional inequality tools link up with the Pigou-Dalton principle? For the dimension-by-dimension approach, the answer is straightforward. A Pigou-Dalton transfer in any dimension will decrease the degree of inequality in that dimension and thus will decrease the *average* of dimension-specific inequality.⁴³ For “sophisticated” multidimensional metrics, the answer is more complex. Such metrics typically satisfy the Pigou-Dalton principle with respect to *coordinated* transfers, which (depending on the metric) can mean either that (1) a *bundle* of Pigou-Dalton transfers between two individuals, in each and every dimension, reduces the degree of inequality assigned to a multidimensional distribution;⁴⁴ or that (2) a Pigou-Dalton transfer in at least one dimension, from an individual who is at a higher level in *all*

42. “Sophisticated” is not a standard term in the literature, but rather one that I use here to mean any approach to measuring multidimensional inequality *other than* dimension-by-dimension assessment. Such an approach is “sophisticated” in the sense that it can take account of interdimensional correlations.

43. It will also decrease the value of any other function monotonically increasing in dimension-specific inequality, *see supra* note 41.

44. More precisely, a standard multidimensional Pigou-Dalton principle says that a bundle of equiproportional dimension-specific Pigou-Dalton transfers between two individuals—equiproportional in the sense that the distance between the individuals on each dimension is reduced by the same fraction—decreases multidimensional inequality. *See* Weymark, *supra* note 38, at 307; ADLER, *supra* note 28, at 131 n.140. The literature on multidimensional metrics also discusses a “majorization” axiom, requiring inequality to decrease if a multidimensional distribution is multiplied by a bistochastic matrix. An inequality metric which satisfies this majorization axiom will also satisfy the equiproportional Pigou-Dalton principle just mentioned. *See* Weymark, *supra*, at 308.

dimensions, to an individual at a lower level in all dimensions, reduces the assigned degree of inequality.⁴⁵

B. Social Welfare Functions

Social welfare functions (SWFs) are much discussed in theoretical welfare economics; are widely used in some fields of applied economics, particularly optimal tax scholarship; and provide the intellectual foundation for cost-benefit analysis with distributional weights. An SWF takes information about the population distribution of individual attainments (be it attainment on a single dimension, or on multiple dimensions), and uses a “utility” function to convert that information into a list of “utility” numbers—one for each individual. Distributions are then ranked on the basis of these utility numbers.⁴⁶

SWFs can assume various functional forms. One such form, the most famous, is utilitarian. The utilitarian SWF compares two distributions by summing utilities. The utilitarian SWF *may* respect the Pigou-Dalton principle with respect to income. This depends on how the translation of income into utility occurs. If this translation is non-linear, and in particular is such that income has “diminishing marginal utility,” the utilitarian SWF will satisfy the Pigou-Dalton principle with respect to income.

But the utilitarian SWF clearly does not respect the Pigou-Dalton principle “in the space of utility.” Assume that there are four individuals. In scenario *x*, the individuals have various attainments, translated into the list of utilities (30, 140, 70, 80). In scenario *y*, the attainments are different, so that the utilities become (60, 110, 70, 80). Thus, in scenario *y*, there has been a Pigou-Dalton transfer in utility (30 units) from Individual 2 to Individual 1. But the sum totals of utility are the same, and hence utilitarianism ranks the scenarios as equally good.

However, there *are* important classes of SWFs that satisfy the Pigou-Dalton principle in the space of utility (for short, “equity-regarding” SWFs). “Continuous prioritarian” SWFs⁴⁷ sum utility numbers


45. See Casilda Lasso de la Vega et al., *Characterizing Multidimensional Inequality Measures Which Fulfill the Pigou-Dalton Bundle Principle*, 35 SOC. CHOICE & WELFARE 319 (2010). On the Pigou-Dalton principle in the multidimensional context, see also sources cited *supra* note 38; Henar Diez et al., *A Consistent Multidimensional Generalization of the Pigou-Dalton Transfer Principle: An Analysis*, 7 B.E. J. THEORETICAL ECON. art. 45 (2007); Kristof Bosmans et al., *Note: A Consistent Multidimensional Pigou-Dalton Transfer Principle*, 144 J. ECON. THEORY 1358 (2009).

46. By “utility,” I mean a measure of well-being that is both intra- and interpersonally comparable, and makes intrapersonal comparisons in conformity with individuals’ preferences. See *infra* Part III. For a comprehensive discussion of social welfare functions, see ADLER, *supra* note 28. On their functional forms, see *id.* ch. 5. For citations to scholarship on SWFs, see *id.* at 83 n.45, 87 n.56. For a recent survey, see John Weymark, *Social Welfare Functions*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015).

47. See ADLER, *supra* note 28, at 356–58.

“transformed” by a concave function, such as the square root. Figure 2 below shows why such SWFs satisfy the Pigou-Dalton principle in the space of utility.

Figure 2: A “Continuous Prioritarian” SWF



Explanation: The continuous prioritarian SWF is the sum of $g(u)$, where u is the utility of a particular individual and $g(\cdot)$ an increasing, concave function, as illustrated by Figure 2. Note that a Pigou-Dalton transfer of Δu units of utility from a person at a higher utility level (here, individual j with utility u_j), to an individual at a lower utility level (here, individual i , with utility u_i) will increase the total sum of $g(u)$. Because $g(\cdot)$ is concave, $g(u_i + \Delta u) - g(u_i)$ is greater than $g(u_j) - g(u_j - \Delta u)$, as the figure shows. This in turn means that $g(u_i + \Delta u) + g(u_j - \Delta u)$ is greater than $g(u_i) + g(u_j)$.

The Pigou-Dalton principle in the space of utility is also satisfied by other SWFs, such as the rank-weighted and leximin SWFs.⁴⁸ Continuous prioritarian, leximin, and rank-weighted SWFs are the prime examples of equity-regarding SWFs.⁴⁹ All three are structured (in somewhat different ways) to give greater weight to utility changes affecting individuals at lower utility levels, as compared to individuals at higher levels.

What is the difference between ranking distributions using an inequality metric with utility as the currency, and ranking such distributions using an equity-regarding SWF? A critical difference is that the first technique ignores the sum total of utility. (For example, if everyone's utility doubles, an inequality metric with utility as its currency will record no change,⁵⁰ while an SWF will see an improvement.) Inequality metrics are *only* measures of inequality. Equity-regarding SWFs are hybrid measures, taking into account both information about total well-being (the sum total of utility) and information about the degree of equality of utility.

However, the key point to stress here is the similarity between equity-regarding SWFs and inequality metrics: namely, that both respect the Pigou-Dalton principle in some fashion. By contrast, with the utilitarian SWF—which does not respect the Pigou-Dalton principle in the space of utility, and does so in the space of income *only if* the translation from income to utility is such that income has diminishing marginal utility—an equity-regarding SWF *always* respects a Pigou-Dalton transfer in the space of utility.⁵¹ Note that, in the case of a non-leaky transfer in utility from a higher utility to a lower utility individual, total well-being is a “wash” and the equality of utility becomes the only consideration driving the equity-regarding SWF. If, in addition, the transfer is a Pigou-Dalton transfer, shrinking the utility gap between the two individuals, the transfer reduces

48. The leximin SWF ranks utility vectors according to the worst-off utilities; if these are equal, the second-worst-off; and so forth. See ADLER, *supra* note 28, at 367–74. It is obvious that this SWF satisfies the Pigou-Dalton principle. The rank-weighted SWF orders utility numbers from lowest to highest, and sums these multiplied by fixed, decreasing weights. See *id.* at 351–56. The rank-weighted SWF can also be shown to satisfy the Pigou-Dalton principle. See *id.* at 353.

49. These are certainly not the *only* SWFs that respect the Pigou-Dalton principle. See, e.g., ADLER, *supra* note 28, at 374–77 (discussing prioritarian SWF with a lexical threshold). However, these are the three such SWFs that are by far the most widely used in scholarly work in the SWF tradition.

50. Strictly speaking, this is true for relative inequality metrics, while absolute inequality metrics will be invariant to increases in overall well-being that occur by increasing everyone's utility by the same amount. On the difference between such metrics, see CHAKRAVARTY, *supra* note 29, at 16. Either type of metric will assign the same value to utility vectors in which everyone's utility is equal, regardless of the total.

51. Note that, because SWFs take utility numbers (not attributes) as their inputs, an SWF can be structured to respect the Pigou-Dalton principle in the space of utility *regardless* of the mapping from individual attributes to utility levels—and this is true of the continuous prioritarian, rank-weighted, and leximin SWFs.

the inequality of utility and is preferred by the equity-regarding SWF, just as it would be by an inequality metric with utility as the currency.

C. Poverty Metrics

Traditionally, poverty measurement focused on *income* poverty (just as inequality metrics were first developed to measure income inequality). Formally, an income-poverty metric is a function of the distribution of individual incomes—with two critical differences from income-inequality measures.⁵² First, poverty measurement also depends on the identification of a *poverty line*. Various approaches are adopted here: the income-poverty line can be specified in an “absolute” fashion (as the amount of income required to meet minimum nutritional requirements or other basic needs); or a “relative” fashion (as some fraction of the mean or median income).⁵³

The poverty line permits the characterization of individuals as “poor” or not: a “poor” individual is someone whose income is below the line. But it also undergirds the so-called “focus” requirement: if Z is the poverty line, and distributions d and d' have the same pattern of incomes below Z , then the degree of poverty for the two distributions is the same.⁵⁴

Because of the focus requirement, the connection between poverty metrics and the Pigou-Dalton axiom is less straightforward than for inequality metrics. A logical consequence of the “focus” requirement is that above-threshold transfers have no effect on the degree of poverty. This is intuitive. Assume that distribution d is (10, 15, 30, 60, 100), while d' is (10, 15, 50, 50, 90). The numbers represent annual incomes with five individuals in total and the poverty line is, let us assume, 25. Note that d' is reached from d via a series of two Pigou-Dalton transfers.⁵⁵ Distribution d' therefore Lorenz-dominates d ; and the Gini coefficient, Theil index, coefficient of variation, or any other standard inequality metric will assign a lower degree of inequality to d' than d . But every standard *poverty* metric (by virtue of the focus requirement) will say that the degree of poverty in d and d' is exactly the same.

52. On income-poverty metrics, see CHAKRAVARTY, *supra* note 29, ch. 2; LAMBERT, *supra* note 29, ch. 6; Dutta, *supra* note 29; Buhong Zheng, *Aggregate Poverty Measures*, 11 J. ECON. SURV. 123 (1997).

53. On the setting of poverty lines, see Martin Ravallion, *Poverty Lines Across the World* (World Bank Dev. Res. Grp., Working Paper No. 5284 April 2010).

54. See CHAKRAVARTY, *supra* note 29, at 50; Zheng, *supra* note 52, at 130. Note that poverty metrics generally can compare two distributions, given the two, possibly different, poverty lines associated with each distribution—but that the focus axiom is expressed with reference to the case in which the same line is used for two distributions.

55. First transfer 10 from the fifth individual to the third, then transfer 10 from the fourth to the third.

Still, the “focus” requirement is consistent with *restricted* versions of the Pigou-Dalton principle, such as the following⁵⁶:

- (1) Below-Threshold Transfers: A Pigou-Dalton transfer between two poor individuals lowers the degree of poverty.
- (2) Across-Threshold Transfers (weak version): If one individual is poor, and a second non-poor,⁵⁷ a Pigou-Dalton transfer between them that does not change their positions relative to the poverty line (i.e., does not change the number of poor and non-poor individuals) lowers the degree of poverty.
- (3) Across-Threshold Transfers (strong version): If one individual is poor, and a second non-poor, a Pigou-Dalton transfer between them lowers the degree of poverty even if it changes the number of poor individuals.

In the theoretical literature on the measurement of income poverty, the first and second axioms are generally endorsed. The third is more controversial;⁵⁸ consider a case in which a Pigou-Dalton transfer increases the number of poor individuals (i.e., the mean income of transferor and transferee is below the poverty line, and the transfer is sufficiently large to bring the transferor below that line).

Whether the third axiom *should* be accepted depends on broader questions about the normative role of poverty measures, mooted below. The basic point, here, is that the “focused” structure of poverty measures precludes their satisfying the Pigou-Dalton principle with respect to transfers above the poverty line—but, reciprocally, *permits* their satisfying restricted versions of the Pigou-Dalton principle (such as (1), (2), or (3)), applicable to a transfer that increases the income of someone who starts out below the poverty line.

56. Statements of the Pigou-Dalton principle as a condition on poverty metrics, such as those immediately below, constrain the poverty measure of two distributions using a common poverty line. For a discussion of different versions of the Pigou-Dalton principle for poverty metrics, see sources cited *supra* note 52. See also Lucio Esposito & Peter J. Lambert, *Poverty Measurement: Prioritarianism, Sufficiency and the T's of Poverty*, 27 *ECON. & PHIL.* 109 (2011); Buhong Zheng, *Poverty Orderings*, 14 *J. ECON. SURVEYS* 427 (2000). Zheng, *supra* note 52, is especially systematic.

57. Strictly, a “non-poor” individual is someone whose income places her at or above the poverty line. See Zheng, *supra* note 52 (noting that a “weak” definition of the poor avoids certain axiomatic difficulties that arise if individuals at the poverty line are counted as “poor”). However, so as to simplify matters, my discussion of the non-poor is focused on the paradigmatic case of individuals whose holdings place them above the poverty line. My claims generalize to “non-poor” individuals who are right at the poverty line. Most fundamentally, note that regardless of whether the term “non-poor” is used to include or exclude those right at the poverty line, poverty metrics are invariably “restricted” in the sense of ignoring Pigou-Dalton transfers among individuals above the poverty line.

58. See CHAKRAVARTY, *supra* note 29, at 53; Zheng, *supra* note 52, at 133.

One of the most popular classes of poverty metrics in academic work is the FGT class.⁵⁹ The FGT approach calculates each poor individual's fractional shortfall from the poverty line (with others assigned a value of zero) and then averages these numbers raised to some power α . If this α parameter is greater than 1, the FGT approach satisfies all three of the above Pigou-Dalton axioms.

Admittedly, some poverty metrics employed in empirical research (or by governmental statistical bureaus) go to the other extreme and fail to satisfy any one of these axioms. This is true, in particular, of the headcount ratio: the percentage of the population that is poor.⁶⁰ But the headcount ratio is often criticized in theoretical work on poverty measurement (beginning with a seminal critique by Amartya Sen, in a 1976 article⁶¹) because it wholly ignores the *distribution* of income among poor individuals and thus is insensitive to transfers (be they from poor to poor or from rich to poor) that change this distribution but do not change the numbers of poor and non-poor.

The recent rise of the “multidimensional” approach to inequality measurement is mirrored by similar work on poverty metrics—in both cases, fueled by a desire to move “beyond income” as an indicator of human well-being.⁶² Assume, now, that d is a *multidimensional* distribution

59. See Udo Ebert & Patrick Moyes, *A Simple Axiomatization of the Foster, Greer & Thorbecke Poverty Orderings*, 4 J. PUB. ECON. THEORY 455 (2002); James Foster, Joel Greer, & Erik Thorbecke, *The Foster-Greer-Thorbecke (FGT) Poverty Measures: 25 Years Later*, 8 J. ECON. INEQUALITY 491 (2010). Let Z be the income poverty line, and y_i the income of individual i or Z , whichever is smaller.

With N individuals, the FGT index is:
$$\frac{1}{N} \sum_{i=1}^N \left(\frac{Z - y_i}{Z} \right)^\alpha.$$

60. See Zheng, *supra* note 52, at 142–44.

61. Amartya Sen, *Poverty: An Ordinal Approach to Measurement*, 2 ECONOMETRICA 219 (1976). See also AMARTYA SEN, INEQUALITY REEXAMINED 102–04 (1992). Sen notes in the latter work that “[t]he need for having distribution-sensitivity in measuring poverty [i.e., a sensitivity of the metric to the distribution of income among the poor] seems to be fairly widely accepted by now.” *Id.* at 106.

62. See Sabina Alkire, *Capabilities*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); CHAKRAVARTY, *supra* note 29, ch. 6; Satya Chakravarty and Maria Ana Lugo, *Multidimensional Indicators*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); THE MANY DIMENSIONS OF POVERTY (Nanak Kakwani & Jacques Silber eds., 2007); QUANTITATIVE APPROACHES TO MULTIDIMENSIONAL POVERTY MEASUREMENT (Nanak Kakwani & Jacques Silber eds., 2008); Sabina Alkire & James Foster, *Counting and Multidimensional Poverty Measurement*, 96 J. PUB. ECON. 476 (2011) [hereinafter Alkire & Foster, *Counting and Multidimensional Poverty Measurement*]; Sabina Alkire & James Foster, *Understandings and Misunderstandings of Multidimensional Poverty Measurement*, 9 J. ECON. INEQUALITY 289 (2011) [hereinafter Alkire & Foster, *Understandings and Misunderstandings*]; Sabina Alkire & Maria Emma Santos, *A Multidimensional Approach: Poverty Measurement & Beyond*, 112 SOCIAL INDICATORS RES. 239 (2013); A.B. Atkinson, *Multidimensional Deprivation: Contrasting Social Welfare and Counting Approaches*, 1 J. ECON. INEQUALITY 51 (2003); Francois Bourguignon & Satya R. Chakravarty, *The Measurement of Multidimensional Poverty*, 1 J. ECON. INEQUALITY 25 (2003); Jean-Yves Duclos et al., *Robust Multidimensional Poverty Comparisons*, 116 ECON. J. 943 (2006); Martin Ravallion, *On*

(describing individual attainments with respect to a multiplicity of welfare-relevant attributes). A multidimensional poverty metric, like its inequality counterpart, will assign some number to this distribution—with the critical difference that this number depends not just on the pattern of individual attainments but also on the identification of dimension-specific thresholds. For example, if d describes individuals' levels of the three attributes income, health, and education, then the researcher will need to specify a separate threshold (poverty line) for income, for health, and for education—e.g., a minimum level of income, health, and education, respectively, required for a decent human life.

What is the structure of multidimensional poverty metrics? Here, as with multidimensional inequality measurement, one can differentiate between (1) dimension-by-dimension and (2) “sophisticated” approaches.⁶³ The dimension-by-dimension approach will employ some metric M traditionally employed to calculate income poverty (for example, the headcount ratio, the FGT metric, or any other); will apply it to each of the dimensions; and will then take an average (or some similar function)⁶⁴ of these dimension-specific poverty values. If M would satisfy a restricted version of the Pigou-Dalton principle, when applied to income, then the dimension-by-dimension approach using M will satisfy the very same principle on an attribute-by-attribute basis.

For example, assume that M , applied to an income distribution, satisfies the axiom of “below-threshold transfers”: a Pigou-Dalton transfer from one person with a below-threshold income to a second with an even lower income reduces the degree of poverty according to M . Now imagine that multidimensional distribution d' is reached from d by a transfer on any dimension from someone below the threshold for that dimension to someone even lower. Then the dimension-by-dimension approach using M will say that the degree of poverty is lower in d' than d .

Research on “sophisticated” multidimensional poverty metrics is quite new, and it is hard to be sure about what will emerge. A leading contender seems to be the “Alkire-Foster” index of multidimensional poverty, which builds upon the FGT formula.⁶⁵ Roughly speaking, Alkire and Foster propose to identify an individual as “poor” if she has below-threshold

Multidimensional Indices of Poverty, 9 J. ECON. INEQUALITY 235 (2011); Kai-yuen Tsui, *Multidimensional Poverty Indices*, 19 SOC. CHOICE & WELFARE 69 (2002).

63. See Alkire & Foster, *Understandings and Misunderstandings*, *supra* note 62, at 303–04 (criticizing “marginal,” i.e., dimension-by-dimension approach because it “does not look across dimensions for the same person and cannot reflect the extent of associations among deprivations”).

64. See *supra* note 41.

65. See Alkire & Foster, *Counting and Multidimensional Poverty Measurement*, *supra* note 62; Alkire & Foster, *Understandings and Misunderstandings*, *supra* note 62; Alkire & Santos, *supra* note 62.

attainments in a sufficient number of dimensions.⁶⁶ Poor individuals are assigned a deprivation value in each dimension, equaling their fractional shortfall from the dimension-specific line. Non-poor individuals are assigned a deprivation value of 0 in each dimension. These values, raised to the power α , are summed across individuals and dimensions.

The Alkire-Foster formula differs from a dimension-by-dimension approach, because its identification of an individual as “poor” depends upon her attainments in all the dimensions, and this identification, in turn, helps determine the deprivation numbers assigned to each individual. Still, the approach satisfies a version of the Pigou-Dalton principle: a transfer on any dimension from one poor individual to a second reduces the degree of poverty.⁶⁷

D. Social-Gradient Metrics

The social-gradient view of equity is adopted in much research on public health and in the literature on environmental justice.⁶⁸ On this view, the population-wide distribution of income, health, utility, happiness, or any other attribute lacks normative relevance as such. Rather, what is relevant is the *association* between valuable or harmful attributes and markers of social status. As a leading researcher explains: “[E]quity in health can be defined as the absence of systematic disparities in health . . . between social groups who have different levels of underlying social advantage/disadvantage—that is, different positions in a social hierarchy.”⁶⁹ Social status is proxied in a number of different ways. “Analytical traditions vary [on how to choose the proxy]: in the United Kingdom, social groups have been defined using five categories of occupation-based social class; in some countries in continental Europe, educational attainment or occupational categories have been used; and in

66. Their approach to characterizing a person as “poor” also allows for differential weighting of dimensions.

67. This assumes (as with the FGT metric for the unidimensional case) that $\alpha > 1$. Note that a Pigou-Dalton transfer in one dimension to a person who is below the dimension-specific threshold, but is not “poor” (*i.e.*, is not below the threshold on a sufficient number of dimensions), will *not* reduce the degree of poverty as per the Alkire-Foster metric. This is true regardless of where the individual ends up relative to the dimension-specific threshold, and of where the transferor starts out or ends up.

By contrast, the dimension-by-dimension application of a unidimensional poverty metric that satisfies some version of the Pigou-Dalton principle *will* record such a transfer as reducing poverty, at least in some cases (depending on the positions of the transferee and transferor relative to the dimension-specific threshold.). In particular, the dimension-by-dimension application of the FGT metric will *always* record such a transfer as reducing poverty.

68. See Matthew D. Adler, *Risk Equity: A New Proposal*, 32 HARV. ENVTL. L. REV. 1, 6–11 (2008) (citing literature).

69. Paula Braveman & Sofia Gruskin, *Defining Equity in Health*, 57 J. EPIDEMIOLOGY & CMTY. HEALTH 254, 254 (2003).

the USA, most research focuses on social categories defined in terms of racial groups.”⁷⁰

A variety of tools are employed to measure the association between social class and harms or benefits.⁷¹ Many such tools satisfy the Pigou-Dalton principle in some form. In outlining this connection, I will focus on the “concentration index”—increasingly favored by researchers as one of the best social-gradient metrics because it integrates information about attainments in all social groups.⁷²

Imagine that the population is divided into G social groups, with Group 1 at the bottom of the social hierarchy, Group 2 in the second-worst social position, etc. There is a distribution, d , which distributes some valuable attribute to the members of each of the G groups.⁷³ The so-called “concentration curve” for this distribution is drawn as follows. For each group, in succession, we plot the percentage of the population belonging to that group and all lower status groups against the percentage of the total amount of the attribute held by those groups. For example, if there are three groups of low, middle, and high status, with respectively 20%, 70%, and 10% of the population; and the total amount of the attribute held by each group is, respectively, 5, 45, and 50; then the concentration curve includes the points (0,0), (20%, 5%), (90%, 50%), and (100%, 100%).

The “concentration index” for a given distribution d , then, is twice the area between its concentration curve and the 45-degree line.⁷⁴ The higher the value for this index, the more the attribute is “concentrated” in higher-status groups. See Figure 3.


70. C.J.L. Murray et al., *Health Inequalities and Social Group Differences: What Should We Measure?*, 77 BULL. WORLD HEALTH ORG. 537, 538 (1999).

71. For descriptions of many of these, see Paula Braveman, *Health Disparities and Health Equity: Concepts and Measurement*, 27 ANN. REV. PUB. HEALTH 167, 172–79 (2006); Johann Mackenbach & Anton E. Kunst, *Measuring the Magnitude of Socio-economic Inequalities in Health*, 44 SOC. SCI. & MED. 757 (1997); Adam Wagstaff et al., *On the Measurement of Inequalities in Health*, 33 SOC. SCI. & MED. 545 (1991).

72. See Nicolas Ziebarth, *Measurement of Health, Health Inequality, and Reporting Heterogeneity*, 71 SOC. SCI. MED. 116, 117 (2010) (noting that the concentration index is “by far the dominant health inequality indicator”). For descriptions of the index and its properties, see Philip Clarke & Tom Van Ourti, *Calculating the Concentration Index when Income is Grouped*, 29 J. HEALTH ECON. 151 (2010); Guido Erreygers, *Correcting the Concentration Index*, 28 J. HEALTH ECON. 504 (2009); Nanak Kakwani et al., *Socioeconomic Inequalities in Health: Measurement, Computation, and Statistical Inference*, 77 J. ECONOMETRICS 87 (1997).

73. The concentration index is sometimes applied to *bad* attributes, e.g., ill health, but the analysis in this Article generally assumes that the currency for equity is a good attribute—be it income, utility, or the multiple goods that go into multidimensional inequality or poverty metrics. Thus, I focus here on the concentration index used to measure social skews in good attributes.

74. It can range in value from 1 to -1 , with lower value indicating *less* concentration of the attribute in higher status groups, i.e., greater equity.

Figure 3: The Concentration Index

Explanation: This figure illustrates the concentration curve for three groups, G1, G2, and G3, ranked from lower to higher social status. The low-status group G1 has 25% of the population and 5% of the total population holdings of the attribute; the middle-status group G2 has 55% of the population and 45% of the attribute; the high-status group G3 has 20% of the population and 50% of the attribute. The points corresponding to each group (those labeled G1, G2, and G3) show, on the x -axis, the percentage of the population in that group and all lower status groups; and, on the y -axis, the percentage of the attribute in that group and all lower status groups. The concentration curve connects these points. The concentration index is twice the area between that curve and the diagonal, i.e., twice the area labeled “CI” in this figure.

The concentration curve has similarities to the Lorenz curve, with the critical difference that individuals are plotted along the x -axis in light of their social position, not their rank in the population distribution of the attribute. For this reason, the concentration index (unlike an inequality metric) satisfies only a social-class-restricted version of the Pigou-Dalton principle.

Pigou Dalton Principle: Social Class Restricted

Imagine that one individual has more of some valuable attribute than a second in distribution d , and that a Pigou-Dalton transfer of the attribute from the first to the second individual occurs, yielding d' . If the transferee has a worse position in the social hierarchy, then d' should be assigned a lower degree of inequity than d .⁷⁵

To see why the concentration index *satisfies* the social-class-restricted version of the Pigou-Dalton principle, observe that a Pigou-Dalton transfer of the attribute from a higher to a lower status individual will tend to “flatten” the concentration curve and bring it close to the diagonal. Conversely, if d' is reached from d via a Pigou-Dalton attribute transfer from someone with more of the attribute to someone who has less, but the transferor has the same or *lower* social status than the transferee, the concentration index will *fail* to see d' as more equitable than d . In other words, the concentration index fails to satisfy the more general version of the Pigou-Dalton principle, requiring any transfer to increase equity regardless of the social classes of the two individuals. Clearly, if the transferee and transferor are in the same group, the value of the concentration index does not change at all. And if we are transferring from a high attribute, low-status individual to a low attribute, high-status individual (for example, transferring income from a high-income, minority person to a low-income Caucasian in a society where race is the marker of social status), we will be tending to increase the concentration of the attribute in the higher status group.

E. A Summary

Several themes emerge from this survey of four widely-used families of equity metrics: inequality metrics, social welfare functions (SWFs),

75. Let there be G groups with increasing social status, $g = 1$ to G , and N individuals, $i = 1$ to N , each falling in one of the groups. Let y_i be the attribute level of individual i , and μ the population mean. Let f_m be the fraction of the population in group m . Define R_g , the so-called “relative rank”

corresponding to group g , as follows: $R_g = \frac{f_g}{2} + \sum_{m=1}^{g-1} f_m$. For each individual i , let her relative rank R_i be the rank of the group she belongs to (lower status individuals will have lower such values). It can

be shown that the concentration index is equal to: $\left[\frac{2}{N\mu} \sum_{i=1}^N y_i R_i \right] - 1$.

This formula makes it apparent that a Pigou-Dalton transfer between two individuals will lower the value of the concentration index iff the transferee has a lower relative rank than the transferor, i.e., iff the transferee is lower social status.

poverty metrics, and social-gradient metrics. First, the “currency” for the Pigou-Dalton principle can vary. Income-inequality metrics and income-poverty metrics, paradigmatically, satisfy the Pigou-Dalton principle with respect to *income* (in the case of income-poverty metrics, a restricted version of the principle). Equity-regarding SWFs satisfy that principle with respect to *utility*. Unidimensional inequality metrics applied to some non-income attribute (e.g., health or happiness) will satisfy the principle with respect to that attribute. Multi-attribute inequality or poverty metrics will satisfy the principle (perhaps in a coordinated and/or restricted form) with respect to transfers in each of the multiple dimensions of individual attainment being measured.

Second, the Pigou-Dalton principle may be restricted to a subset of transfers. Poverty metrics will not satisfy the principle with respect to transfers between non-poor individuals (those above the poverty line). The concentration index, a paradigmatic social-gradient metric, does not satisfy the principle except for transfers from higher to lower social status individuals.

Table 1: Understanding Equity Metrics⁷⁶

	What is the Currency for the Pigou-Dalton Principle?	Is the Principle Restricted in Scope?
Inequality Metrics <i>Unidimensional</i> <i>Multidimensional</i> <u>Simple</u> <u>Sophisticated</u>	Income or any other single attribute (health, happiness, utility, etc.) Each of a multiplicity of attributes (a transfer along any dimension improves equality) A transfer along one dimension, if appropriately coordinated with the other dimensions, improves equality	<u>No</u> : A Pigou-Dalton transfer is approved regardless of whether the transferee is non-poor, of high social class, or responsible for being worse off than the transferor
Equity-Regarding SWFs (most notably the continuous prioritarian, leximin and rank-weighted SWFs)	Utility (a measure of well-being that is both interpersonally comparable and respects individual preferences)	<u>No</u>
Poverty Metrics <i>Unidimensional</i> <i>Multidimensional</i> <u>Simple</u> <u>Sophisticated</u>	Income or any other single attribute (health, happiness, utility, etc.) Each of a multiplicity of attributes (a transfer along any dimension reduces poverty) A transfer along one dimension, if appropriately coordinated with the other dimensions, reduces poverty	<u>Yes</u> : Transfers to individuals who are above the poverty line (or “non-poor” in a more complicated sense for sophisticated multidimensional metrics) do not reduce poverty

76. Note that the table entry for “sophisticated” multidimensional inequality metrics requires appropriate “coordination” among the transfers. As discussed earlier, the versions of the Pigou-Dalton principle discussed in the literature on multidimensional inequality build in this “coordination” requirement in various ways. One such requirement requires equiproportional transfers on all dimensions; another, that the transferor be better off in all dimensions. *See supra* text accompanying notes 44–45.

Similarly, the table entry for “sophisticated” multidimensional poverty metrics has a coordination requirement. Here, recall that the Alkire-Foster metric (currently the leading multidimensional poverty metric) incorporates such a requirement because it *doesn't* record a dimension-specific transfer as reducing poverty if the transferee is above a sufficient number of thresholds in the *other* dimensions. *See* sources cited *supra* note 65.

<p>Social Gradient Metrics</p> <p>(The analysis focuses on the concentration index, the best-developed social-gradient metric.)</p>	<p>Income or any other single attribute (health, happiness, utility, etc.)</p>	<p><u>Yes</u>: Transfers improve equity only if the transferee is in a lower-status social group than the transferor</p>
--	--	--

These divergent choices with respect to the specification of the Pigou-Dalton principle help to explain the diversity of these metrics—to explain why their mathematical expression varies, and why they can disagree in their comparisons of situations as more or less equitable. To be sure, divergence in the specification of the Pigou-Dalton principle is not the *only* important difference between these metrics. (For example, the Gini coefficient and the coefficient of variation, applied to compare income distributions, will each satisfy the unrestricted Pigou-Dalton principle with income as its currency—and yet these metrics *are* still distinct, namely in how they rank two income distributions with crossing Lorenz curves.) Still, I think we make substantial progress in understanding the “topography” of equity metrics by seeing that the four popular families I have described all embrace the Pigou-Dalton principle *in some form*, with disagreement about what that form should be.⁷⁷

II. WHY THE PIGOU-DALTON PRINCIPLE? A GENERIC JUSTIFICATION

The Pigou-Dalton principle, evidently, is pervasive in the measurement of equity. But can this pervasive role be justified? Why, as a normative matter, should we endorse it?

Here I offer a general argument for the Pigou-Dalton principle—an argument encapsulated in the form of a “Generic Justification” for the principle, stating conditions under which the principle seems to have

77. Two other choices with respect to the specification of the Pigou-Dalton principle, orthogonal to those just summarized, should be mentioned. First, under conditions of uncertainty regarding individual attributes, equity metrics may or may not satisfy an “ex ante” Pigou-Dalton principle. Second, in a multiperiod framework, the Pigou-Dalton principle may be applied to individuals’ lifetime attribute holdings, or instead to their period-by-period holdings. See ADLER, *supra* note 28, chs. 5, 6 (discussing these choices in context of equity-regarding SWFs). How equity metrics make these choices, too, is an important aspect of their structure—one with real consequence for empirical work. See, e.g., Matthew D. Adler, *Well-Being, Inequality and Time: The Time-Slice Problem and Its Policy Implications* (Inst. for Law and Econ. Research Paper 07–17, 2007), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1006871 (documenting empirical differences between lifetime and sublifetime approaches to equity); *supra* note 46 (describing debate about whether inequality of consumption—a rough proxy for lifetime income—has increased to the same degree, in the U.S., as inequality of lifetime income).

However, for the sake of simplicity, I ignore time and uncertainty in this Article and focus on exploring the structure of equity metrics as regards (1) “currency,” (2) whether the Pigou-Dalton principle is “restricted” to adjust for poverty or social class, and (3) individual responsibility, see Part V.

considerable normative force.⁷⁸ The Generic Justification, if accepted, implies that *some* Pigou-Dalton transfers, in *some* currency, are normatively attractive.

The Pigou-Dalton Principle: A Generic Justification

Assume that (a) one individual (the “transferee”) has a lower level of some currency C than a second individual (“the transferor”); (b) the transferee is worse off than the transferor; (c) the transferee is not responsible for being worse off than the transferor; (d) a Pigou-Dalton transfer in the currency from the transferor to transferee will produce a well-being improvement for the transferee which is *greater than or equal to* the well-being loss for the transferor; (e) after the transfer, the transferee will not be better off than the transferor; (f) there are no externalities (no one else’s well-being is affected). Then the transfer is normatively justified.

Assume that Toni, a transferee, starts off with amount C_{Toni} of some currency, and Gregor, a transferor, starts off with amount C_{Gregor} , which is greater than C_{Toni} . Society engineers a Pigou-Dalton transfer from Gregor to Toni. (Toni increases her holdings of C by the same amount, ΔC , that Gregor loses; and this amount is less than or equal to half the distance between their initial holdings of C , so that Toni ends up at $C_{\text{Toni}} + \Delta C$, which is less than or equal to $C_{\text{Gregor}} - \Delta C$.) Add now the critical facts about *well-being* and *responsibility*, built into the Generic Justification. C_{Toni} is not merely a smaller amount of C than C_{Gregor} . In addition, Toni gains at least as much from the transfer, in well-being terms, as Gregor loses. Moreover, before the transfer, Toni is at a lower well-being level than Gregor and, even afterwards, is at a lower (or at least not a higher) level. In such a case, doesn’t Toni have a stronger normative claim in favor of the transfer, than Gregor’s normative claim in opposition?

Plausibly, the strength of an individual’s normative claim for or against some social policy is a function, first, of the *impact* of the policy upon her (the magnitude of her well-being difference pre- and post-policy); and, second, her well-being *level* (whether she is better or worse off than others in her society).⁷⁹ *Ceteris paribus*, individuals more substantially affected by a given policy have stronger claims regarding it; and, *ceteris paribus*, worse-off individuals who stand to be affected by a given policy have stronger claims regarding it. But in the case at hand *each* of these two

78. Cf. ADLER, *supra* note 28, at 307–51 (defending Pigou-Dalton principle for case where C is well-being).

79. See *id.* at 339–40; THOMAS NAGEL, EQUALITY AND PARTIALITY (1995); THOMAS NAGEL, *Equality*, in MORTAL QUESTIONS 106 (1979).

factors weighs in favor of Toni. Were Toni to benefit from the transfer by less than Gregor loses, the first factor would point to Gregor; and were she to start and/or finish better off than him, the second might. But (by hypothesis) neither is true here.

To be sure, if Toni were (in some sense) at fault for being at a lower pre-transfer level of C than Gregor, that might deflate her normative claim to the transfer. And if the transfer harmed other individuals, then their claims in opposition—together with Gregor's—might defeat Toni's claim in favor. But the Generic Justification does not apply in such cases; it purports to justify a Pigou-Dalton transfer in some currency only absent externalities, and absent transferee responsibility for being comparatively worse off.

The Generic Justification presumes that well-being *levels* and *differences* are interpersonally comparable, at least to some extent. This is not to say that well-being levels and differences are *universally* comparable, merely (and much less controversially) that there are *some* cases in which one individual is genuinely better off than or equally well off as a second, and *some* cases in which an individual is more affected by or equally affected by a choice as compared to a second—as opposed to these levels or changes being non-comparable.⁸⁰

One objection to the Generic Justification has to do with Pigou-Dalton transfers, in some currency C , that increase the gap between the transferee and individuals with an even lower level of C . Assume that there are three individuals in the population, Able, Baker and Charlie. Able is at level 10, with much less C than Baker, at 80, and Charlie, at 100. Ten units of C are transferred from Charlie to Baker. So we have moved from the distribution (10, 80, 100) to the distribution (10, 90, 90). This is indeed a Pigou-Dalton transfer in C —but it leaves Able “stranded” at the bottom of the distribution, even further below the next best off, Baker, than Able was before the transfer. Isn't it intuitively plausible that the distribution (10, 80, 100) is actually more equitable than the distribution (10, 90, 90)?⁸¹

A person's position in the population distribution of some attribute might *itself* affect her well-being. For example, someone's welfare might be (in part) a function of her income *relative* to those around her—perhaps because low relative income tends to cause unhappiness. By virtue of such “positional” effects, the distribution (10, 90, 90) might be worse *for Able* than the distribution (10, 80, 100).⁸² But, if so, the Generic Justification

80. See ADLER, *supra* note 28, at 185–92.

81. See, e.g., Ernesto Savaglio, *Inequality as Differences: A Simple Characterization*, 61 RES. IN ECON. 31, 32 (2007); Buhong Zheng, *Utility-Gap Dominances and Inequality Orderings*, 28 SOC. CHOICE & WELFARE 255, 256 (2007).

82. The revisionary inequality metrics mentioned immediately below, which fail to respect the Pigou-Dalton principle, see *infra* note 83, are regularly justified with reference to the *feelings* of

does not apply. It claims that a Pigou-Dalton transfer in *C* is justified *on the condition* that there are no externalities. Positional effects, or other kinds of effects on the well-being of third parties, certainly could override the normative case for a transfer.

Thus, in considering the plausibility of the Generic Justification, we should imagine that Able is equally well off with (10, 80, 100) and (10, 90, 90). If so, isn't (10, 90, 90) a normative improvement over (10, 80, 100)? Able is unaffected by the choice. Since he is unaffected, why should it matter that he is at 10 rather than, say, 150? Presumably the distribution (150, 90, 90) is better than the distribution (150, 80, 100)—on the condition that the shrinking of the gap between Baker and Charlie has no external impact on the well-being of Able. But, if so, why reach a different verdict about the comparison of (10, 90, 90) and (10, 80, 100), if that condition still holds true?

The critic of the Generic Justification might respond that (10, 80, 100) is a more *equal* distribution than (10, 90, 90), while (150, 90, 90) is a more equal distribution than (150, 80, 100). This change in the degree of inequality makes (10, 90, 90) worse than (10, 80, 100) even if Able is genuinely unaffected by the transfer from Charlie to Baker.

But what *is* the criterion for measuring equality which says that (10, 90, 90) is less equal than (10, 80, 100)? The first distribution Lorenz-dominates the second, and thus every standard inequality metric will say that *it* is more equal.

To be sure, the critic might try to construct a non-standard inequality metric which says that (10, 80, 100) is more equal than (10, 90, 90) while (150, 90, 90) is more equal than (150, 80, 100).⁸³ But any such metric will be contestable. The Lorenz criterion has real intuitive force (which of course explains why the criterion is the linchpin for traditional inequality metrics). If one distribution of some attribute Lorenz-dominates a second, then—in a straightforward and intuitive way—it is closer to perfect equality.

In short, a strong case can be presented for accepting the Generic Justification. Like most normative arguments, this case is not irresistible. But it shows why *some* version of the Pigou-Dalton principle is very plausible.

resentment and deprivation that worse-off persons have in virtue of their relative position. See CHAKRAVARTY, *supra* note 29, at 85; Brice Magdalou & Patrick Moyes, *Deprivation, Welfare, and Inequality*, 32 SOC. CHOICE & WELFARE 253, 254 (2009).

83. For discussion of inequality metrics that fail to respect the standard Pigou-Dalton principle, see CHAKRAVARTY, *supra* note 29, ch. 3; Udo Ebert, *Taking Empirical Studies Seriously: The Principle of Concentration and the Measurement of Welfare and Inequality*, 32 SOC. CHOICE & WELFARE 555 (2009); Magdalou & Moyes, *supra* note 82; Savaglio, *supra* note 81.

The Generic Justification, in turn, is a tool that we can use to evaluate different equity metrics. We can now ask: do the transfers approved by that metric approximate the conditions identified by the Generic Justification? First, does the metric employ an appropriate currency (given that Justification)? Second, insofar as it restricts the transfers (as do poverty and social-gradient metrics), is *that* appropriate? It is to these questions that we now turn.

III. WHAT IS THE BEST CURRENCY FOR THE PIGOU-DALTON PRINCIPLE?

What is the best currency for the Pigou-Dalton principle? This Part grapples with the question and shows why it is a contestable one—at least at present, absent a consensus about how to construct an interpersonally comparable⁸⁴ measure of well-being.

Inequality and poverty metrics traditionally focus on the distribution of income, and, paradigmatically, satisfy the Pigou-Dalton principle (at least a restricted version thereof) in the space of income.⁸⁵ While income has attractive features as a currency for the Pigou-Dalton, it also possesses major deficits. The Generic Justification provides an analytic lens that brings the deficits of an income currency into clear view. But it remains disputable how to specify an alternative currency that will redress all of these imperfections.

The advantage of an income currency is that income is both readily measurable and widely useful in improving well-being (given the existence of markets)—so that income and well-being levels *correlate*. What are its disadvantages? First, because of positional effects, a Pigou-Dalton transfer in income may lower the well-being of third parties and thus, on balance, be inequitable. Second, quite apart from positional effects, a Pigou-Dalton transfer in income may be inequitable if the transferee has a higher level than the transferor of some non-economic good (health, environmental quality, a happy disposition, social life, etc.)—sufficiently so that the transfer is *regressive* in terms of well-being, making a better-off (if lower income) transferee yet better off, and a worse-off (if higher income) transferor yet worse off. The point here is that income and well-being levels are only *imperfectly* correlated. Third, a pure transfer in income may be a “leaky” transfer in well-being: the transferee may gain less, in well-

84. A plausible measure of well-being may have some “gaps” with respect to interpersonal comparisons of well-being levels and differences, but it should not wholly eschew such comparisons, or make them so rarely as to contradict our intuitions about hypothetical or real-world cases in which interpersonal comparisons seem to be possible. *See supra* text accompanying note 80.

85. I say “paradigmatically” because there are some income poverty metrics (albeit disapproved in the theoretical literature) that fail to satisfy the Pigou-Dalton principle even restricted to transfers between poor individuals. *See supra* text accompanying notes 60–61.

being terms, than the transferor loses. This is possible (although empirically unlikely) even if both individuals are identical in non-income attributes. It is possible (and empirically more plausible) where their non-income attributes differ.⁸⁶

The three observations in the previous paragraph all point to different ways in which a Pigou-Dalton transfer *in income* might fail to meet the conditions identified by the Generic Justification. A different observation is this: income-poverty and income-equality metrics fail to satisfy the Pigou-Dalton principle with respect to non-income currencies. A Pigou-Dalton transfer in health, for example, is not *registered* by such metrics (unless it happens to change incomes). But a Pigou-Dalton transfer in health, without any impact on income, obviously could be a net positive for equity.

The deficits of using income (or some other economic resource) as the “currency” for the Pigou-Dalton principle are, clearly, *not* resolved by using another, single, currency—if that attribute is simply one *aspect* of well-being, rather than an inclusive measure of well-being. For example, measuring equity by applying a unidimensional inequality or poverty metric to the distribution of longevity is vulnerable to all the criticisms just leveled against income-based inequality and poverty metrics.

Multidimensional inequality and poverty metrics, recall, have *multiple* currencies for the Pigou-Dalton principle. Such measures quantify equity as a function of individual attainment on a multiplicity of well-being relevant dimensions. Interest in them has been fueled by the (correctly) perceived deficits of the traditional, income-based approach to measuring poverty and inequality.⁸⁷ Critiques of income as the *distribuendum* for equity are standard fare in the “multidimensional” literature.⁸⁸

However, the multidimensional approach is hardly a panacea. Consider, first, dimension-by-dimension metrics. This methodology applies an inequality or poverty metric to each dimension and then calculates multidimensional inequality or poverty as the average (or some similar function) of the dimensional values.⁸⁹ Thus, a Pigou-Dalton transfer in *any one* of the dimensions will reduce multidimensional inequality (and multidimensional poverty, depending on where the dimension-specific poverty lines are set).

86. A further objection to applying the Pigou-Dalton principle to income (or any other single attribute which is only one aspect of well-being) is that this can conflict with the Pareto principle. See ADLER, *supra* note 28, at 114–17; *see also infra* text accompanying notes 93–94 (discussing conflicts between multidimensional Pigou-Dalton principle and Pareto principle).

87. See A.B. Atkinson & F. Bourguignon, *The Comparison of Multi-Dimensioned Distributions of Economic Status*, 49 REV. ECON. STUD. 183, 183–85 (1982).

88. See, e.g., *id.*; Bourguignon & Chakravarty, *supra* note 62, at 25–27; Erik Thorbecke, *Issues Related to the Concept of Multidimensional Poverty*, in THE MANY DIMENSIONS OF POVERTY, *supra* note 62, at 3–5; Tsui, *supra* note 62, at 71–72.

89. See *supra* text accompanying note 41.

But a Pigou-Dalton transfer in any single dimension, taken alone, may be *regressive* in terms of well-being. To illustrate, assume that there are five dimensions taken account of by the metric, and a Pigou-Dalton transfer in the first dimension occurs. Before the transfer, the transferee has bundle (a, b, c, d, e) , while the transferor has bundle $(a^*, b^*, c^*, d^*, e^*)$. The transferor's holdings of the first attribute, a^* , are greater than the transferee's, a . After the transfer, the bundles become $(a + \Delta a, b, c, d, e)$ and $(a^* - \Delta a, b^*, c^*, d^*, e^*)$. Since $b, c, d,$ and e may well be greater than, respectively, $b^*, c^*, d^*,$ and e^* , the transfer on the first dimension may have the perverse effect of making a better-off person yet better off.

The proponent of the dimension-by-dimension approach might resist this attack by denying the comparability of well-being levels and differences or the relevance of interpersonal well-being comparisons to equity—in effect, rejecting the Generic Justification and offering some alternative, non-welfarist “story” in favor of the Pigou-Dalton principle.⁹⁰ Even so, the dimension-by-dimension approach is hardly impeccable. A Pigou-Dalton transfer on any one dimension might, on balance, be unfair to the transferor if the transferee has a higher level on other dimension(s).⁹¹ Even the non-welfarist should be able to see the force of this objection to the approach.

“Sophisticated” multidimensional inequality or poverty metrics do not merely average dimension-specific measures. They may therefore have the feature of preferring a transfer in one dimension only if appropriately coordinated with transfers in other dimensions. For example, a “sophisticated” metric might only see a transfer as increasing equity if the transferor is better off than the transferee in all dimensions.⁹² This stipulation would allow the “sophisticated” approach to avoid the objections to the dimension-by-dimension approach offered in the previous two paragraphs.

However, other objections remain. Perhaps the most serious one is that multidimensional inequality and poverty metrics—sophisticated as well as dimension-by-dimension—can come into conflict with the Pareto principle.⁹³ Let d and d' be two multidimensional distributions. The Pareto

90. Although a strong case can be made that well-being (adjusted for responsibility) is the appropriate currency for assessments of fair distribution, this remains controversial. See generally Richard Arneson, *Welfare Should be the Currency for Justice*, 30 CAN. J. PHIL. 497 (2000).

91. See Daniel M. Hausman, *What's Wrong with Health Inequalities?*, 15 J. POL. PHILOSOPHY 46 (2006).

92. See Lasso de la Vega, *supra* note 45.

93. On conflicts between the Pigou-Dalton and Pareto principles in the multidimensional context, see Marc Fleurbaey, *Social Welfare, Priority to the Worst-Off and the Dimensions of Individual Well-Being*, in *INEQUALITY AND ECONOMIC INTEGRATION*, *supra* note 38, at 225, 238–342; Marc Fleurbaey & Alain Trannoy, *The Impossibility of a Paretian Egalitarian*, 21 SOC. CHOICE & WELFARE 243 (2003).

principle says: (1) If everyone is indifferent between d and d' , the two distributions are equally good; and (2) if everyone either prefers d' to d or is indifferent, and at least one individual prefers d' to d , then d' is better. One version of a “coordinated” multidimensional Pigou-Dalton principle says that if d' is reached from d via a Pigou-Dalton transfer on any single dimension, and the transferor starts off with higher levels on all dimensions, then d' is better than d .⁹⁴

It turns out that these two principles are logically inconsistent. It is impossible to come up with an integrated procedure for ranking distributions that satisfies both principles—as shown by the example in Table 2 immediately below.

Table 2: Multidimensional Pigou-Dalton versus Pareto

	Distribution d		Distribution d^*		Distribution d^{**}		Distribution d^{***}	
	A	B	A	B	A	B	A	B
Individual 1	15	1	16	2	3	17	4	18
Individual 2	18	4	17	3	2	16	1	15

Explanation: The A and B columns show the individuals’ levels of each attribute in distributions d , d^* , d^{**} , and d^{***} . Note that, in distribution d , Individual 2 has more of each attribute than Individual 1, and d^* is reached from d by a Pigou-Dalton transfer in both attributes. Similarly, Individual 1 has more of each attribute in distribution d^{***} , and distribution d^{**} is reached by a Pigou-Dalton transfer in both attributes.

But imagine that the individuals’ preferences are such that: first, they are each indifferent between their bundles in d and their bundles in d^{**} ; and, second, they are each indifferent between their bundles in d^* and their bundles in d^{***} .

With these preferences, it is impossible to rank the distributions in a manner that respects both the Pareto principle and the Pigou-Dalton

A different objection to the multidimensional approach is that one or more dimension-specific transfers may yield a leaky rather than pure transfer in well-being (thus departing from the conditions identified by the Generic Justification), where one individual gains less in well-being than a second loses. Note that this is true even for a sophisticated approach which requires the transferor to be better off in all dimensions.

94. Not only is this principle satisfied by some “sophisticated” multidimensional inequality metrics; note that it is also satisfied by the dimension-by-dimension approach, which prefers a Pigou-Dalton transfer in any dimension regardless of the comparative position of transferee and transferor on other dimensions.

What about the other standard multidimensional Pigou-Dalton principle—preferring an equiproportional transfer in all dimensions? ADLER, *supra* note 28, at 117–18, illustrates how this principle can conflict with the Pareto principle.

principle (either the dimension-by-dimension version, or the version limited to transfers where the transferor has more of each attribute). The Pareto principle requires that d is equally good as d^{**} , and that d^* is equally good as d^{***} . The Pigou-Dalton principle requires that d^* is better than d and that d^{**} is better than d^{***} . These can't both be true.

The Pareto principle has strong normative appeal. It lies at the core of welfare economics (as readers surely know), just as the Pigou-Dalton principle is foundational to equity measurement. It would be nice to show that these two grand normative projects are *compatible*: that we can endorse Pigou-Dalton transfers (in an appropriate currency), without inevitably running afoul of the Pareto principle. Unfortunately, the multidimensional “currency” for Pigou-Dalton does not enable such happy reconciliation.

What about using an *inclusive measure of individual well-being* as the currency for Pigou-Dalton? For any bundle of attributes someone might possess (including both economic resources, such as income, and any non-economic attributes the equity researcher might wish to include), an inclusive measure of individual well-being assigns that bundle a single number quantifying its *well-being* value, integrating information about all the attributes. Equity is then measured as a function of the distribution of these well-being values.

As discussed above, equity-regarding social welfare functions (SWFs) follow this strategy. In particular, they employ *utility* as an inclusive well-being indicator. The SWF approach assumes that “utility” has the following features: First, utility covaries with preferences, meaning that any given individual has greater utility in one outcome as compared to a second iff⁹⁵ she prefers the first outcome to the second. Second, utility mirrors interpersonal comparisons of well-being levels and differences. One person's utility is greater than a second's iff the first person is better off than the second (and similarly for well-being differences).⁹⁶

Assume that a utility indicator with these features can be constructed. Then we can readily achieve the desired integration of efficiency and equity—of the Pareto principle and the Pigou-Dalton principle. Given any set of possible distributions—distributions of bundles of both economic and non-economic attributes among a population—our utility measure will translate each distribution into a list of individual utilities. An equity-

95. “Iff” means “if and only if.”

96. See ADLER, *supra* note 28, at 61–88. On a broader view, “utility” is a measure of well-being (even a non-preference-based account of well-being). More traditionally, economists (including those in the SWF tradition) use “utility” as a measure of preferences—and because I believe a strong case can be made for analyzing well-being in terms of preferences (perhaps appropriately “laundered” preferences), I stick to that usage here. See *id.* ch. 3.

regarding SWF will rank the lists and corresponding distributions in a logically coherent (transitive) fashion. Moreover, if two distributions are such that some individuals have more utility in the first, and no one has less utility in the first, any equity-regarding SWF will rank the first distribution higher.⁹⁷ This assures consistency with the Pareto principle (on the assumption that we have a measure of utility that covaries with preferences). Finally, if one distribution is reached from a second by a Pigou-Dalton transfer in utility, an equity-regarding SWF will rank the first distribution higher. Thus (on the assumption that utility mirrors interpersonal comparisons of well-being levels and differences) a Pigou-Dalton transfer in *utility* will always meet the conditions set forth by the Generic Justification.⁹⁸

However, producing a measure of utility with the desired features is not straightforward. The difficulty arises because of the heterogeneity of preferences. Consider, first, the case in which everyone has the same preferences, including risk preferences: the same ranking of attribute bundles and the same ranking of lotteries over bundles. Then there will be a function, u , the mathematical expectation of which represents these common preferences. Clearly, such a u does secure consistency with the Pareto principle. Moreover, if B is a bundle held by Jim, and B^* held by Sue, $u(B) > u(B^*)$ means that both individuals (and everyone else) prefer the first bundle to the second—so it seems very plausible that Jim with B does indeed have a higher level of well-being than Sue with B^* . Finally, u in the case of common preferences is at least a plausible interpersonally comparable measure of well-being differences.

But where preferences are heterogeneous, it is less clear how to arrive at utility.⁹⁹ Where Jim, Sue, Randy, etc., rank bundles differently, there is a function u_{Jim} that represents Jim's preferences; a different function, u_{Sue} , that represents Sue's; yet a different function, u_{Randy} , that represents Randy's; and so forth. If we pick Jim's function as our measure of utility (say), and use that as the input to our SWF, the Pareto principle may be violated, since u_{Jim} does not track Sue's preferences or Randy's—and, in any event, it seems arbitrarily "dictatorial" to give priority to one person's preferences in assigning a well-being value to attribute bundles for the purposes of social decision making.

Thus, in the case of heterogeneous preferences, in order to assign a "utility" number to attribute bundles—where these numbers will, in turn, serve as the input for an equity-regarding SWF or some other equity

97. Throughout the Article, I focus on standard, "Paretian" SWF that have the monotonicity feature just described. See ADLER, *supra* note 28, at 70-71, 307-08.

98. See *id.* at 119-24, 339-51.

99. See *id.* at 199-200, 279-84.

metric—we will somehow need to take account of the *plurality* of individual functions representing each individual’s preferences over the bundles. Although there *are* intelligible proposals for how to do this, these proposals have arguable flaws or lacunae.¹⁰⁰ The SWF literature often tends to overlook the problem of measuring utility in the case of heterogeneous preferences, and certainly has not arrived at a plausible, consensus solution to that problem.¹⁰¹

The concept of “equivalent” income is one proposal for handling heterogeneous preferences in the construction of an inclusive measure of individual well-being—a proposal recently much discussed by economic theorists.¹⁰² The idea here is to *normalize* someone’s income so that it reflects her level of non-income attributes. More precisely, some “reference bundle” of non-income attributes is specified. The equivalent income of a given individual is such that she is indifferent between that amount and the reference bundle of non-income attributes, as compared to her actual income and actual non-income attributes.¹⁰³ “Equivalent income” can then function as the input to an inequality or poverty metric, SWF, or social-gradient metric.

“Equivalent income” has substantial attractions as a currency for the Pigou-Dalton principle. Unlike actual income, it reflects non-income attributes. And there is no logical inconsistency between requiring the ranking of distributions to satisfy both the Pareto principle and the Pigou-Dalton principle in terms of equivalent income.¹⁰⁴ But this proposed currency is hardly perfect. The Pigou-Dalton principle framed in terms of

100. In my own work, I have developed the concept of “extended preferences” as a basis for interpersonal comparisons. A given “spectator” or “deliberator” develops a well-being ranking of bundles of attributes (I also use the term “history”), and in so doing may choose to respect the preferences of the individuals (“subjects”) who possess the attributes. I use the term “extended preferences” to refer to such a ranking. Different spectators may have different extended preferences, and one way to accommodate such diversity is to pool them. *See id.* at 201–36; Matthew Adler, *Extended Preferences*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); Matthew Adler, *Extended Preferences and Interpersonal Comparisons: A New Account*, 30 ECON. AND PHIL. 123 (2014).

An objection to the pooling of extended preferences is that this may yield too much incomparability with respect to well-being levels and differences. *See* Alex Voorhoeve, *Book Review*, 42 SOC. CHOICE & WELFARE 245 (2014). Even leaving aside pooling, a lacuna in my extended-preferences framework is the absence of a specific theory for scaling the von Neumann-Morgenstern functions representing the preferences of subjects with different tastes. The theory allows for various such scalings, but gives no definitive recommendations about which to choose. *See* Adler, *Extended Preferences*. Further refinement of the framework will, I hope, close this gap.

101. *See* ROBIN BODWAY, FROM OPTIMAL TAX THEORY TO TAX POLICY 199–217 (2012).

102. *See generally* MARC FLEURBAEY & DIDIER BLANCHET, BEYOND GDP: MEASURING WELFARE AND ASSESSING SUSTAINABILITY ch. 4 (2013); Marc Fleurbaey, *Equivalent Income*, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY (Matthew D. Adler and Marc Fleurbaey eds., forthcoming 2015).

103. *See* Fleurbaey, *supra* note 102.

104. *See* ADLER, *supra* note 28, at 124–30.

equivalent income may *fail* to prefer a policy that meets the Generic Justification, and vice versa.¹⁰⁵

“Happiness” is yet another possible currency—as in recent empirical work that quantifies the inequality of happiness, based on responses to happiness surveys.¹⁰⁶ On one view, someone’s happiness is simply *one* dimension of her well-being. More ambitiously, a response to an appropriately constructed happiness survey might be seen as an inclusive measure: someone who expresses a higher happiness number is, all things considered, better off.¹⁰⁷

The viability of happiness as an inclusive measure is a complex question, which I have elsewhere discussed at length.¹⁰⁸ In brief, the proponent of this approach faces a dilemma. (1) Someone’s answer to a happiness question *might* be seen as an indicator of the quality of her experiences: whether she feels pain, avoids pleasure, has a sense of satisfaction about what she does, etc. But such an indicator can only be a reflection of someone’s all-things-considered well-being *if* well-being is mentalistic: if someone’s experiences are the only intrinsic source of well-being. And *that* is a highly controversial view about well-being. (2) Alternatively, someone’s answer to a happiness survey might be taken as an indicator of how well her preferences are satisfied. In short, it indicates her “utility.” But, even in the case where individuals have identical preferences (so that the construction of a utility metric is relatively straightforward), it need not be true that a higher happiness score signals higher utility, or that differences in happiness scores mirror differences in

105. To see this in a simple way, imagine that in d all individuals actually have the reference bundle of non-income attributes. If well-being is an increasing, concave function of income with non-income attributes fixed at the reference bundle, then there will be cases where d' is produced from d via a *leaky* transfer of income from a higher- to a lower-income individual, thus is not a Pigou-Dalton transfer in equivalent income, but *is* a Pigou-Dalton transfer in well-being. Conversely, if well-being is an increasing, convex function of income with non-income attributes fixed at the reference bundle, then there will be cases where d' is produced from d via a Pigou-Dalton transfer in equivalent income but is *not* a Pigou-Dalton transfer in well-being.

The proponent of equivalent-income might respond that well-being is *linear* in income at the reference bundle of non-income attributes. However, it is hard to see why this would be true (for any choice of reference bundle). It is a truism about well-being that income has non-constant (e.g., diminishing) marginal utility.

A second and more subtle difficulty with using equivalent income as the currency for the Pigou-Dalton principle is that it fails to differentiate between individuals with the same ordinal preferences but different risk preferences. See Adler, *Extended Preferences*, *supra* note 100.

106. See Clark et al., *supra* note 15; Stevenson & Wolfers, *supra* note 15, at S33; Ruut Veenhoven & Wim Kalmijn, *Inequality-Adjusted Happiness in Nations*, 6 J. HAPPINESS STUD. 421 (2005). Happiness has also been used as the currency for poverty measurement. See G.G. Kingdon & J. Knight, *Subjective Well-Being Poverty vs. Income Poverty and Capabilities Poverty*, 42 J. DEV. STUD. 1199 (2006).

107. See Matthew D. Adler, *Happiness Surveys and Public Policy: What’s the Use?* 62 DUKE L.J. 1509 (2013)

108. See *id.*

utility.¹⁰⁹ Nor is it clear that using happiness scores is an attractive solution to the much harder problem of constructing a utility metric when individuals have heterogeneous preferences.

In summary, from a welfarist perspective, the ideal currency for the Pigou-Dalton principle would be an inclusive measure of well-being. More specifically—for anyone who endorses a preference-based view of well-being and, therewith, the Pareto principle—the ideal currency would be an interpersonally comparable measure of individual utility. But it remains controversial how to construct such utilities once preference heterogeneity enters the picture. Reciprocally, currencies *other* than an inclusive well-being measure have clear imperfections. At least some such imperfections, summarized in Table 3 below, characterize not only traditional income and poverty metrics but also metrics that employ another single currency as well as multidimensional metrics, a newer tool. However, the continuing presence of equity metrics with these sorts of currencies is understandable and, to some extent, justified as a “second-best” matter, absent a clearly viable inclusive well-being measure.

109. The problem is that different individuals may use different utility scales to express the common preferences, or give a response to a happiness question that expresses their feelings rather than “utility” in the sense of preference-satisfaction. *See id.*

Table 3: The Pros and Cons of Different “Currencies” for the Pigou-Dalton Principle

	Transferee might be better off (even though at lower level of currency)?	Transfer might be “leaky” in terms of well-being (well-being increase of transferee less than loss of transferor)?	Potential conflict with Pareto principle?
<i>Income</i>	Yes	Yes	Yes
<i>Any other single dimension of individual well-being (e.g., health)</i>	Yes	Yes	Yes
<i>Multiple Dimensions</i>			
<u>Simple approach</u> (a transfer along any dimension improves equity)	Yes	Yes	Yes
<u>“Sophisticated” approach</u> (only coordinated transfers improve equity)	No* (if coordination requirement appropriately specified)	Yes	Yes
<i>Equivalent Income</i>	No*	Yes* (if unlikely)	No
<i>Utility</i>	No	No	No

Explanation: The “No” entries marked by an asterisk assume that well-being is increasing in each attribute (both income and the other attributes). It is possible but unlikely that a Pigou-Dalton transfer in equivalent income will be leaky in well-being, for reasons discussed in note 105 above.

IV. SHOULD THE PIGOU-DALTON PRINCIPLE BE “RESTRICTED”? HEREIN OF POVERTY AND SOCIAL-GRADIENT METRICS

Traditional inequality metrics and equity-regarding SWFs illustrate the Pigou-Dalton principle in its simplest, unrestricted form. The Gini coefficient, applied to the distribution of income, is such that a Pigou-Dalton transfer in income from *anyone* at a higher level of income, to *anyone* at a lower level, reduces the degree of inequality. Similarly, equity-regarding SWFs prefer a Pigou-Dalton transfer in utility from *anyone* at a higher level of utility, to *anyone* at a lower level.

A “restricted” version of the Pigou-Dalton principle, in some currency, prefers only a subset of transfers in that currency. In this Part, I normatively evaluate the two types of restrictions embodied by equity metrics that are currently in wide use: a restriction defined by the transferee’s and transferor’s location relative to the poverty line (embodied by poverty metrics), and a restriction defined by their social class (embodied by social-gradient metrics).

A. Poverty Metrics

As we saw in Part I, there are different ways in which the poverty line might function to restrict the scope of the Pigou-Dalton principle. For example, an income poverty metric might prefer every Pigou-Dalton transfer in income from a poor person (someone who lies below the income-poverty line) to someone even poorer, but not necessarily every Pigou-Dalton transfer in income from a non-poor person to a poor person. By contrast, the widely favored “FGT” income-poverty metric is fairly expansive in its sensitivity to transfers. If the FGT parameter α is above 1, *any* Pigou-Dalton transfer in which the transferee is poor reduces the measured degree of poverty.¹¹⁰

However, *every* income-poverty metric (including the FGT metric and every other) is insensitive to transfers between non-poor individuals. Income-poverty metrics simply do not take account of the income distribution above the poverty line; this is a universal, constitutive feature of such metrics. Similarly, *every* multidimensional poverty metric is such that a transfer (in one or more dimensions) between two individuals who are decisively non-poor—who have above-threshold attainments in all dimensions—will have no effect on the measured degree of poverty.

In short, all poverty metrics are insensitive (at least) to transfers between the non-poor. Why would this restriction be justified?

One argument for the restriction is based upon a normative view, “sufficientism,” defended by the contemporary moral philosopher Roger Crisp.¹¹¹ “Sufficientists” contend that a Pigou-Dalton transfer *in well-being* is a matter of normative indifference—not a normative improvement—if the transferee and transferor are both sufficiently well-off. Crisp supports this contention by describing a hypothetical case—the “Beverly Hills” case—in which society has to choose between giving fine wine to rich individuals or to super-rich individuals, where the welfare benefit of drinking the wine is the same for both types. He argues that it makes no difference who gets the wine.

110. See *supra* text accompanying notes 56–59.

111. Roger Crisp, *Equality, Priority, and Compassion*, 113 *ETHICS* 745 (2003).

It seems somewhat absurd to think that the *Rich* should be given priority over the *Super-rich*. . . . [W]hat the Beverly Hills case brings out is that, once recipients are at a certain level, any prioritarian concern for them disappears entirely. This implies that any version of the priority view must fail: when people reach a certain level, even if they are worse off than others, benefiting them does not, in itself, matter more. And this seems to me to be true even if, in a Beverly Hills case, the utilities are equal. That is, even if the benefits to each of the *Rich* and the *Super-rich* are identical and their numbers are the same, there still seems to me nothing to be said for giving priority to the “worse off.” At this level, only [the sum total of] utilities matter, so there would be nothing to choose between the two distributions.¹¹²

The Beverly Hills case seems to provide intuitive support for sufficientism. However, such intuitions may weaken if we think carefully about the case and keep in mind that (as Crisp constructs the case) having fine wine is supposed to *improve* the well-being of the rich and super-rich, rather than being a frivolous gesture that make no welfare difference at all.¹¹³

Assume that Richie is indeed very well-off but that Super-Richie is yet better off. It's not *merely* that Super-Richie's life is piled up with more material luxuries than Richie's. (After all, someone could have more material luxuries without being better off.) Rather, Super-Richie's life is really better for well-being than Richie's. Moreover, Super-Richie (and, *a fortiori*, Richie) is not yet at the maximum of human well-being.

As it turns out, Society has some indivisible item which *would* improve either Richie's life or Super-Richie's—indeed, by the same amount—and which no one else cares about. Society could just discard the item, but anyone who thinks Society should do *that* will run headlong into the Pareto principle. Now, if both Richie and Super-Richie were badly off, the choice between them would *not* be a matter of indifference (as Crisp would concede). Richie would be the appropriate recipient. Why would *that* be the case? Plausibly, for the reason I sketched earlier, arguing for the Generic Justification: Richie's *claim* to the item (understood as a function both of Richie's well-being level, and of the difference the item would make to

112. *Id.* at 755. Although Crisp focuses specifically here on prioritarianism, his critique is really a broader critique of any scheme for ranking outcomes (formalized by some type of “equity-regarding” SWF) that satisfies the Pigou-Dalton principle in the currency of well-being without regard to whether the transferee is above some well-being threshold.

113. For critiques of sufficientism, see Adler, *supra* note 28, at 345–51; Paula Casal, *Why Sufficiency Is Not Enough*, 117 *ETHICS* 296 (2007); Larry Temkin, *Egalitarianism Defended*, 113 *ETHICS* 764 (2003).

him) would be stronger than Super-Richie's (similarly understood). Why, now, does Richie's comparatively stronger claim evaporate if both he and Super-Richie are above some well-being threshold? If the Pareto principle doesn't disappear above a threshold, why does the Pigou-Dalton principle?

Relatedly, it is unclear how the sufficientist means to specify the limitation on the scope of the Pigou-Dalton principle. Supposedly, a Pigou-Dalton transfer, in well-being, to someone whose well-being level is above some threshold, is a matter of normative indifference—not an improvement in equity. But where *is* that normative threshold?

Sufficientism offers one—controversial—justification for the restriction in scope of the Pigou-Dalton principle embodied by poverty metrics. The traditional argument for poverty lines is different. This argument sees the poverty line with respect to some individual attribute (be it income, or some non-income attribute) as the threshold for satisfying minimum human *needs* (whether defined in absolute terms, for example as what someone needs for physical survival, or in relative terms, as what someone needs to be accorded basic respect in her society).¹¹⁴ In effect, the poverty line marks a “discontinuity” in the relation between that attribute and an individual's well-being.¹¹⁵ Below a certain level of income, someone's life is *qualitatively* worse: she becomes socially excluded because of her inability to buy the material goods that society values, may lack money to purchase sufficient food, etc. Similarly, someone who lacks reliable shelter, or minimal education, or companionship, is living a categorically different (and worse) kind of life than someone whose basic needs are satisfied.

This line of analysis helps to justify poverty metrics *as one component of an overall assessment of equity*, if we lack a workable inclusive measure of well-being (in particular, an interpersonally comparable utility indicator). Absent such a measure, it seems quite plausible to bifurcate the assessment of equity. First, what progress is society making in reducing the very grave inequity that occurs when some individuals lack minimally decent lives? Second, how inequitable is the distribution of economic and non-economic resources more generally—including, now, inequities between those who have enough in all dimensions to be non-poor, and others yet better off? The first step in this bifurcated analysis would involve a *poverty metric*. It would mean focusing on the number and attainment of individuals who have below-threshold incomes or non-income attainments,

114. See, e.g., Amartya Sen, *Issues in the Measurement of Poverty*, 81 SCANDINAVIAN J. ECON. 285, 288–89 (1979).

115. Sufficientism sees a discontinuity within moral assessment itself; there is some level of *well-being* at which the Pigou-Dalton principle is no longer morally compelling. By contrast, the view now under discussion sees a discontinuity in the *determination of individual well-being*—in how attributes produce well-being at the level of each individual.

and ignoring Pigou-Dalton transfers among non-poor individuals—for purposes of *this first step* (not the second).

It is less clear why such bifurcated assessment would be warranted given a workable, interpersonally comparable utility function.¹¹⁶ Such an indicator (to be fully adequate) would need somehow to reflect, in its mathematical structure,¹¹⁷ the discontinuity with respect to well-being that occurs around a poverty line. Utility (or the rate of increase of utility) would “jump” at that point. An equity-regarding SWF could then be applied to the distribution of utility in a single, unified assessment of equity—with the Pigou-Dalton principle taking utility (unrestrictedly) as its currency, and with no need for a poverty metric at all.

But, as explained in Part III, constructing an inclusive measure of well-being continues to pose challenges. Absent such a measure, the use of poverty metrics as one component of equity analysis—and the attendant insensitivity of this component to Pigou-Dalton transfers among non-poor—has substantial appeal.

B. Social-Gradient Metrics

Social-gradient metrics look at the comparative attainment of individuals with different social statuses—using some proxy for status.¹¹⁸ The concentration index is a dominant such metric, which often is employed to quantify social-status skews in *health*, but in principle can be used to quantify skews in any currency *C*. With some *C* as its input, the concentration index satisfies the Pigou-Dalton principle with respect to *C* where the transferor (the individual with more *C*) has higher social status than the transferee (the individual with less *C*). But the Pigou-Dalton principle is not satisfied (1) in cases where the transferor has *lower* social status than the transferee, nor (2) in cases where the two individuals have the same social status.¹¹⁹ Why would these restrictions be justified?

A straightforward argument for status-based restriction (1) is that having a higher status is itself a substantial boost to well-being. Thus, the

116. Cf. Martin Ravallion, *Measuring Social Welfare with and Without Poverty Lines*, 84 AEA PAPERS & PROCEEDINGS 359, 360 (1994) (suggesting that empirical difficulties in measuring well-being and uncertainties about the parameters of a social welfare function help to justify poverty metrics).

117. See John Creedy, *Labour Supply and Social Welfare when Utility Depends on a Threshold Consumption Level*, 73 ECON. RECORD 159 (1997); G.W. Lewis & D.T. Ulph, *Poverty, Inequality and Welfare*, 98 ECON. J. 117 (1988).

118. The desirability of social-gradient versus population-wide measures of health inequality is one of the themes of an important recent book: *INEQUALITIES IN HEALTH: CONCEPTS, MEASURES, AND ETHICS* (Nir Eyal, Samia A. Hurst, Ole F. Norheim & Dan Wikler eds., 2013). This Article was written prior to the publication of the book, and the analysis of social-gradient metrics that follows has not been refined to take account of it.

119. See *supra* note 75 and accompanying text.

Generic Justification often does not apply where the transferee has higher status, since her status advantage may compensate for her comparative deficit in *C* and make her better off than the transferor, all-things-considered. This argument, note, hinges on the absence of an inclusive measure of well-being that itself incorporates status effects.¹²⁰ Moreover, the argument does nothing to explain status-based restriction (2).

A different kind of argument—one that helps to explain *both* status-based restrictions—has to do with individual responsibility. Such an argument can be drawn from the work of Iris Marion Young, a philosopher who has provided one of the fullest defenses to date of the social-gradient format for conceptualizing equity. Young argues:

[I]f we simply identify some inequality of condition or situation between individuals at a particular time we have no account of the causes of this unequal condition. It is the causes and consequences of some pattern of inequality, rather than the pattern itself, that raise issues of justice. If the causes of an inequality lie in the uncoerced and considered decisions and preferences of the less well-off persons, for example, then the inequality is probably not unjust

. . . [However], the causes of many inequalities of resources or opportunities among individuals lie in social institutions, their rules and relations, and the decisions others make within them that affect the lives of the individuals compared.¹²¹

Young concludes that an inequality must be a “structural inequality” to be a central concern of distributive justice. “Structural inequality . . . consists in the relative constraints some people encounter in their freedom and material well-being as the cumulative effect of the possibilities of their social positions, as compared with others who in their social positions have more options or easier access to benefits.”¹²²

Young’s suggestion dovetails with the “responsibility-sensitive” strain of modern egalitarianism, pioneered by Ronald Dworkin and Richard Arneson.¹²³ Dworkin and Arneson’s insight was that equality of welfare

120. Cf. Thomas Aronsson & Olof Johansson-Stenman, *When the Joneses’ Consumption Hurts: Optimal Public Good Provision and Nonlinear Income Taxation*, 92 J. PUB. ECON. 986 (2008) (discussing utility functions where an individual’s well-being depends upon her relative as well as absolute income).

121. Iris Marion Young, *Equality of Whom? Social Groups and Judgments of Injustice*, 9 J. POL. PHIL. 1, 8. (2001).

122. *Id.* at 15.

123. Richard J. Arneson, *Equality and Equal Opportunity for Welfare*, 56 PHIL. STUD. 77 (1989); Ronald Dworkin, *What is Equality? Part 1: Equality of Welfare*, 10 PHIL. & PUB. AFF. 185 (1981); Ronald Dworkin, *What is Equality? Part 2: Equality of Resources*, 10 PHIL. & PUB. AFF. 283 (1981).

could not be the criterion of distributive justice, since well-being inequalities might be the result of poor individual choices, rather than bad luck or unfair conditions.

Individuals can arrive at different welfare levels due to choices they make for which they alone should be held responsible. A simple example would be to imagine two persons of identical tastes and abilities who are assigned equal resources by an agency charged to maintain distributive equality. The two then voluntarily engage in high-stakes gambling, from which one emerges rich . . . and the other poor In . . . [another] example, one person may voluntarily cultivate an expensive preference . . . , while another person does not. In . . . [these] examples it would be inappropriate to insist upon equality of welfare when welfare inequality arises through the voluntary choice of the person who gets lesser welfare.¹²⁴

Note that the Generic Justification is *alive* to the responsibility-egalitarian insight by including the transferee's lack of responsibility for being worse off as one condition for a justified transfer.

Now, Young's insight is that inequality caused by poor social position is sufficient to meet the responsibility-egalitarian test. If Leela has less income, or life expectancy, or happiness, or health than Morris *because of* the fact that Leela is lower status, that inequality is not Leela's fault. This insight is clearly correct, *if* social position itself largely flows from immutable characteristics such as race or gender. But what of the case where someone's social position may be partly determined by her own prior choices—for example, in societies where income substantially shapes social position?

Even leaving aside the mutability of social position itself, Young's analysis can be criticized for conflating a *sufficient* condition for unjust inequality with a *necessary* condition. Assume that Max, like Leela, is worse off than Morris, but (unlike her) shares Morris's social position. It hardly follows from responsibility-sensitive egalitarianism that the Max/Morris disparity is normatively neutral. Obviously, factors outside Max's control, *other than his social position*, might account for his being worse off than Morris. He might have been mistreated by his parents, victimized in a home invasion, hit by a negligent driver, systematically harassed by malicious coworkers, etc.

Much of the subsequent philosophical literature on responsibility-sensitive egalitarianism is cited in Nicholas Barry, *Reassessing Luck Egalitarianism*, 70 J. POL. 136 (2008).

124. Arneson, *supra* note 123, at 83–84.

Still, the proponent of the social-gradient approach to measuring equity might repackage Young's analysis—seeing social position not as a strictly necessary or sufficient condition for injustice but as a workable proxy for responsibility. If Leela is worse off than Morris, and has a lower social position than him, it is *much more likely* that she lacks responsibility for this divergence than in the cases of Max and Neil, who are worse off than Morris but—respectively—have equal or higher social position.

This observation, in turn, helps to rationalize both of the status-based restrictions inherent in the social-gradient approach—status-based restrictions (1) and (2), above—if finer-grained techniques for determining Leela's, Max's, and Neil's responsibility are not available. Let us now, therefore, consider “equality of opportunity” metrics—a recent development in scholarship on equity—which purport to offer a general methodology for parsing between responsibility and non-responsibility factors in the measurement of equity.

V. EQUALITY OF OPPORTUNITY METRICS

Systematic work on adjusting equity metrics to take account of individual responsibility begins, in the 1990s, with John Roemer's scholarship, itself inspired by the philosophical explorations, starting yet a decade earlier, of responsibility-sensitive egalitarians such as Dworkin and Arneson.¹²⁵ Roemer proposes to separate between an individual's “circumstances” (the characteristics for which she lacks responsibility) and her “effort” (the characteristics for which she bears responsibility).¹²⁶ An individual's level of attainment, with respect to some currency (be it health, income, utility, etc.), given some government policy, would be a function *both* of her “circumstances” *and* of her “effort.”¹²⁷

The population can, in principle, be separated into different “circumstance” classes, defined by a particular cluster of “circumstance” attributes. For example, if the individual's race, urban or rural place of birth, and parental education (defined, e.g., as whether both parents

125. For presentations and theoretical discussions of Roemer's approach, see JOHN E. ROEMER, *EQUALITY OF OPPORTUNITY* (1998); John E. Roemer, *On Several Approaches to Equality of Opportunity*, 28 *ECON. & PHIL.* 165 (2012); John E. Roemer, *Defending Equality of Opportunity*, 86 *MONIST* 261 (2003); John E. Roemer, *Equality of Opportunity: A Progress Report*, 19 *SOC. CHOICE & WELFARE* 455 (2002); MARC FLEURBAEY, *FAIRNESS, RESPONSIBILITY & WELFARE* ch. 8 (2008); Matthias Hild & Alex Voorhoeve, *Equality of Opportunity and Opportunity Dominance*, 20 *ECON. & PHIL.* 117 (2004); Erwin Ooghe et al., *Equality of Opportunity versus Equality of Opportunity Sets*, 28 *SOC. CHOICE & WELFARE* 209 (2007); Mathias Risse, *What Equality of Opportunity Could Not Be*, 112 *ETHICS* 720 (2002). Empirical applications are cited *infra* note 134.

126. John E. Roemer, *Equality of Opportunity: A Progress Report*, 19 *SOC. CHOICE & WELFARE* 455 (2002).

127. *Id.*

graduated high school) are each “circumstance” attributes, then white individuals born in cities whose parents have both graduated high school is one circumstance class, white individuals born in rural areas whose parents have both graduated high school is another circumstance class, etc.

Similarly, the population can be divided into different “effort” classes. “Effort” as well as “circumstance” may be difficult to observe—raising econometric questions—but in principle each cluster of “effort” attributes defines an “effort” class. For example, if the *individual's* education level (as contrasted with the level of education of her parents), and whether she engages in regular exercise, are “effort” attributes, then high school-educated individuals who engage in regular exercise would be one “effort” class, high school-educated individuals who don’t regularly exercise a second, etc.

A governmental policy, then, corresponds to a *matrix*.¹²⁸ The rows of the matrix are “circumstance” classes. The columns are “effort” classes. The entry in each “cell” of the matrix is the attainment level, in the chosen currency, of individuals in the corresponding effort class and circumstance class. A policy will yield a particular such attainment level for each cell (circumstance-effort pairing); there will also be a particular number of individuals at each circumstance-effort pairing. For simplicity, my presentation of the Roemer approach and variations will assume that the number of such individuals is the same.¹²⁹

How to rank matrices? Roemer proposes to do so in a manner that will be sensitive to equity, but more specifically will see inequality of attainment due to variation in circumstance—*not* inequality of attainment due to variation in effort—as inequitable. Yet more specifically, Roemer suggests the following responsibility-egalitarian rule for ranking matrices. (1) For a given policy P , and a given effort class e , determine the *lowest* level of attainment of any circumstance class making effort e . Call this value $\min(P, e)$. (2) Assign P an overall value by summing the $\min(P, e)$ values across all the effort classes. (3) Choose the policy with the greatest such value.

An obvious objection to the Roemer approach is that it focuses solely on the *worst-off* individual within each effort grouping. Assume that there are three or more circumstance classes, and that policy P' —within each effort class—improves the attainment of the second-worst-off individual, at the expense of the best-off individual. Then P' is, plausibly, both (1) a more

128. Or, more generally, to a probability distribution over such matrices—but issues of uncertainty are ignored here. See *supra* note 77.

129. See Ooghe et al., *supra* note 125, at 211, 214 (adopting this assumption for Roemer-leximin and discussing how it might be relaxed).

egalitarian policy than P , and (2) a better policy than P , all things considered. But the Roemer rule will rank P' and P equally good.

However, there is a rule closely related to the Roemer approach—call it Roemer-leximin—which has the effect of preferring P' to P . Roemer-leximin compares two policies by summing the lowest attainments in each effort class (if there are E effort classes, P is assigned the sum of the E lowest-in-effort-class values resulting from P , and P' the sum of the E lowest-in-effort-class values resulting from P'). If these sums are equal, the policies are compared by summing the second-lowest attainment levels within each effort class; if those are equal, by summing the third-lowest attainment levels within each effort class; etc.¹³⁰

How, exactly, does Roemer-leximin achieve the “trick” of seeing inequality due to circumstance, but not effort, as inequitable? To see this, intuitively, let us consider an especially simple case. Assume that policy P is such that circumstance classes can be unambiguously ranked from worst to best.¹³¹ Imagine, first, a policy P' that erases all inequality within each effort class. P' is such that each effort class (each column of the matrix) has exactly the same attainment level—namely, the average value of the corresponding column in P . Imagine, now, a different policy P^* which erases all inequality within each *circumstance* class. Each such class (each row of the matrix) has exactly the same attainment level—namely, the average value of the corresponding row in P . Finally, P^+ achieves full equality: every cell of the matrix has the same value (namely, the average value of all cells in P).

Note, now, that Roemer-leximin will see P^* (merely erasing inequality due to effort variation) as no improvement at all on P . By contrast, P' (erasing inequality due to circumstantial variation) is ranked as an improvement. And P^+ (complete equality) is no further improvement: P^+ and P' are ranked equally good by Roemer-leximin. See Table 5.

130. See *id.*

131. In other words, there are circumstance classes c_1, c_2, c_3, \dots ; and, with policy P , the highest effort in c_1 yields a lower attainment than the lowest effort in c_2 , the highest effort in c_2 yields a lower attainment than the lowest effort in c_3 , etc. Again, this is an especially simple case; but it serves to illustrate in a dramatic way how Roemer-leximin separates between inequality due to circumstance and effort.

Table 5: The Roemer-Leximin Approach to Equity:

Policy <i>P</i>				Policy <i>P'</i>			
	Effort:				Effort:		
	Low	Moderate	High		Low	Moderate	High
Circumstances T_1	50	80	110	→	25	40	55
T_2	20	30	40		25	40	55
T_3	5	10	15		25	40	55

Policy <i>P*</i>				Policy <i>P'+</i>			
	Effort:				Effort:		
	Low	Moderate	High		Low	Moderate	High
Circumstances T_1	80	80	80	→	40	40	40
T_2	30	30	30		40	40	40
T_3	10	10	10		40	40	40

More formally, and generally, the way in which Roemer-leximin bifurcates between effort and circumstances can be expressed in terms of our old friend: the Pigou-Dalton principle. First, Roemer-leximin *fails* the ordinary Pigou-Dalton principle. If two individuals are in different effort classes, one reaching a higher attainment level than a second, a Pigou-Dalton transfer between them is *not* necessarily preferred by Roemer-leximin.¹³² Second, Roemer-leximin *satisfies* a restricted Pigou-Dalton principle—with the restriction, now, expressed not in terms of poverty lines or social groups but in terms of effort classes.

Pigou-Dalton Principle: Effort Class Restricted

Assume that policy *P* produces some distribution of currency *C* across the entire population, and policy *P'* produces an alternative distribution, which represents a Pigou-Dalton transfer within *one* effort class *e*, between everyone within that class at some higher level of *C*, to everyone at some lower level. Then policy *P'* is a normative improvement over *P*.¹³³

132. Assume for simplicity that there is one individual within each circumstance-effort grouping, and that the transfer does not change either the transferee's or the transferor's rank within their effort classes. Then such a transfer will be seen as an improvement (yielding a higher-ranked matrix) by the Roemer-leximin rule if the transferee's rank within her effort class is lower than the transferor's rank within *her* effort class; it will produce no change with respect to that rule if the ranks are equal; and it will be seen as a *worsening* (yielding a lower-ranked matrix) by the Roemer-leximin rule if the transferee's rank within her effort class is higher than the transferor's rank within hers.

133. The proviso that a Pigou-Dalton transfer must occur between everyone at the two currency levels is to maintain consistency with the simplifying assumption that, in all policies, the number of individuals in each circumstance-effort grouping is the same. See *supra* note 129 and accompanying

Since Roemer's pioneering contributions, the literature on responsibility-sensitive equity metrics has blossomed. Such metrics are often, now, referred to as "equality-of-opportunity" metrics, although that is something of a misnomer—since the unifying feature of these metrics is not "opportunity" per se, but rather the "effort"/"circumstance" bifurcation—the distinction between those sources of differential attainment for which the affected individuals bear responsibility, versus those for which they don't. Work in this vein includes empirical scholarship by Roemer and collaborators,¹³⁴ as well as the development of alternative approaches for constructing responsibility-sensitive metrics.¹³⁵ Important alternatives to the Roemer methodology include these:

(1) Equalizing the "options" of each type. As in the Roemer approach, each policy corresponds to a matrix, with circumstance groupings corresponding to rows and effort classes to columns. We assign each row a single "option" value, as a function of every cell in that row. (Most simply, this value is just the average of the entries—meaning the attainment of an individual in that circumstance grouping if she is equally likely to exert any of the effort levels.) Then each policy corresponds to a list of "option" values, one for each circumstance. And policies are then ranked by applying an inequality metric or SWF to these lists.¹³⁶

(2) Calculating what inequality of attainment would be, if efforts were identical. For a given matrix of circumstances, efforts, and resulting attainments, we can choose a particular effort class and predict what the distribution of attainments *would* be if everyone's effort attributes were in that particular class, with their circumstances unchanged. And we can then

text. For purposes of real-world policy assessment, the Roemer-leximin rule would of course need to relax that assumption, and (depending on how the rule is specified for the different-number case) the proviso might be relaxed as well.

134. See, e.g., Humberto G. Llavador & John E. Roemer, *An Equal-Opportunity Approach to the Allocation of International Aid*, 64 J. DEV. ECON. 147 (2001); John E. Roemer et al., *To What Extent Do Fiscal Regimes Equalize Opportunities for Income Acquisition Among Citizens?* 87 J. PUB. ECON. 539 (2003).

135. Discussions of alternative approaches include: FLEURBAEY, *supra* note 125; Rolf Aaberge et al., *Measuring Long-Term Inequality of Opportunity*, 95 J. PUB. ECON. 193 (2011); François Bourguignon et al., *Inequality of Opportunity in Brazil*, 53 REV. INCOME & WEALTH 585 (2007); Daniele Checchi & Vito Peragine, *Inequality of Opportunity in Italy*, 8 J. ECON. INEQUALITY 429 (2010); Hild & Voorhoeve, *supra* note 125; Francisco H.G. Ferreira and Vito Peragine, *Individual Responsibility, in OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY* (Matthew D. Adler & Marc Fleurbaey eds., forthcoming 2015); Arnaud Lefranc et al., *Inequality of Opportunities vs. Inequality of Outcomes: Are Western Societies All Alike?*, 54 REV. INCOME & WEALTH 513 (2008); Juan D. Moreno-Ternero, *On the Design of Equal-Opportunity Policies*, 31 INVESTIGACIONES ECONÓMICAS 351 (2007); Ooghe et al., *supra* note 125; Nicolas Pistolesi, *Inequality of Opportunity in the Land of Opportunities, 1968-2001*, 7 J. ECON. INEQUALITY 411 (2009); Giuseppe Pignataro, *Equality of Opportunity: Policy and Measurement Paradigms*, 26 J. ECON. SURVEYS 800 (2012) (reviewing many empirical studies implementing Roemer's approach or alternatives); Juan Gabriel Rodríguez, *Partial Equality-of-Opportunity Orderings*, 31 SOC. CHOICE & WELFARE 435 (2008); Roemer, *On Several Approaches*, *supra* note 125.

136. See Ooghe et al., *supra* note 125.

ascribe a degree of inequality to the matrix by applying an inequality metric to this counterfactual distribution of attainments. (The counterfactual distribution, of course, depends on which effort grouping everyone is hypothesized as falling into.)¹³⁷

(3) Calculating what inequality of attainment would be, if circumstances were identical. Symmetrically, starting with a matrix of circumstances, efforts, and resulting attainments, we can choose a particular circumstance class and predict what the distribution of attainments would be if everyone's circumstances were in that class, with their effort attributes unchanged. The resulting counterfactual distribution is not unfair: it reflects *just* diversity in attainment due to diverse effort. Thus, by applying an inequality metric to the original distribution of attainments, and then *subtracting* the inequality number assigned to this counterfactual distribution, we arrive at a responsibility-adjusted inequality number for the original matrix.¹³⁸

Such alternative methodologies, like the Roemer-leximin approach, *fail* to satisfy the ordinary Pigou-Dalton principle: a transfer between individuals in different effort classes will not, necessarily, decrease the measured degree of inequity. However, like Roemer-leximin, they may well satisfy some appropriately restricted principle (focused, in some way, on transfers between individuals with different circumstances).

Responsibility-adjustment is an important development in scholarship on equity measurement. First, it seeks concrete tools by which to express the compelling, but quite abstract, normative insights of Dworkin, Arneson, and other philosophers in their camp—tools more fine-grained than social-gradient metrics. Second, (like any substantially new equity-measurement methodology) it may well offer a new “take” on any specific question of equity. Consider, for example, the question mooted in the Introduction: has the United States become a less equitable society over the last thirty-five years? Pistolessi addressed this question using data from the Michigan Panel Study of Income Dynamics, which surveys a large group of U.S. individuals annually—ascertaining both the respondent's earned (labor) income and certain other characteristics, in particular: the respondent's age, the educational level of his parents, his father's occupation, the respondent's race, his region of birth, his own educational attainment, and the number of hours worked annually.¹³⁹ Pistolessi classified the last two characteristics as “effort” attributes, and the others as “circumstance” attributes. He then estimated responsibility-adjusted inequality of U.S. income, on an annual basis, using the sample data and the two

137. See Pistolessi, *supra* note 135, at 414–15.

138. See Bourguignon et al., *supra* note 135; Pistolessi, *supra* note 135, at 414–15.

139. See Pistolessi, *supra* note 135.

“counterfactual” methodologies described in the previous paragraph. Pistoletti concluded that, while inequality of income increased dramatically between 1980 and 2000, responsibility-adjusted income inequality hardly increased at all.¹⁴⁰

To be sure, the day still seems far off when the responsibility-adjustment approach becomes the dominant method for quantifying equity—displacing non-adjusted inequality metrics, poverty metrics, social welfare functions, and social-gradient metrics. The approach (at least as presently developed) has various limitations:

(1) Distinguishing between circumstances and efforts. The approach, at its core, categorizes some attributes as efforts and others as circumstances. Such differentiation will inevitably be controversial. (For example, Pistoletti’s conclusion about the time trend of U.S. income inequality hinges on the quite contestable view that an individual is responsible for his education level and labor hours.) The responsibility-adjusted egalitarian might counter that such controversies are a *virtue* of the approach—that the approach brings to the surface normative questions that are both difficult, but also vital, for any truly nuanced assessment of equity, and that are ignored or “buried” by other methodologies. For example—she might say—the social-gradient methodology implicitly counts social class as a “circumstance” and all other attributes as an “effort.” Still, it is relevant to note—as a predictive, not normative matter—that the responsibility-sensitive approach may fail to gain widespread traction just because any specific result will be vulnerable to criticism by those who reject the attendant effort/circumstance differentiation.

(2) The diversity of approaches. A diverse array of distinct techniques exist for normalizing equity to “wash out” differential effort. I have described four, and the literature has explored yet further methods. Progress now needs to be made in clarifying the pros and cons of the different formats, and (ultimately) achieving some consensus on their relative preferability—akin to the convergence of the inequality-measurement literature on a relatively short list of standard metrics, or the poverty-measurement literature on a preference for the FGT class.

(3) Effort as a residual. Some approaches (not all) effectively define effort as a residual. For example, in much of his work, Roemer actually “infers” an individual’s effort from her attainment—by observing her percentile of attainment in the distribution of attainments for her circumstance—rather than using some independent attributes (labor hours,

140. A more recent empirical study of responsibility-adjusted inequality in the U.S. is Gustavo A. Marrero & Juan G. Rodriguez, *Inequality of Opportunity in the United States: Trends and Decomposition*, in *INEQUALITY OF OPPORTUNITY: THEORY AND MEASUREMENT* 217 (Juan Gabriel Rodriguez ed., 2011).

educational attainment, etc.) as a direct marker of effort.¹⁴¹ But doing so tends to overstate the role of differential effort in producing observed inequality, and (conversely) to understate the degree of genuine inequity associated with any given outcome or policy.¹⁴²

CONCLUSION

This Article has surveyed a wide range of equity-sensitive metrics, and has brought to light their shared nexus to the Pigou-Dalton principle. All satisfy that principle in *some* form. Traditional income-inequality metrics do so with respect to income: a pure, non-rank-switching transfer of income, from someone with more income to someone with less, leaving everyone else's income unchanged, reduces the degree of income inequality. Similarly, equity-regarding SWFs satisfy the Pigou-Dalton principle with respect to "utility"; multidimensional inequality and poverty metrics, with respect to a multiplicity of well-being relevant attributes; poverty metrics, in a "restricted" fashion defined by a poverty line; social-gradient metrics, with a different sort of "restriction," involving the relative social positions of transferee and transferor; and responsibility-sensitive approaches, with a normalization that takes account of individual "effort."¹⁴³ Risk and time create further orthogonal possibilities.¹⁴⁴

Differences in the specification of the Pigou-Dalton principle do not account for *all* differences in equity metrics,¹⁴⁵ but go a long way. The Pigou-Dalton principle is truly the heart of equity, just as the Pareto principle is the heart of efficiency. But Pigou-Dalton has more "open texture" than the Pareto principle¹⁴⁶—it seems to invite a wider range of plausible specifications—and it is this multiplicity that largely accounts for the diversity of equity metrics.

The Article has not only taxonomized equity metrics, using the Pigou-Dalton principle as the key organizing principle, but I have also taken a normative perspective—a welfarist one.¹⁴⁷ Welfarism is *not* utilitarianism.

141. See Hild & Vorhooeve, *supra* note 125, at 122–24; Ooghe et al., *supra* note 125, at 211; Roemer, *Defending Equality of Opportunity*, *supra* note 125, at 267–69; Roemer, *Equality of Opportunity*, *supra* note 125, at 461–63. See also Checchi & Peragine, *supra* note 135, at 432 (discussing how their framework defines effort as a residual).

142. Indeed, Roemer concedes as much. See Roemer, *Equality of Opportunity*, *supra* note 125, at 463.

143. See *supra* Part I.

144. See *supra* note 77.

145. See *supra* text accompanying note 77.

146. As regards the open texture of the Pareto principle, see ADLER, *supra* note 28, at 52–53. Note that the present Article employs economists' specific version of the principle, as a principle concerning preferences, and not the more general version which John Broome has termed the "principle of personal good" and which *Well-Being and Fair Distribution* employs.

147. See *supra* Part II.

Welfarism and a concern for fair distribution can be—indeed, very plausibly *should* be—combined.¹⁴⁸

From this perspective, many of the current formats for measuring equity are imperfect. A better approach, in principle, would be to employ an equity-regarding SWF, suitably adjusted for responsibility.¹⁴⁹ A well-designed utility function would be sensitive to multiple sources of well-being (unlike income, which ignores non-market sources);¹⁵⁰ would covary with preferences (and thus avoid the conflicts with the Pareto principle that afflict multidimensional metrics);¹⁵¹ would have a discontinuity around the level of basic needs (obviating the need for poverty metrics);¹⁵² and would use social status as one determinant of utility (obviating one argument for social-gradient metrics). Systematic adjustment for differential effort would obviate a second argument for social-gradient metrics.¹⁵³

The key obstacles to this “first best” approach are twofold: (1) overcoming the heterogeneity of preferences in designing an interpersonally comparable measure of utility,¹⁵⁴ and (2) arriving at a consensus format for the “effort” adjustment,¹⁵⁵ so that utility is *appropriately* normalized to take account of individual responsibility. Some progress has been made on these two fronts, but much more remains to be done. Putting a number on injustice is not a quixotic or hopeless task, but neither is it smooth sailing. Doing so means patiently refining tools and methodologies to mirror underlying normative commitments—whether those commitments center around human well-being or something else.

-
148. See generally ADLER, *supra* note 28; Arneson, *supra*, note 90.
149. See *supra* Part V; ADLER, *supra* note 28, at 579–84.
150. See *supra* text accompanying note 86.
151. See *supra* text accompanying note 94.
152. See *supra* text accompanying notes 114–117.
153. See *supra* Section IV.B.
154. See *supra* text accompanying notes 99–105.
155. See *supra* Part V.