When Inequalities Diverge

By MICHAEL C. WOLFSON*

A significant innovation in discussions of income inequality is the addition, since the early 1980's, of the "disappearing middle class" (e.g., Robert Kuttner, 1983; Lester Thurow, 1984). This concept is typically equated with the concept of increased income inequality. However, such an equation entails a fundamental conceptual error. With apologies to Frank Levy and Richard J. Murnane (1992), a few quotes from their recent survey of trends in U.S. earnings inequality indicate the problem:

> ...a polarization of the earnings distribution means a decline in middle class jobs.... Despite the variety of scalar (inequality) measures, none seems well suited to the proposition of a vanishing middle class. That proposition refers to a polarization in which observations move from the middle of the distribution to both tails. Standard inequality measures cannot distinguish this polarization from other kinds of inequality.... If the middle of the male earnings distribution was hollowed out, that fact would be registered by scalar inequality measures.
>
> (pp. 1338, 1339, 1351)

Levy and Murnane thus recognize, correctly, that polarization (a shorthand for the phenomenon of the disappearing middle class) is not like the usual notion of inequality. However, they continue to use conventional scalar measures of inequality to assess the extent and trend in polarization.

Is this a problem? I shall argue that it is, first on theoretical and then on empirical grounds. Theoretically, Figure 1 should dispel any doubts. This graph shows two hypothetical income-distribution density functions. The first is a uniform or rectangular density over the interval 0.25–1.75, shown by a dashed line. The second density, shown by a solid line, is clearly bimodal and has a somewhat depleted middle. I would argue that, according to any sensible definition of polarization or disappearing middle, this latter density is the more polarized. Is it also more unequal?

The answer is unequivocally no. The second bimodal density has been constructed such that, according to *any* inequality measure that is consistent with the Lorenz criterion (the "gold standard" for the concept of inequality), it is more equal. In other words, the bimodal density has a Lorenz curve that is closer to the 45-degree line than the Lorenz curve for the uniform density. The formal proof follows simply from the fact that the bimodal distribution can be "derived" from the uniform distribution (in several ways, one of which is) by two sets of progressive mean-preserving redistributive transfers in the sense of A. B. Atkinson (1970), as indicated by the arrows in Figure 1. By

†*Discussants:* Daniel T. Slesnick, University of Texas; Lynn Karoly, Rand Corporation; James Heckman, University of Chicago.

*Statistics Canada and Canadian Institute for Advanced Research, Ottawa, K1A 0T6, Canada. I am greatly indebted to Tony Atkinson for suggesting a collaboration with James Foster to probe more deeply into the question of measuring polarization, to James Foster for our joint work in developing the measurement concepts, and to Brian Murphy, Geoff Rowe, and Milorad Kovacevic for support on the empirical and statistical work. I remain solely responsible for any errors or omissions, and for the views expressed. An extended version of this paper is available in Wolfson (1994).

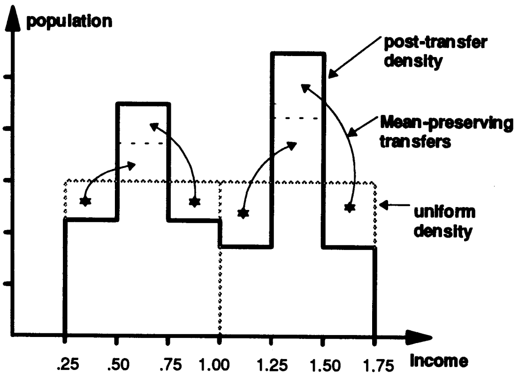

FIGURE 1. POLARIZATION AND INEQUALITY

Sources: R. Love and M. C. Wolfson, 1976; OECD, 1993.

construction, the bimodal distribution is at the same time more polarized and more equal than the uniform distribution from which it was derived.

Polarization and inequality are thus demonstrably different concepts, so what statistics should be used to measure polarization? In the literature on the disappearing middle, in addition to inequality measures, some authors have used quintile income shares, while others have used the fraction of the population in various income ranges defined in terms of the mean or median income. In fact, Figure 1 has been constructed in a particularly nasty way for these kinds of statistics.

Since the distribution is symmetric, the mean is equal to the median, which is 1.0. It can be shown that the *income* share of the middle third of the bimodal distribution is lower than the income share of the middle third of the uniform distribution, while the income share of the middle two-thirds rises in the transition to the bimodal distribution. Thus, the income shares of various middle-quantile groups are *not* necessarily consistent with any sensible formalization of the concept of polarization or disappearing middle. In turn, this means that the large number of papers purporting to analyze the disappearance of the middle class which use inequality indicators such as quintile shares (e.g., Levy, 1987) are simply unable to detect the phenomenon they claim to be studying.

Moreover, the share of the *population* with "middle-level incomes" is similarly perverse in the example of Figure 1, going up or down depending on how "middle" is defined. The population with incomes within 25 percent of the mean = median clearly falls, but the population with incomes within 50 percent of the mean = median rises. Thus, statistics that count the share of the population with "near middle" incomes are also *not* necessarily consistent with a sensible definition of polarization. For example, Thurow (1984) considered the proportion of the population with incomes between 75 percent and 125 percent of the median in his analysis of the disappearing middle class, while McKinley L. Blackburn and David E. Bloom (1985) in a similar analysis focused on the proportion with incomes between 60 percent and 225 percent of the median.

This is an unsatisfactory situation. There is an expanding literature seeking to analyze the phenomena of inequality and the disappearing middle class, accompanied by an incoherent variety of statistical indicators. These analyses are only rarely accompanied by a justifiable sense of unease about what precisely the researchers are measuring. Moreover, I have just shown that perhaps the most basic axiom underlying the formal theory of inequality measurement, the Pigou-Dalton condition of transfers which is formally equivalent to the Lorenz curve criterion, is inconsistent with the concept of polarization, or the roughly equivalent notions of "spreadoutness" from the middle or bimodality that lie at the heart of the disappearing middle-class phenomenon.

In this context, James E. Foster and Wolfson (1992) sought to provide a formalization of these latter concepts to give them the same rigor as for inequality measures. There is space here only to sketch the lines of development. It turns out that there is quite a nice duality or complementarity between polarization and inequality. The strand of development in both cases starts with a cumulative density function (cdf) for the distribution of income. For inequality measures, one can think graphically of one intermediate step to arrive at the Lorenz curve. This step involves "exchanging" the axes of the cumulative density function so that population percentiles are ranged along the horizontal axis and incomes are along the vertical axis. The result is Jan Pen's (1973) "parade of dwarfs (and a few giants)." This "parade" curve (after dividing through by the mean income) is then integrated moving right from the origin to obtain the usual Lorenz curve.

Formalizing the concept of polarization can follow a similar path of graphical transformations of the initial cdf, but with a few key differences. Imagine that, after exchanging the axes of the cdf as above, individuals' incomes along the vertical axis are divided by the median (rather than the mean) income. The resulting median-normalized "parade" is next cut at the midpoint of the horizontal axis, the 50th population percentile. The horizontal axis is then shifted up to touch the curve at this point, which is (by definition) the median income. The portion to the left (i.e., the first half of the parade curve for the 50 percent of the population with incomes below the median) is then rotated around the horizontal axis. The result is a curve looking a bit like a lopsided gull. It shows, for any population percentile along the horizontal axis, how far its income is from the median, thus giving an indication of how "spread out" from the middle (50th percentile) the distribution of income is. A less spread-out distribution (i.e., one with a larger middle class) will have a curve that is lower.

For reasons that are intuitively similar to the notion of second-order stochastic dominance, one can integrate this curve out from the midpoint along the horizontal axis (where by construction the height of the curve is zero) to get what Foster and I (1992) have called the "polarization curve." We have shown that one distribution has an unequivocally smaller middle class (is more spread away from the median) if and only if its polarization curve is everywhere higher. This polarization curve thus plays the same role for the concept of polarization as the Lorenz curve plays for inequality.

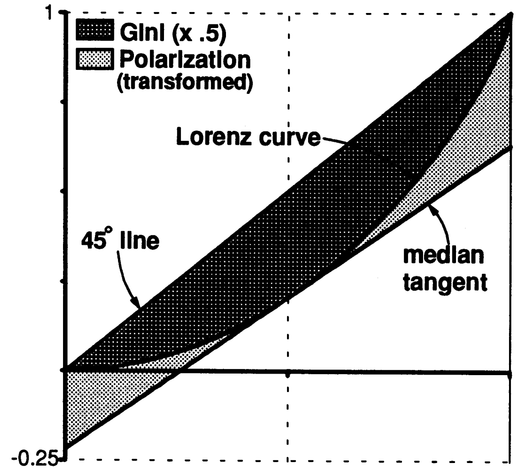

FIGURE 2. A NEW MEASURE OF POLARIZATION BASED ON THE LORENZ CURVE

It can then be shown that the area under this polarization curve is a scalar index of polarization, just as the Gini coefficient, as (twice) the area between the 45-degree line and the Lorenz curve, is a scalar index of inequality. However, as with the Gini and Lorenz curves, it is still possible to have crossing polarization curves. Thus, polarization curves induce only a partial ordering over income-distribution densities with respect to the sizes of their middle classes, while the area under the polarization curve induces a complete ordering.

These two strands of development, with their common starting point in the cumulative income-distribution density function, can now be brought together again in a very nice extension of the Lorenz curve. Figure 2 shows a typical Lorenz curve. The key addition is the tangent line to the Lorenz curve at the 50th population percentile, with the vertical axis extended down to meet this tangent. It turns out that the polarization curve just described is closely related to the Lorenz curve. If the vertical axis of the polarization curve is first renormalized by multiplying by the ratio of the mean to the median, and then the horizontal axis is rotated until it has the same slope as the tangent to the Lorenz curve at the 50th population percentile, the result is that this transformed polarization curve is identical to the Lorenz curve!

In turn, the area under the polarization curve, call it P^*, is a scalar indicator of the extent of polarization or the size of the middle class. The lightly shaded area in Figure 2 between the tangent line and the Lorenz curve is $T - \text{Gini}/2$, so that $P^* = (T - \text{Gini}/2)/\text{mtan}$; where mtan = "median tangent" = m/μ = the slope of the tangent to the Lorenz curve at the 50th population percentile; m = median; μ = mean; and T = the area of the trapezoid defined by the 45-degree line and the median tangent = the vertical distance between the Lorenz curve and the 45-degree line at the 50th percentile = $0.5 - L(0.5)$ = the difference between 50 percent and the income share of the bottom half of the population [which latter, $L(0.5)$, I refer to as the "median share" or mshare]. P^* has a minimum of 0 for a perfectly equal distribution of income, and a value of 0.25 for a perfectly bimodal distribution with half the population at zero income and the other half at 2μ.[1] In order to have an index with a similar range to the Gini (i.e., in the [0, 1] interval if there are no negative incomes), I arbitrarily define the scalar polarization index $P = 4P^* = 2(2T - \text{Gini})/\text{mtan}$.

Figure 2 makes it immediately clear where the conflicts between inequality and polarization arise, and why the concepts have so often been confused. If there is an "equalizing transfer" of income (in the sense of the Pigou-Dalton condition of transfers) from an individual above the median to an individual with income below the median (and the transfer is not so large that it causes either to cross the median), then both inequality and polarization decline: the Lorenz curve moves closer to the 45-degree line as does the tangent line at the 50th population percentile. By virtue of this class of examples, there are clearly many situations where inequality and polarization rankings will agree.

The two concepts will disagree, however, when there are equalizing transfers entirely on one side of the median, exactly as in Figure 1 earlier. In these cases, the median tangent curve is unaffected by the transfer, but the portion of the Lorenz curve on the affected side of the median moves closer to the 45-degree line. Such a shift in the Lorenz curve necessarily reduces the Gini coefficient and correspondingly increases the polarization measure P.

Such a divergence between inequality and polarization could, of course, be merely a theoretical curiosity. An important question is whether in practice one may see divergent trends in the two kinds of attributes of income distributions. The answer is yes, and illustrations are provided shortly.

The demonstration that inequality as formalized is not always in accord with the concept of polarization reopens the question of the axiomatic foundation of inequality measures. Specifically, it raises questions about the Pigou-Dalton condition of transfers. As noted by Yoram Amiel and Frank Cowell (1989 footnote 14), Pigou was doubtful about its validity. Moreover, in their survey of nearly 1,000 undergraduate economics students (most before they had studied this topic), a majority rejected this axiom as part of their concept of inequality. At the very least, this suggests that in order to capture the concerns of the general public, summary measures based on concepts like polarization should be given equal space along with Lorenz-consistent inequality measures when describing trends in income distribution. Indeed, polarization as formalized here may be closer to the general public's vernacular concept of inequality than formal measures of inequality based on Pigou-Dalton-Lorenz-Gini concepts.

Empirical results from a time series of the Canadian Surveys of Consumer Finance

[1]Note that 0.25 is not necessarily the maximum. P^* could exceed 0.25 if half the population had a negative average income. For the same reason, the Gini can exceed 1.0. For any given mshare and mtan both positive, P^* is minimized and approaches 0 for a trimodal distribution where one individual has a very large negative income, another has a very large positive income, and everyone else has the same income in between.

TABLE 1—SELECTED INEQUALITY AND POLARIZATION INDICATORS, CANADA, 1967–1991

Indicator	Percentages				
	1967	1973	1981	1986	1991
Inequality					
Squared CV	58.2	60.9	54.4	62.0	66.7
Gini	36.3	37.9	37.8	39.6	40.3
Exp	44.6	45.2	45.0	45.8	46.1
VL	87.1	78.2	79.0	85.7	86.3
Polarization					
Population share in ranges of median income:					
75–150%	42	37	36	32	32
60–225%	66	64	62	59	59
Range of income/median covering middle population:					
40–60%	33.8	40.0	39.5	43.4	43.7
30–70%	69.8	79.2	81.5	90.9	88.7
20–80%	112.5	125.8	131.3	140.4	141.1
Median share	24.3	22.9	22.6	21.3	21.0
Median/mean	89.7	86.9	88.2	86.0	85.2
Polarization	33.8	37.6	38.5	41.5	41.7

have been used to illustrate the divergence between inequality and polarization. Table 1 presents data for the distribution of labor income for all individuals aged 18–64 with annual labor income of at least 5 percent of the average wage.[2]

Two sets of statistics are given. The first is a set of inequality measures: the top-sensitive (in Cowell's [1977] sense) squared coefficient of variation (CV), the middle-sensitive Gini coefficient, and the bottom-sensitive Exponential measure (Exp).[3] The fourth statistic is the variance of logarithms (VL), which is included only to show that it is *not* an inequality measure.[4]

The second set of statistics is related to the concept of polarization. The first five statistics count the proportion of the population with "middle-class" labor incomes, though from two different perspectives. The first pair, denoted "population share in ranges of median income," give the numbers of individuals with incomes between 75 and 150 percent of the median, and those with incomes between 60 and 225 percent of the median, respectively. These statistics measure the size of the middle class defined in terms of an arbitrary range of median-normalized incomes. The next three statistics effectively exchange the axes of the cdf by defining the middle class in "people space" rather than "income space." These statistics are based on symmetric percentile ranges of the population—within 10, 20, and 30 percent of the 50th percentile, denoted "40–60%," "30–70%," and "20–80%," respectively. For each of these people-space ranges, the corresponding range of incomes they span, divided by the median, is the statistic given. Thus, for example, if the figure for 40–60% population is 33.8 percent (as shown for 1967), this is the 60th-percentile income minus the 40th-percentile income times 100 divided by the median. Even though Figure 1 shows that any one of these statistics may be misleading by itself, agreement among a set is more likely to indicate an unambiguous change in polarization as I have formalized the concept.

[2]The source is special tabulations by the author of the Survey of Consumer Finance internal working files. Labor income includes wages and salaries, military pay and allowances, and self-employment income (which may be negative). The *de minimus* exclusion of those with less than 5 percent of the average wage was based only on employment income. Note that 95 percent confidence intervals estimated for the various measures (taking account of the complex sample design) suggest that at most two digits are statistically significant, and for top-sensitive measures like the squared CV, only one digit is significant (M. S. Kovacevic, 1993).

[3]This measure, introduced in Wolfson (1986), is $\Sigma p_i \exp(-y_i/\mu)$, where p_i is the proportion of the population in the ith income group, y_i is the average income in that group, and μ is the overall mean income. It was introduced precisely because of the advantage over other bottom-sensitive measures like the Theil-Entropy, Theil-Bernouilli (also referred to as the mean logarithmic deviation), and members of the Atkinson (1970) family in that it does not explode with zero or near-zero incomes.

[4]I have included the variance of logarithms only because it continues to be widely used in the literature even though it has been known for decades (Love and Wolfson, 1976; Cowell, 1977) that it is *not* consistent with Lorenz curve rankings: it moves in the wrong direction for transfers above $\sim 2.7\mu$ (often about the 95th percentile). It should be banished from serious inequality analysis.

The last three polarization-related statistics all derive from the polarization/Lorenz curve shown in Figure 2. The first of these is the median share mentioned earlier: the share of income accruing to the bottom half of the population. This in turn is exactly the height of the Lorenz curve halfway along the horizontal axis (i.e., at the 50th percentile). Also, 0.5(median share) is the distance between the 45-degree line and the Lorenz curve at the 50th percentile, hence the area T of the trapezoid enclosing the Lorenz curve in Figure 2. The second statistic is the ratio of the median to the mean income, m/μ. In addition to its graphical interpretation (the slope of the tangent to the Lorenz curve at the 50th percentile), this ratio is also an indicator of the skewness of the distribution. Finally, the last statistic is the polarization measure P defined above. Higher P means more polarization, and a smaller middle class.

Leaving aside the variance of logarithms, measures indicating the concepts of inequality and polarization generally move in the same direction. From 1973 to 1981, however, all of the Lorenz-consistent inequality measures decline or are constant; at the same time, almost all of the polarization measures increase. The divergence between polarization and inequality is not merely a theoretical curiosity; it occurs in practice as well.[5] Research in income distribution and economic inequality is therefore well advised to include measures related to polarization.

REFERENCES

Amiel, Yoram and Cowell, Frank A. "Measurement of Inequality: Experimental Test by Questionnaire." Discussion Paper No. TIDI/140, Suntory-Toyota International Centre for Economics and Related Disciplines, London School of Economics, October 1989.

Atkinson, A. B. "On the Measurement of Inequality." *Journal of Economic Theory*, September 1970, *2*(3), pp. 244–63.

Blackburn, McKinley L. and Bloom, David E. "What is Happening to the Middle Class?" *American Demographics*, January 1985, *7*(1), pp. 18–25.

Cowell, Frank A. *Measuring inequality*. Oxford: Philip Allan, 1977.

Foster, James E. and Wolfson, Michael C. "Polarization and the Decline of the Middle Class: Canada and the U.S." Mimeo, Vanderbilt University, 1992.

Kovacevic, M. S. "Variance Estimation for Measures of Income Inequality and Polarization." Mimeo, Methodology Branch, Statistics Canada, Ottawa, 1993.

Kuttner, B. "The Declining Middle." *Atlantic Monthly*, July 1983, *252*, pp. 60–71.

Levy, Frank. "Changes in the Distribution of American Family Incomes, 1947 to 1984." *Science*, 22 May 1987, *236*, pp. 923–27.

Levy, Frank and Murnane, Richard J. "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." *Journal of Economic Literature*, September 1992, *30*(3), pp. 1333–81.

Love, R. and Wolfson, M. C. *Income inequality: Statistical methodology and Canadian illustrations*, Catalogue 13-559 Occasional. Ottawa: Statistics Canada, March 1976.

OECD. *Employment outlook*. Paris: Organization for Economic Cooperation and Development, 1993.

Pen, Jan. "A Parade of Dwarfs (and a Few Giants)," in A. B. Atkinson, ed., *Wealth, income and inequality*. Middlesex, U.K.: Penguin, 1973, pp. 73–82.

Thurow, Lester. "The Disappearance of the Middle Class." *New York Times*, 5 February 1984, section 3, p. 2.

Wolfson, Michael C. "Stasis Amid Change—Income Inequality in Canada 1965–1983." *Review of Income and Wealth*, December 1986, *32*(4), pp. 337–69.

______. "Divergent Inequalities—Theory, Empirical Results and Prescriptions." Mimeo, Analytical Studies Branch Research Papers, Statistics Canada, and Economic Growth and Technical Change Working Paper, Canadian Institute for Advanced Research, Toronto, 1994.

[5]The variance of logs is also clearly inconsistent with Lorenz inequality in the changes from 1967 to 1973.