

ON THE LIMITATIONS OF SOME CURRENT USAGES OF THE GINI INDEX

BY LARS OSBERG*

Dalhousie University

This note constructs a simple two class example in which the Gini index is held constant while the size of the rich and poor populations change, in order to illustrate how very different societies can have the same Gini index and produce very similar estimates of standard inequality averse Social Welfare Functions. The rich/poor income ratio can vary by a factor of over 12, and the income share of the top one per cent can vary by a factor of over 16, with exactly the same Gini index. Focussing solely on the Gini index can thus obscure perceptions—e.g. of important market income trends or large changes in the redistributive impact of the tax and transfer system. Hence, analysts should supplement the use of an aggregate summary index of inequality with direct examination of the segments of the income distribution which they think are of greatest importance.

JEL Codes: D63, D30, D31, H23

Keywords: inequality, Gini, redistribution effort, Social Welfare

Discussions of economic inequality often begin with the seemingly simple question: “Has economic inequality increased or decreased?” Since every society has many different types of economic resources, used by many different people at different points in time, answering this question requires both measurement choices about *what* economic resource is being distributed among *whom*, and *when* (i.e. over what period of time) and conceptual choices about indices of inequality. In practice, most analysts choose to focus on the distribution of annual disposable money income among households and it is very common to summarize the inequality in that distribution using the Gini index.

Although Atkinson (1970) showed long ago that Lorenz dominance is a necessary condition for unambiguous inequality ranking and that, because Lorenz dominance is rare, different indices of inequality will often disagree in the ranking of inequality, the potential ambiguities of inequality indices have very often been forgotten in the recent literature. In popular and professional discussion, it has become common to observe authors referring to “trends in inequality” rather than using the more cumbersome, but more accurate, language of “trends in the Gini index of household income inequality.” Indeed, some recent authors—e.g. Gale, Kearney, and Orszag (2015)—use the size of changes in the Gini index of household disposable income as synonymous with the size of changes in

Note: Thanks to Peter Burton, Casey Warman and two anonymous referees for helpful comments on earlier drafts. Remaining errors are mine alone.

*Correspondence to: Lars Osberg, Economics Department, Dalhousie University, 6214 University Avenue, Halifax, Nova Scotia B3H 3J5, Canada (lars.osberg@dal.ca)

“inequality” and derive important policy conclusions from the small size of changes in the Gini index. The purpose of this note, therefore, is to illustrate some of the limitations of using the Gini index in this way.

In recent years, Canada has provided an empirical example of the importance of distinguishing between changes in an income distribution and changes in an index number which tries to summarize an income distribution. In Canada, the Gini index of annual disposable household money income increased significantly between 1980 and 2000, but remained fairly flat from approximately 2000 to 2011. Over the same period, until the Great Recession of 2008, the income share of the top 1 percent in Canada increased strongly. Nevertheless, the approximate constancy of the Gini index produced, in the popular press (e.g. Coyne, 2013), statements that since 2000 “inequality has not increased in Canada.” Burkhauser *et al.* (2016, p. 2) have noted that similar statements have appeared in the British press, based on similar observation of a recently flat Gini index for household income and similar disregard of a continued increase in the top one percent income share. Such journalistic statements assume that there is no important difference between change in the Gini index of inequality and changes in the inequalities which that index tries to summarize—and implicitly some professional economists seem to agree. However, as this note illustrates, societies that are intuitively very different in income distribution can have the same Gini index (and generate very similar standard inequality averse Social Welfare Function estimates). This note therefore argues that discussion of trends in an aggregate index of inequality, calculations of aggregate Social Welfare and estimates of the redistributive effort of the tax/transfer system should always be supplemented by explicit analysis of changes in the distribution of income in the regions of the income distribution which are of primary concern to the researcher.

Mathematically, the Gini index (G) is calculated as the average of the absolute value of the relative mean difference in incomes between all possible pairs of individuals, as in Equation (1).

$$(1) \quad G = \frac{1}{2\bar{y} \cdot n \cdot (n-1)} \cdot \sum_{i \neq j} \sum_j^n |y_i - y_j|$$

where y_i, y_j = income of individuals i and j , n = total population size, \bar{y} = average income of all individuals.

Subramanian (2015) and Majumdar (2015) have recently suggested extensions to the Gini index but the enduring appeal of the original Gini index as a summary measure of inequality undoubtedly owes a great deal to its easy graphical representation. In Figure 1, the horizontal X axis measures the cumulative percentage of income recipients while the vertical Y axis measures the cumulative percentage of income received. If all income recipients are ordered by income from lowest to highest, then the curved solid line is the well-known Lorenz curve, which plots $N(Y)$ the cumulative share of income (Y) received by the poorest proportion N of the population. If everyone had identical income the Lorenz curve would be the 45 degree straight line od —sometimes called the line of perfect equality. If we define the area between od and the Lorenz curve as A and if B is defined




Figure 1. The Lorenz Curve and the Gini Index [Colour figure can be viewed at wileyonlinelibrary.com]

as the area within triangle Δocd but outside the Lorenz curve, the Gini index is equal to the ratio of A to the area of the triangle Δocd , which can be expressed as in [2]. [Note that when oc and cd are normalized to 1, then $\Delta ocd = A + B = 1/2$.]

$$(2) \quad G = A / (A + B) = A / \Delta ocd$$

However, there are many different possible income distributions and corresponding Lorenz curves for which area A remains constant. In particular, Osberg (1981, p. 14) invented a very simple example—“Adanac”—in which the rich and the poor were clear types. Using lower case letters to refer to magnitudes, Adanac has n_1 poor people all earning an identical income y_1 and n_2 rich people all receiving income y_2 . Using upper case letters to denote proportions, the income share of the poor in Adanac is Y_1 and the population share of the poor is N_1 , while the income share of the rich is $Y_2 [= 1 - Y_1]$ - and the population share of the rich is $N_2 [= 1 - N_1]$. One might intuitively think that Adanac would be a very different sort of place to live than Canada—but the Gini index for Adanac and for Canada can easily be made the same. The Lorenz curve for such an imaginary society is defined by (N_1, Y_1) and would look like line obd in Figure 2. Figure 2 constructs the triangle Δobd by drawing a 45 degree line ac starting from point c , which will be perpendicular to line od at point a . If one then selects point b along line ac such that $ab / ac = G$, and constructs triangle Δobd , the area of Δobd will equal A .

$$(3) \quad \Delta obd / \Delta ocd = A / \Delta ocd = G$$

The Lorenz curve of income in real societies is curved and continuous, because it represents the distribution of incomes derived from different sources by




Figure 2. The Lorenz Curve: Adanac and Canada Compared [Colour figure can be viewed at wileyonlinelibrary.com]

many individuals who vary in many dimensions, to degrees both small and large. However, in economics it is often useful to abstract from the messy shadings of real world distributions and discuss ideal types. For example, as Acemoglu and Autor (2010, p. 1) note, the “canonical model” for analysis of the impacts of technical change on inequality divides workers into two categories—skilled and unskilled—with corresponding wage rates, which implies a two class society and a Lorenz curve qualitatively similar to that of Adanac.

This note uses the labelling convention that when the size of the poor population of Adanac is set at 80 percent, one calls this example $Adanac_{80}$. In $Adanac_{80}$, when the bottom 80 percent of the population share 40 percent of total income (i.e. each bottom quintile gets 10 percent of total income), the top 20 percent all have an income exactly six times higher, and thus get 60 percent of total income, and the Gini index is 0.4. In real world Canada in 1981, the poorest 20 percent of households actually got 4.6 percent of total pre-tax income (compared to 10 percent in $Adanac_{80}$) and the richest 20 percent of Canada got 41.6 percent

(much less than their 60 percent share in Adanac₈₀). Hence, the 6:1 rich/poor income ratio implies that in Adanac₈₀, both the richest and the poorest income quintiles would get a substantially bigger share of national income than in real world 1981 Canada. However, although in 1981 Canada the second quintile got 11 percent of income, which is roughly the same as its 10 percent in Adanac₈₀, the middle and upper middle class of Adanac₈₀ would be considerably worse off in income share. In real world 1981 Canada, the third quintile got 17.7 percent of income (compared to a 10 percent share in Adanac₈₀) and the fourth quintile of real-world Canada got 25.1 percent of income (much more than their 10 percent of Adanac₈₀ income).

If the richest and the poorest quintiles get much more, while the lower middle class gets roughly the same and the middle and upper middle quintiles get significantly less, is society more equal or more unequal? Real world 1981 Canada had a Gini index for household money income before tax (unadjusted for household size) of 0.4. These numbers for Adanac₈₀ were thus picked to generate exactly the same Gini index—and the purpose was to illustrate the fact that the Gini index alone cannot say whether it is Adanac₈₀ or real world 1981 Canada which is more unequal.

In fact, there are many “Adanacs” which can have the same Gini index as real world Canada. In Figure 2, if one draws a straight line *egbf* through *b* which is parallel to *od*, then any point on that line could be the vertex of a triangle whose base is *od*, and all of these triangles would have an area equal to Δobd . Since we have defined units of measurement such that $oc = cd = 1$, one can show by similar triangles that $df = oe = og = ab / ac = G$. Hence, the line *ebf* is defined by $Y = X - G$. In a two class society with a Gini index equal to *G*, the population share (N_1) and Income share (Y_1) of the poor are bound by the relationship that:

$$(4) \quad Y_1 = N_1 - G$$

As well, in a two class society populated only by n_2 identical rich people with income y_2 and n_1 identical poor people with income y_1 , a simple rich/poor income ratio relation holds (see Appendix A for proof):

$$\text{Since } y_1 = Y_1 / N_1 = (N_1 - G) / N_1 \text{ and } y_2 = [1 - (N_1 - G)] / (1 - N_1)$$

$$(5) \quad \frac{y_2}{y_1} = \frac{N_2 + G}{N_1 - G} \cdot \frac{N_1}{N_2}$$

Using equation (5), one can calculate what the income ratio between rich and poor would be, for given *G*, if the affluent became a smaller or larger proportion of the population—i.e. as a two class society moves along the line *egbf*. In Figure 2, the dotted line *ob'd* shows the Lorenz curve for Adanac₉₉, which has 99 percent poor people and 1 percent rich. The point *b'* on the line segment *gbf* is the vertex of triangle $\Delta ob'd$ with area also equal to *A*. The line *ob'd* plots the Lorenz curve for this two class society, and the co-ordinates of *b'* define the percentage of the population that is poor (i.e. $N_1 = 0.99$), implicitly establishing the percentage that is rich [$N_2 = (1 - N_1) = 0.01$].

TABLE 1
THE INCOMES OF ADANAC WHEN THE RICH GET FEWER AND RICHER

Gini	Percentage Population Poor	Rich/Poor Income Ratio	If weighted average = \$45,000		Quintile shares					
			Income Of Poor \$	Income Of Rich \$	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	TOP 1% SHARE
0.317	60%	3.8	21,225	80,663	9.4%	9.4%	9.4%	35.9% ¹	35.9%	1.8% ²
	65%	3.7	23,054	85,757	10.2%	10.2%	10.2%	31.1%	38.1%	1.9%
	70%	3.8	24,621	92,550	10.9%	10.9%	10.9%	26.0%	41.1%	2.1%
	75%	3.9	25,980	102,060	11.5%	11.5%	11.5%	20.0%	45.4%	2.3%
	80%	4.3	27,169	116,325	12.1%	12.1%	12.1%	12.1%	51.7%	2.6%
	85%	5.0	28,218	140,100	12.5%	12.5%	12.5%	12.5%	49.8%	3.1%
	90%	6.4	29,150	187,650	13.0%	13.0%	13.0%	13.0%	48.2%	4.2%
	95%	11.0	29,984	330,300	13.3%	13.3%	13.3%	13.3%	46.7%	7.3%
	99%	48.1	30,591 ³	1,471,500	13.6%	13.6%	13.6%	13.6%	45.6%	32.7%
ACTUAL CANADA 2000⁴					7.1%	13%	17.6%	23.3%	39.1%	12.7%

As Appendix B (Table S1) shows, in Canada adjusting for household size and taking account of income taxes, the Gini index of inequality in the distribution across individuals of equivalent disposable income was 0.285 in 1981, increasing to 0.317 in the year 2000 but after that remaining rather flat. If one holds the Gini index constant at 0.317 and keeps average disposable equivalent income unchanged at \$45,000 (which is the 5 year average 2007–2011), one can calculate income distribution statistics for Adanac with varying percentages poor, as Table 1 reports. [Each row in Table 1 thus corresponds to a different Lorenz curve *obd* as *b* moves along the line *egbf*.]

One way of thinking about Table 1 is as the statistical reports from a hypothetical two-class society (Adanac) in which successive 5 percent segments of the population are downwardly mobile in market income (perhaps due to the “canonical model” of skill-biased technical change), but the Gini index remains constant and average disposable equivalent income remains unchanged.

Alternatively, one could imagine a series of tax and transfer policy changes in which everyone’s market income is unchanged and income is taken from successive 5 percent groups of the affluent and redistributed to both the previously poor and to the remaining rich in such a manner as to keep the Gini index of post-tax, post-transfer equivalent income constant. [For present purposes, it is unimportant whether the unlucky 5% are selected randomly at each stage or by some other, non-income criterion—e.g. height or IQ test score.]

Either way, as Adanac₆₀ becomes Adanac₆₅ and morphs into Adanac₇₀ and then into Adanac₇₅ and Adanac₈₀, continuing right up to Adanac₉₉, the 5 percent

¹Income share of middle class maximized (9.4 percent + 9.4 percent + 35.9 percent = 54.8 percent).

²Income share of top 1 percent minimized (1.8 percent).

³Absolute income of least well-off maximized (\$30,591).

⁴Source: CANSIM Table 202-0707.

who are downwardly mobile each time lose the differential in income between the rich and the poor. Their income loss is transferred to the other 95 percent of the population—i.e. partly to previously poor people and partly to the now smaller number of rich people. Hence, at each stage in this income redistribution process, income gainers vastly outnumber (19:1) income losers.

Cumulatively, as Adanac₆₀ becomes Adanac₉₅, the bottom 60 percent of the population experience a 41 percent increase in their incomes (from \$21,225 to \$29,984). As Adanac₉₅ then turns into Adanac₉₉ the income share of the top 1 percent and the rich/poor income ratio both more than quadruple (top 1 percent income share increasing from 7.3 percent to 32.7 percent and rich/poor income ratio increasing from 11:1 to 48.1:1). However, the income of the poorest quintile continues to increase (to \$30,591). Since the least well-off are always becoming better off, a Rawlsian social analyst would surely approve each stage of the transition from Adanac₆₀ to Adanac₉₅ or Adanac₉₉—and journalists and non-specialist academics could appeal to the unchanged Gini index as evidence that throughout this process there has “really” been no change in economic inequality in Adanac.

However, it would be more accurate to say that the Adanac scenarios provide examples of the insensitivity of the Gini index to changes in top one percent income share. In Table 1 the Gini index is constant at 0.317 but the income share of the top one percent is 1.8 percent in Adanac₆₀ and 32.7 percent in Adanac₉₉—varying by a factor of over 16—which illustrates concretely how much the very top end of the income distribution can change without affecting the Gini index.

Downward mobility of those at the margin of affluence is sometimes thought of in terms of “the disappearing middle class,” and the Adanac scenarios are a narrative of affluence becoming rarer and more extreme. But as Adanac₆₀ turns into Adanac₉₉, is Adanac becoming more unequal or less unequal or remaining just the same? By construction, the Gini index is absolutely constant but our intuition may say that Adanac is changing in fundamental ways. Adanac₆₀ is a place in which the income differential between the 60 percent who are poorer and the 40 percent who are richer is under \$60,000. In Adanac₉₉ there is an income differential of \$1.4 Million between the top 1 percent elite and everyone else. Adanac₆₀ is a place where two large groups exist in society, with a differential in standard of living (an income ratio of 3.8:1) which is miniscule compared to the income ratio (48.1:1) in Adanac₉₉, where a huge gulf separates a very small elite from everyone else.

However, the winners at each stage in the progression from Adanac₆₀ to Adanac₉₉ vastly outnumber the losers and the incomes of the least well-off always unambiguously increase. Hence, a clear majority would, if motivated solely by personal income, vote in favour of all of these changes. So what’s not to like about Adanac₆₀ democratically becoming Adanac₉₅ or even Adanac₉₉?

By construction, this note has thus far only considered inequalities of outcome, with no discussion whatsoever of the fairness or equity of the processes which generate unequal incomes. The Adanac example is, however, simple enough that it can also easily illustrate how the processes that generate income inequality matter to moral judgments about it. The top income group in any of the Adanac variations conceivably might have received their income in a lottery (which could be annual or once in a lifetime). If every ticket in the income lottery has the same

chance of getting the prize and each individual only has one ticket, all individuals are equal in an *ex ante*, before the income draw, “expected value of income” sense, even if their incomes become very unequal immediately after the tickets have been drawn. Alternatively, in an “age-set” society in which everyone has the same low income while young and the same high income when they are old, the affluent are just those people who have finally become old enough to receive higher incomes. A pure age-set society would have complete equality of annual income within age cohorts and complete equality in lifetime income (if everyone had the same lifespan)—but at any point in time there is inequality of annual income (i.e. inequality between age cohorts). However, any of the Adanac variations could also be a caste society in which the affluent have inherited their economic status from their parents and will pass it on exactly to their children. In such a society, there would be some inequality of outcome and complete inequality of opportunity. Finally, it is possible that Adanac might be a society in which people compete for elite membership and every year it is the hardest-working N_2 percent who all get identical high incomes—i.e. complete equality of opportunity, but inequality of effort and rewards.

If a general definition of “economic inequality” is “differences among people in their command over economic resources” Osberg (1981, p. 7), then economic inequity can be defined as “morally unjustifiable economic inequality.” Whatever the facts of economic inequality are, the causes of those facts—i.e. *why* some people are rich and others are poor—matter fundamentally for moral judgements about economic inequity. All these processes—chance, age, inheritance or effort—could conceivably generate the same facts of annual income inequality (i.e. the same observed distribution of annual income), but most observers will agree that these processes are not morally equivalent. Since our moral judgments often depend heavily on process, *inequality* is not the same concept as *inequity*.

However, this note concerns the measurement of inequality, and there is a long tradition within economics of abstracting from the processes that generate incomes and just comparing distributions of income. Formally, as the textbook presentation of Lambert (1989—especially Chapters 4 and 5) makes particularly clear, the standard method within neoclassical economics for assessing the welfare implications of inequality starts from the assumption that the Social Welfare Function (SWF) can be written as depending only on the vector of incomes $[y_1, y_2, y_3, y_4, \dots, y_n]$. Cowell (1995, p. 38) then summarizes nicely the standard requirements, for income inequality measurement purposes, that the Social Welfare Function is: [1] individualistic and non-decreasing (i.e. social welfare increases if any individual’s income increases); [2] symmetric (i.e. social welfare is unaffected if individuals trade places in the distribution of income); [3] additive ($SWF = \sum_i U(y_i)$ i.e. social well-being is the weighted sum of individual well-being across all individuals); [4] strictly concave ($\partial U(y_i)/\partial y_i < 0$ - i.e. the social welfare weight attached to an individual’s income decreases as income increases) and [5] has constant relative inequality aversion (which implies that $U(y_i) = [y_i^{1-\varepsilon} - 1]/(1-\varepsilon)$ where ε is the inequality aversion parameter and $\varepsilon > 0$).

A calculation of Social Welfare under these restrictions will, like the Gini index, produce a single index number which tries to summarize an entire distribution of income. As a little calculation with the incomes reported in Table 1 can

easily verify, as Adanac_{60} turns into Adanac_{99} there is very little change in Social Welfare for a constant relative inequality aversion Social Welfare Function, whatever the value of ε chosen within the range $0.1 \leq \varepsilon \leq 5$. Because the Gini index and average income are both held constant in all Adanac variants, the Social Welfare Function calculation yields the evaluation that Adanac_{60} is much the same as Adanac_{99} .

But even if the Gini index is the same and a standard Social Welfare calculation produces much the same number, many people might have the intuition that living in a society divided between a top 1 percent who get 48 times higher income than a homogeneously poor 99 percent (i.e. \$1,471,500 compared to \$30,591) would be a profoundly different experience from living in a society where 60 percent get a \$21,000 income and 40 percent make roughly \$80,000. Which is correct? Should one conclude that the Gini and Social Welfare calculations are right in concluding that all Adanac variants are much the same or is our intuition right in thinking that these calculations are missing something?

Of course, our intuitions come from our life experiences in the real world, which is not neatly divided into two homogeneous social classes and in which economic and political issues are not neatly separated. Our intuitions may thus be partly driven by real world experiences which suggest that the distribution of political power may be affected by the distribution of income, that status and social well-being may depend on economic differentials and that the utility which individuals experience may be affected by relative income and by consumption comparisons with others—issues which are all ignored in the standard Social Welfare Function methodology.

Thinking about the differences between Adanac_{60} and Adanac_{99} , and why the former might be preferred over the latter, may thus be useful as a gateway to analysis of aspects of economic inequality which may matter, even if these aspects might not show up in standard measures, like the Gini index or an inequality averse Social Welfare Function. Even if the Gini index is the same, and the Social Welfare Function is the same, some people might reasonably prefer to live in a society with relatively small differentials between large groups of people compared to living in a society with a gigantic income differential between a very small elite and everyone else. Although the Adanac examples force the realization that this preference is not Rawlsian, Adanac_{60} might quite plausibly be preferred to Adanac_{99} , particularly if the least well-off are not absolutely deprived. Such a preference for Adanac_{60} might be because of a hunch that small elites with very large income advantages tend to accumulate disproportionate political power, or it might be due to a hunch that very large income differentials can generate unattainable consumption norms which depress the well-being of those who cannot “keep up” in consumption.⁵

As well, the Adanac examples might suggest a closer examination of blanket assertions about changes in the redistributive role of taxes and transfers.

⁵This note uses the term “hunch” as a way of advertising the fact that no empirical evidence on these possible political and social implications of economic inequality is provided here. The only claim made here is that individuals with a subjective belief that Adanac_{60} and Adanac_{99} would differ politically and socially might reasonably choose one or the other.

Although Adanac₆₀ could turn into Adanac₉₉ as the result of a series of changes in market income, an alternative possible narrative is that an unchanging distribution of market income is combined with a series of tax and transfer policy changes in which income is taken from successive 5 percent groups of the affluent and is redistributed to both the previously poor and to the remaining rich. Since Aronson, Johnson and Lambert (1994a, 1994b), the redistributive effect of taxes and transfers has often been measured as the difference ($G_m - G_{pt}$) between the Gini index of household market income (G_m) and the Gini index of household income after taxes and transfers (G_{pt}). Monti, Pellegrino and Vernizzi (2015) extend this approach. In the recent empirical literature, Heisz and Murphy (2016) and the OECD (2015, Chapter 7) provide examples. However, since the Gini index of after-tax, after-transfer income in all Adanac scenarios is constant, $G_m - G_{pt}$ is also constant, and this method for assessing the redistributive impact of taxes and transfers would conclude that throughout the progression from Adanac₆₀ to Adanac₉₉, there was no change at all in the redistributive role of taxes and transfers—which seems intuitively quite wrong.

Piketty (2014, p. 266) has argued against the use of single summary indices of inequality, such as the Gini index, on the grounds that “The social reality and economic and political significance of inequality are very different at different levels of the distribution, and it is important to analyse these separately.” This note has used an artificial example to illustrate that perspective. As Adanac₆₀ morphs into Adanac₉₉, the rich/poor income ratio varies by a factor of over 12, and the income share of the top one percent varies by a factor of over 16, with exactly the same Gini index. Holding constant the Gini index, in Adanac₆₀, the income share of the middle 60 percent of the population is maximized and the income share of the top 1 percent is minimized. In Adanac₉₉, the absolute income of the least well-off is highest *and* the income gap between rich and poor is also maximized.

Some people worry about income inequality trends because they think that the well-being and size of the middle class is crucial to political stability. Some people are concerned about inequality because they care about the income share and the absolute income of the poorest. For others, the rich/poor income gap and the income share of the top one percent are the most important aspects of inequality, perhaps because of concerns about political democracy. These are all important aspects of economic inequality—but they are different aspects of inequality, have different implications, raise different policy issues and do not always coincide. Fundamentally, this note argues that analysts should be clear about which aspect of economic inequality they think matters most—e.g. elite concentration or middle class inclusion or the share or income level of the disadvantaged—and supplement the use of any single inequality index with direct examination of the relevant segment of the income distribution. One’s choice of inequality measures should always depend on *why* one wants to analyse inequality.

REFERENCES

- Acemoglu, D. and D. Autor, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” Working Paper 16082, National Bureau of Economic Research, 2010.

- Anwar, S. and S. Sun, "Taxation of Labour Income and the Skilled-Unskilled Wage Inequality," *Economic Modelling*, 47, 18–22, 2015.
- Aronson, J., P. Johnson and P. Lambert, "Redistributive Effect and Unequal Income Tax Treatment," *The Economic Journal*, 104, 262–70, 1994a.
- , "Decomposing the Gini Coefficient to Reveal Vertical, Horizontal and Reranking Effects of Income Taxation," *National Tax Journal*, 47, 273–94, 1994b.
- Atkinson, A., "The Measurement of Inequality," *Journal of Economic Theory*, 2, 244–63, 1970.
- Burkhauser, R., N. Héroult, S. Jenkins, and R. Wilkins, "What has Been Happening to UK Income Inequality since the Mid-1990s? Answers from Reconciled and Combined Household Survey and Tax Return Data," Working Paper 21991, National Bureau of Economic Research, 2016.
- Cowell, F., *Measuring Inequality*, 2nd ed., LSE Handbooks in Economics Series, Prentice Hall/Harvester Wheatsheaf, London, 1995.
- Coyne, A., "The Myth of Income Inequality: Since the Bleak 90s Things Have Actually Gotten Better," *The National Post*, 2 September 2013.
- Gale, W. G., M. S. Kearney, and P. R. Orszag, "Would a Significant Increase in the Top Income Tax Rate Substantially Alter Income Inequality?" Economic Studies at Brookings, 2015 <<http://www.brookings.edu/~media/research/files/papers/2015/09/28-taxes-inequality/would-top-income-tax-alter-income-inequality.pdf>>.
- Heisz, A. and B. Murphy, "The Role of Taxes and Transfers in Reducing Income Inequality," in D. A. Green, W. C. Riddell and F. St-Hilaire (eds), *Income Inequality: The Canadian Story*, The Institute for Research on Public Policy, Montreal, 435–78, 2016.
- Lambert, P., *The Distribution and Redistribution of Income: A Mathematical Analysis*, Basil Blackwell, Cambridge and Oxford, 1989.
- Majumdar, A., "An alternative measure of economic inequality under the Lorenz curve framework in analogue to the index of refraction of geometrical optics", *Economics Bulletin*, 35, 1076–86, 2015.
- Monti, M. G., S. Pellegrino and A. Vernizzi, "On Measuring Inequity in Taxation Among Groups of Income Units," *Review of Income and Wealth* 61, 43–58, 2015.
- Organisation for Economic Co-operation and Development, *In It Together: Why Less Inequality Benefits All*, Paris, OECD Publishing, 2015.
- Osberg, L., *Economic Inequality in Canada*, Butterworths, Toronto, 1981.
- Piketty, T., *Capital in the Twenty-First Century*, Harvard University Press, Cambridge, 2014.
- Subramanian, S., "More Tricks with the Lorenz Curve," *Economics Bulletin*, 35, 580–89, 2015.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix A: Proof

Appendix B: Figure S1 and Table S1