USING THE P90/P10 INDEX TO MEASURE U.S. INEQUALITY TRENDS WITH CURRENT POPULATION SURVEY DATA: A VIEW FROM INSIDE THE CENSUS BUREAU VAULTS

by Richard V. Burkhauser*

Cornell University

Shuaizhang Feng

Shanghai University of Finance and Economics

and

Stephen P. Jenkins

University of Essex

The March Current Population Survey (CPS) is the primary data source for estimation of levels and trends in U.S. earnings and income inequality. However, time-inconsistency problems related to top coding lead many CPS users to measure inequality via the ratio of the 90th to the 10th percentile (P90/P10) rather than by more traditional summary measures. With access to public use and restricted-access internal CPS data, and by applying bounding methods, we show that using P90/P10 does not completely obviate time-inconsistency problems, especially in capturing household income inequality trends. Using internal data, we create consistent cell mean values for all top-coded public use values that, when used with public use data, closely track inequality trends in earnings and household income using internal data. But estimates of longer-term inequality trends with these corrected data based on P90/P10 differ from those based on the Gini coefficient. The choice of inequality measure still matters.

1. Introduction

The vast majority of the studies of trends in U.S. labor earnings and income inequality since the 1970s have been based on public use files of the March Current Population Survey (CPS). Yet time-inconsistency problems related to top coding in these data and in other national data sets lead many researchers to use the ratio

of the 90th and 10th percentile of a distribution (P90/P10) rather than a more traditional summary measure of inequality such as the Gini coefficient, Theil index, or coefficient of variation, each of which uses information about all income values, rather than only two. In the U.S. labor economics literature, P90/P10 is the most commonly used measure of wage or labor earnings dispersion; see, e.g. Juhn *et al.* (1993), Danziger and Gottschalk (1993), DiNardo *et al.* (1996), Gottschalk and Smeeding (1997), Gottschalk and Joyce (1998), Katz and Autor (1999), Blau and Kahn (2005), Lemieux (2006), Pencavel (2006), Autor *et al.* (2008), and Lemieux (2008). In the U.S. income inequality literature, the P90/P10 is also a standard measure of inequality in the distributions of size-adjusted family or household income; see, e.g. Danziger and Gottschalk (1993), Gottschalk and Smeeding (1997), Gottschalk and Danziger (2005), and Daly and Valletta (2006).

In the cross-national comparative literature, CPS data are also commonly used to compare both labor earnings and income inequality levels and trends in the U.S. with other industrialized countries. See Smeeding (2004) for a review of this literature using the CPS. Other recent examples include Alderson *et al.* (2005), Prus and Brown (2006), Atkinson (2007), Brandolini (2007) and Burkhauser *et al.* (2008). The most important source of standardized cross-sectional micro data on industrialized countries—the Luxembourg Income Study (LIS)—uses the public use version of the CPS data for the U.S. On its website (http://www.lisproject.org/keyfigures/ineqtable.htm), LIS (2008) summarizes income inequality using P90/P10 and Gini coefficient estimates that do not adjust for the top coding issues discussed below. Nor does LIS (2008) provide its users with information on the degree that top coding issues of this type are found in its data from other countries. The public use CPS data as developed by LIS is also a major source of information about U.S. inequality in the United Nation University–World Income Inequality Database (UNU-WIDER, 2008). UNU-WIDER (2008) also does not alert its users of top coding issues on its website (http://www.wider.unu.edu/research/Database/en_GB/database).[1]

Other things being equal, any of the traditional summary measures of inequality are likely to be better measures of the inequality of the entire distribution, and hence of its trends over time, than P90/P10 which only captures two points in that distribution. But other things are often not equal. The public use March CPS is the best source of annual information on trends in the labor earnings and income of U.S. households available to the research community. However, all sources of income in the public use CPS are top coded, which makes accurate calculations of traditional summary measures of the distribution impossible and comparisons of these values over time difficult (Levy and Murnane, 1992; Gottschalk and Smeeding, 1997). Moreover, even the internal CPS data, which are not subject to top coding, have been censored to various degrees over time (Welniak, 2003).

---

[1]The U.S. Census Bureau is not alone in imposing topcodes on its income data. Issues of confidentiality make top coding of income values a necessity for most national data sets. For instance, in the U.S., the National Longitudinal Survey of Youth (NLSY97) top codes some of its income sources as does the Panel Study of Income Dynamics (PSID). In Great Britain, in order to comply with the new Statistics and Registration Services Act of 2007, the Annual Population Survey and the Quarterly Labour Force Survey have introduced top codes on earnings data in their main public release files.

The impact of censoring on Gini coefficients estimated with both the public use and internal CPS data has been documented in previous research; see, e.g. Burkhauser *et al.* (2003–2004, 2004) and Feng *et al.* (2006). But no similar scrutiny has been given to the impact of censoring on quantile ratio measures such as P90/P10. Researchers have implicitly assumed that P90/P10 is not affected by censoring, reasoning that the fraction of observations affected by censoring of total wages and salaries, labor earnings or income is less than 10 percent. While this is true, in the CPS data, censoring takes place at the level of each income source, not for income totals, so some values below the 90th percentile of total labor earnings and especially the 90th percentile of the household income distribution are censored. As a result, even what are apparently modest amounts of censoring in the population as a whole may affect estimates of P90/P10.

To address the issues raised by censoring requires use of internal March CPS data, and we have been able to gain access to them for this purpose.[2] Our analysis considers data for income years 1975–2004. We examine three distributions of income that are commonly assessed in the labor and income inequality literatures: (i) wages and salaries income among individuals working full-time full-year for wages; (ii) total earnings income among full-time, full-year workers (wage and salaries plus farm and non-farm self-employment earnings); and (iii) household income among all individuals.

Our paper makes three contributions. First, using innovative bounding methods, we show that calculating P90/P10 with public use CPS data—even when Census Bureau cell means are used for top coded values—does not completely obviate the problem of time-inconsistency, especially for those interested in trends in the inequality of individuals' size-adjusted household income. Second, we offer a means by which researchers may reduce problems caused by censoring. With access to the internal CPS data, we were able to create consistent cell mean values for all top-coded values in all years of internal data made available to us (1975–2004) that offer a plausible correction for time inconsistency problems in the public use CPS data when integrated with them. These cell means were cleared for distribution to the public by the Census Bureau and are available in Larrimore *et al.* (forthcoming) along with a fuller discussion of how they were created.

Our third contribution concerns the assessment of longer-term U.S. inequality trends. Comparing estimates of P90/P10 based on our adjusted public use CPS data with estimates of Gini coefficients based on either the internal or public use CPS data consistently top-coded to control for time inconsistencies, we find our P90/P10 trends differ significantly from trends in either of the two Gini coefficient series. Hence, researchers should be cautious in making inference about trends in the inequality of the distributions of wages and salaries income, labor earnings income, or size-adjusted household income over the last three decades based on changes in the relative position of only two points in each of those distributions.

[2]To gain access to the internal CPS data, two of us (Burkhauser and Feng) became Special Sworn Status researchers of the U.S. Census Bureau at the New York Census Research Data Center, Cornell University, in 2005.

2. Censoring Problems in the Current Population Survey

The CPS, based on a large representative sample of the U.S. population, interviews about 57,000 households each month. Each March, the CPS collects detailed information about each source of income in the previous year for every household member. To protect the confidentiality of respondents, top codes are imposed on all sources of income above a specific value. Less well known to the research community is that even the internal data that the Census Bureau uses to calculate various official statistics, including inequality measures, are also subject to censoring. In earlier years this was primarily because of restrictions on computer tape space. Although such constraints are substantially relaxed nowadays, CPS internal income data are still censored for various Census Bureau considerations, including minimizing the possible impact of recording (keying) errors, helping to maintain respondents' confidentiality, and preventing volatility and distortion of annual statistics (Welniak, 2003; Feng *et al.*, 2006).

The precise Census Bureau variable names, and definitions of the three sources of income that we analyze, and how they have changed over time, are shown in Table 1. For income years 1975–86, the Census Bureau reported three sources of labor earnings and eight other sources of income. From 1987 onwards they have used a finer categorization, reporting four sources of labor earnings and 20 other sources of income. For all income components, both the internal and public use CPS censoring points have changed over time. Public use CPS censoring points are shown for income years 1975–86 and 1987–2004 in appendix tables 1 and 2 of Burkhauser *et al.* (2007). Corresponding internal CPS censoring points for the two periods are provided in appendix tables 3 and 4 of Burkhauser *et al.* (2007).

Because censored values start at different points in the distribution each year, any inequality estimate not taking account of this variation is time-inconsistent. This includes Census Bureau published estimates using internal CPS data. Past researchers have recognized this problem and, for the most part, used some rule-of-thumb adjustment procedure to control for it; see, e.g. Juhn *et al.* (1993) and Trejo (1997). More recently, Burkhauser *et al.* (2003–2004, 2004) consistently top code values for each source of income in the CPS. They do so for each source of income by choosing the highest point in its income distribution available for all years. In so doing they are able to estimate Gini coefficients that, while lower in level, captured the long-term trends in inequality relatively well. They argue that their Gini coefficient estimates from the public use CPS data better captured long-term trends in labor earnings for this population than even Census Bureau estimates based on uncorrected internal CPS data.[3]

3. Bounds for P90/P10 and Seven Series of Estimates

In this section, first we describe our method for putting bounds on estimates of P90/P10 from censored data series, and second, we define seven series of CPS-based estimates that arise from application of the methods and from ignoring censoring.

[3]For examples of the use of consistent top coding to control for time inconsistency in the public use CPS data, see Karoly and Burtless (1995), Burkhauser *et al.* (2003–2004, 2004), Gottschalk and Danziger (2005), Feng *et al.* (2006), and Burkhauser *et al.* (2008).

TABLE 1

Income Items Reported in the Current Population Survey

Name	Name in Public Files	Name in Internal Files	Definition
1975–86			
Labor earnings			
INCWAG	I51A	WSAL_VAL	Wages and salaries
INCSE	I51B	SEMP_VAL	Self employment income
INCFRM	I51C	FRSE_VAL	Farm income
Other sources			
INCSS	I52A	I52A_VAL	Income from Social Security and/or Railroad Retirement
INCSEC	I52B	SSI_VAL	Supplemental Security Income
INCPA	I53A	PAW_VAL	Public Assistance
INCINT	I53B	INT_VAL	Interest
INCDIV	I53C	I53C_VAL	Dividends, rentals, trust income
INCOMP	I53D	I53D_VAL	Veteran's, unemployment, worker's compensation
INCRET	I53E	I53E_VAL	Pension income
INCALC	I53F	I53F_VAL	Alimony, child support, other income
1987–2004			
Labor earnings			
INCER	ERN_VAL	ERN_VAL	Primary earnings
INCWG1	WS_VAL	WS_VAL	Wages and salaries—second source
INCSE1	SE_VAL	SE_VAL	Self employment income—second source
INCFR1	FRM_VAL	FRM_VAL	Farm income—second source
Other sources			
INCSS	SS_VAL	SS_VAL	Social Security Income
INCSEC	SSI_VAL	SSI_VAL	Supplemental Security Income
INCPA	PAW_VAL	PAW_VAL	Public Assistance and Welfare Income
INCINT	INT_VAL	INT_VAL	Interest
INCDV2	DIV_VAL	DIV_VAL	Dividends
INCRNT	RNT_VAL	RNT_VAL	Rental income
INCALM	ALM_VAL	ALM_VAL	Alimony income
INCHLD	CSP_VAL	CSP_VAL	Child Support Income
INCUC	UC_VAL	UC_VAL	Unemployment income
INCWCP	WC_VAL	WC_VAL	Worker's compensation income
INCVET	VET_VAL	VET_VAL	Veteran's benefits
INCRT1	RET_VAL1	RET_VAL1	Retirement income—source 1
INCRT2	RET_VAL2	RET_VAL2	Retirement income—source 2
INCSI1	SUR_VAL1	SUR_VAL1	Survivor's income—source 1
INCSI2	SUR_VAL2	SUR_VAL2	Survivor's income—source 2
INCDS1	DIS_VAL1	DIS_VAL1	Disability income—source 1
INCDS2	DIS_VAL2	DIS_VAL2	Disability income—source 2
INCED	ED_VAL	ED_VAL	Education assistance
INCONT	FIN_VAL	FIN_VAL	Financial assistance
INCOTH	OI_VAL	OI_VAL	Other income

3.1. *Bounds on Estimates of P90/P10 from Top Coded Data*

Let the true income distribution be denoted by the random variable x, which has a cumulative distribution function $F(x)$. The p-th population income quantile ξ_p is defined by:

$$(1) \qquad p = F(\xi_p) = \Pr(x \leq \xi_p), \quad 0 \leq p \leq 1.$$

Suppose we have a random sample s comprising N income units, with the distribution of their incomes described by the vector $x = \{x_1, x_2, x_3, \ldots, x_N\}$. The sample estimate of the p-th quantile of the distribution is:

$$(2) \qquad \hat{\xi}_p = \sup\{x_i \in s | \hat{F}(x_i) \le p\},$$

derived by solving the equation $p = \hat{F}(\xi_p)$, where the sample estimate of the cumulative distribution function for x is:

$$(3) \qquad \hat{F}(x) = \sum_s w_i I(x_i \le x)\big/\hat{N}, \quad \text{with } \hat{N} = \sum_s w_i.$$

$I(.)$ is the indicator function and the sample weight for unit i is w_i.

The problem for researchers is that x is not fully observed. Top coding (or right censoring in general) means that some incomes at the top of the income distribution are not observed. Instead, two other vectors are observed in the sample data: censored incomes $y = \{y_1, y_2, y_3, \ldots, y_N\}$ and censoring indicators $c = \{c_1, c_2, c_3, \ldots, c_N\}$, with $y_i = x_i$ if $c_i = 0$ and $y_i < x_i$ if $c_i = 1$, for each $i = 1, \ldots, N$. In addition, because we are trying to model incomes that are aggregates of several income sources, but censoring occurs at the level of each individual income source, some lower-valued incomes might be censored while higher-valued ones are not censored.

The sample estimate of the proportion of censored observations is $\hat{\theta}$ where:

$$(4) \qquad \hat{\theta} = \sum_s w_i I(c_i = 1)\big/\hat{N}.$$

Although income values may be censored, we can place lower and upper bounds on the quantiles that we are trying to estimate (Manski, 1994). The lower bound is derived from distribution y, assuming that the true (unobserved) value of each censored observation is equal to the observed censored value. The upper bound is derived by assuming that the true income value of each censored observation is equal to positive infinity, i.e. estimated from a distribution $z = \{z_1, z_2, z_3, \ldots, z_N\}$, with $z_i = x_i$ if $c_i = 0$ and $z_i = +\infty$ if $c_i = 1$, for each $i = 1, \ldots, N$. In general, the ranking by income of units differs between distributions y and z and hence lower and upper bound estimates of the quantiles of the true distribution differ.

More formally, the estimate of the lower bound is:

$$(5) \qquad \hat{\xi}_p^L = \sup\{y_i \in s | \hat{F}_y(y_i) \in p\},$$

where the empirical cumulative distribution function (CDF) of the censored distribution y is:

$$(6) \qquad \hat{F}_y(y) = \sum_s w_i I(y_i \le y)\big/\hat{N}.$$

The estimate of the upper bound is:

$$(7) \qquad \hat{\xi}_p^U = \sup\{z_i \in s | \hat{F}_z(z_i) \le p\},$$

where the empirical CDF of the distribution z is:

$$(8) \qquad \hat{F}_z(z) = \sum_s w_i I(z_i \leq z) \big/ \hat{N}.$$

It is straightforward to show that $\hat{\xi}_p^L \leq \hat{\xi}_p \leq \hat{\xi}_p^U$ for $0 \leq p \leq 1$, because $y_i \leq x_i \leq z_i$ for each $i = 1, \ldots, N$. Moreover, when $p \leq 1 - \hat{\theta}$, the upper and lower bounds are both informative. If, instead, $p > 1 - \hat{\theta}$, censoring bites: the p-th quantile lies within the censored income range. In this case, the lower bound estimate of the p-th quantile derived from y remains well-defined, but the upper bound estimate is uninformative—it is infinity.

To illustrate how the lower and upper bounds of order statistics such as quantiles are derived, we give a simple numerical example. Assume the distribution of observed incomes is {2,000, 1,000, 4,000, 5,000} and the first income is censored. (Recall that a censored value need not be the maximum value observed in sample data.) Suppose the aim is to estimate the income corresponding to the upper quartile (the income of the second highest earner in this simple case). Only one income is censored, and so we have the case corresponding to $p \leq 1 - \theta$. The lower bound estimate of the upper quartile is 4,000, and the upper bound estimate is 5,000. Now suppose instead that income 4,000 is also censored. This takes us to the case $p > 1 - \theta$. The lower bound estimate of the upper quartile is again 4,000, but the upper bound estimate is uninformative.

If the total income for any income-recipient unit (e.g. a household) is the aggregate of incomes across individuals belonging to the same unit, the same estimation methods may be applied. The greater the aggregation across income sources, or across individuals, the further down the distribution of total income that censoring is likely to occur. There is a range of top coded values interspersed along the range of non-top coded values. This dispersion means that the adjustment for top coding in the CPS proposed by Fichtenbaum and Shahidi (1988) for estimation of the Gini coefficient, based on fitting a Pareto distribution to incomes above a single critical value, is not practical in the current context.

3.2. *Seven Series of P90/P10 Estimates*

Using these bounding methods, we calculate lower and upper bound estimates for P90/P10 based on public use CPS data files, which we will call the *Public-Lower* and *Public-Upper* series respectively. Because we have access to the internal CPS data files, we are also able to calculate *Internal-Lower* and *Internal-Upper* series of P90/P10 estimates from the internal CPS data in a similar way. Because internal data contain more information than public use data (the internal censoring point is greater than or equal to the public use censoring point), the *Public-Lower* estimates will be less than or equal to the corresponding *Internal-Lower* estimates and the *Public-Upper* estimates will be greater than or equal to the corresponding *Internal-Upper* estimates.

We also calculate three other P90/P10 series from the CPS for comparison purposes. The first, labeled *Public*, is calculated from public use files using the top coded value assigned by the Census Bureau to individuals' sources of income for all years. For each income year before 1995, estimates are the same as

Public-Lower estimates for the same year. They are greater thereafter because, from income year 1995 onwards, the Census Bureau assigned an estimated cell mean to each top coded value based on the person's characteristics rather than the top code cutoff value. For these years, because the *Public* series is based on a distribution in which income values are more accurately observed than in the distribution including top coded values, it should yield P90/P10 estimates closer to those based on internal data.

The second additional series, labeled *Cell-Mean*, assigns a cell-mean that we consistently calculate over all the years of internal data available to us (1975–2004) for each person top coded. Because we were given permission to use the internal data, we were able to construct a data file similar to the one discussed below that the Census Bureau has, since 1995, used to assign cell means to top coded values in the public use files. For the same reasons discussed above, the P90/P10 estimates in this series should more closely track the estimates derived from the internal data in all years.

In income year 1995, the Census Bureau began providing cell mean values rather than the top coded cutoff value for wages and salaries, self-employment earnings, and farm earnings from sex/race/work experience cells. That is, rather than reporting the top code cutoff value, the public use file reports the average value for those with the same sex/race/work experience characteristics with values above the top code cutoff point. In income year 1998, the Census Bureau extended its provision of cell means to other non-governmental sources of income. However, to date the Census Bureau has not provided such cell means for earlier years. Hence for reasons of consistency, researchers interested in comparing trends in labor earnings or income before 1995 with those after 1995 are not able to take advantage of the cell mean option available in the public use data. However, using our access to the internal data, we were able to create a consistent set of cell mean values for each income source for every person in the public use files for income years 1975–2004.[4]

The third additional series, labeled *Rule-of-Thumb*, assigns a value of 150 percent of the top code cutoff value to all top coded values. This popular rule-of-thumb approach to assigning top code values has been used in the labor economics literature by Katz and Murphy (1992), Lemieux (2006) and Autor *et al.* (2008).

4. Trends in Wages and Labor Earnings Inequality for Full-Time, Full-Year, Workers

Seven series of P90/P10 estimates were calculated for the distribution of wages and salaries of full-time, full-year workers, the most typical definition of labor earnings and of a worker in the labor economics literature tracking the inequality

---

[4]For every income source, we calculate a single mean value for all top coded values. But we do not provide cell-means for subgroups of the population defined by, for example, sex, race, and experience. In contrast, the Census Bureau provides cell means based on sex/race/work experience cells for labor earnings but only single cell means for non-governmental sources of non-labor incomes and they do not provide cell means at all for governmental sources of non-labor income. In addition, our series provides consistent cell-mean values for earlier years, something the Census Bureau has not provided to the research community yet. See Larrimore *et al.* (forthcoming) for a fuller discussion of how they were created. It also contains the full set of cell means.

TABLE 2

P90/P10 Estimates for Wages and Salaries of Full-Time Full-Year Workers

Income Year	Public-Lower	Public-Upper	Public	Rule-of-Thumb	Cell-Mean	Internal-Lower	Internal-Upper
1975	4.00	4.00	4.00	4.00	4.00	4.00	4.00
1976	3.97	3.97	3.97	3.97	3.97	3.97	3.97
1977	4.08	4.08	4.08	4.08	4.08	4.08	4.08
1978	4.17	4.17	4.17	4.17	4.17	4.17	4.17
1979	3.97	3.97	3.97	3.97	3.97	3.97	3.97
1980	4.12	4.12	4.12	4.12	4.12	4.12	4.12
1981	4.00	4.00	4.00	4.00	4.00	4.00	4.00
1982	4.33	4.33	4.33	4.33	4.33	4.33	4.33
1983	4.29	4.29	4.29	4.29	4.29	4.29	4.29
1984	4.36	4.36	4.36	4.36	4.36	4.36	4.36
1985	4.44	4.44	4.44	4.44	4.44	4.44	4.44
1986	4.49	4.49	4.49	4.49	4.49	4.49	4.49
1987	4.40	4.40	4.40	4.40	4.40	4.40	4.40
1988	4.50	4.50	4.50	4.50	4.50	4.50	4.50
1989	4.66	4.66	4.66	4.66	4.66	4.66	4.66
1990	4.55	4.55	4.55	4.55	4.55	4.55	4.55
1991	4.57	4.57	4.57	4.57	4.57	4.57	4.57
1992	4.50	4.50	4.50	4.50	4.50	4.50	4.50
1993	4.58	4.58	4.58	4.58	4.58	4.58	4.58
1994	4.92	4.92	4.92	4.92	4.92	4.92	4.92
1995	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1996	4.81	4.89	4.85	4.81	4.89	4.81	4.81
1997	4.92	5.00	5.00	5.00	5.00	5.00	5.00
1998	4.81	5.00	4.96	4.89	4.96	4.89	4.89
1999	5.00	5.07	5.00	5.00	5.00	5.00	5.00
2000	4.90	5.03	5.03	5.00	5.03	5.00	5.00
2001	5.00	5.13	5.07	5.00	5.07	5.00	5.00
2002	5.33	5.33	5.33	5.33	5.33	5.33	5.33
2003	5.13	5.26	5.19	5.13	5.26	5.13	5.13
2004	5.19	5.25	5.25	5.25	5.25	5.22	5.22

Note: The definitions of the series are provided in the main text.

of labor earnings[5] (see Table 2). The first five columns provide estimates based on public use CPS data including our now publically available cell means series. The last two columns are derived from internal CPS data. We show below that, although censoring is a potential problem in estimating inequality trends for wage and salary income of this population, it is not a very important one, since the internal data has no censoring problem and the public use data has only a small potential problem.

Prior to income year 1987, wages and salaries income came from only one source (INCWAG) (see Table 1). Hence top coding was not a problem since none of the workers with wage and salary top codes in these years had incomes below the 90th percentile of the wage and salary distribution. Since then, the 90th

[5]For our analysis of full-time, full year workers' income from wage and salaries, we excluded individuals who had non-positive income from wage and salaries or whose primary source of labor earnings income was farm income or non-farm self-employment income. For our analysis of full-time, full year workers' income from labor earnings and our analysis of all individuals' size-adjusted household income, we allow non-positive values for specific income sources but assign a value of $1 if the sum of all these income sources is non-positive to avoid including negative incomes in the any of our calculations of labor earnings or size-adjusted household income.

percentile value could be affected by top coding, at least in principle, since the Census Bureau began reporting wage and salary income from two sources, one primary (INCER) and one secondary (INCWG1). Hence it is possible that workers below the 90th percentile of the distribution of wages and salaries formed by the sum of these two sources could be top coded in one of them. As columns 1 and 2 of Table 2 show, top coding is not a problem for estimation of P90/P10 for any income year prior to 1987 and is only a potential problem after 1995—where *Public-Lower* does not equal *Public-Upper*. And in none of these years is the difference between these two values very large.

Columns 6 and 7 of Table 2 show that the internal CPS data provide accurate P90/P10 estimates for all years since the *Internal-Lower* (column 6) values equal *Internal-Upper* values (Column 7) in all years and are, in fact, the same value as reported in columns 1 and 2 in all the years prior to 1996. Hence with respect to wage and salaries, P90/P10 estimates are relatively free of top coding problems. This pattern of no difference in values prior to 1996 and only small differences thereafter with the internal values holds for all of the other series in Table 2. The *Public* series values (column 3) and the *Cell-Mean* series values (column 5) are almost identical. This is the case prior to 1995 because top coding was not a problem for estimation of P90/P10 from either the internal or public use data, so not correcting for top coding by adjusting the cell means in the *Public* series in these years does not matter. Thereafter our consistently measured *Cell-Mean* series is so close to the *Public* series that there is almost no difference. Both the *Public* and *Cell-Mean* series are slightly higher than the *Internal-Upper* one in most years since 1995, showing that using either of the cell mean adjustments slightly overestimates values derived from internal data. In contrast, the *Rule-of-Thumb* series yields virtually the same P90/P10 estimates as the series based on internal data.

Table 2 confirms that, whereas in theory top coding could affect both internal and public use P90/P10 estimates for wages and salaries income, in practice it has no effect on P90/P10 estimates from internal data and only a minor effect on estimates from public use data after 1995. The table also suggests that the rule-of-thumb method common in the wage and salaries literature is at least as effective as using cell means to control for the effects of inconsistent top coding.

Table 3 reports trends in P90/P10 for the distribution of the total earnings of full-time, full-year workers, for each of the seven series. Prior to income year 1987, the Census Bureau summed income from three different sources to create the total earnings variable: wages and salaries (INCWAG), self-employment earnings (INCSE), and farm earnings (INCFRM). Since then, four sources have been combined: primary earnings (INCER), second wages and salaries (INCWG1), secondary self-employment earnings (INCSE1), and secondary farm earnings (INCFR1). As in Table 2, censoring does not matter for any year prior to 1987 or for the years up to 1996 in the public use data (columns 1 and 2). However, in more recent years, top coding has become more of a potential problem. But even in these years, the differences between the series of estimates are small. In the years for which we have access to the internal files, censoring has not been a problem, with *Internal-Lower* estimates equaling *Internal-Upper* estimates in all years (columns 6 and 7). Once again, the *Public* estimates (column 3) and the *Cell-Mean* estimates (column 5) produce series that differ little after 1995 because they use a similar cell

TABLE 3
P90/P10 Estimates for Total Earnings of Full-Time Full-Year Workers

Income Year	Public-Lower	Public-Upper	Public	Rule-of-Thumb	Cell-Mean	Internal-Lower	Internal-Upper
1975	4.27	4.27	4.27	4.27	4.27	4.27	4.27
1976	4.40	4.40	4.40	4.40	4.40	4.40	4.40
1977	4.63	4.63	4.63	4.63	4.63	4.63	4.63
1978	4.18	4.18	4.18	4.18	4.18	4.18	4.18
1979	4.45	4.45	4.45	4.45	4.45	4.45	4.45
1980	4.29	4.29	4.29	4.29	4.29	4.29	4.29
1981	4.58	4.58	4.58	4.58	4.58	4.58	4.58
1982	4.61	4.61	4.61	4.61	4.61	4.61	4.61
1983	4.65	4.65	4.65	4.65	4.65	4.65	4.65
1984	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1985	4.72	4.72	4.72	4.72	4.72	4.72	4.72
1986	4.86	4.86	4.86	4.86	4.86	4.86	4.86
1987	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1988	4.84	4.84	4.84	4.84	4.84	4.84	4.84
1989	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1990	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1991	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1992	4.92	4.92	4.92	4.92	4.92	4.92	4.92
1993	5.09	5.09	5.09	5.09	5.09	5.09	5.09
1994	5.45	5.45	5.45	5.45	5.45	5.45	5.45
1995	5.00	5.00	5.00	5.00	5.00	5.00	5.00
1996	5.17	5.19	5.17	5.17	5.19	5.17	5.17
1997	5.32	5.32	5.32	5.32	5.32	5.32	5.32
1998	5.20	5.38	5.31	5.23	5.31	5.23	5.23
1999	5.38	5.54	5.54	5.46	5.54	5.46	5.46
2000	5.36	5.36	5.36	5.36	5.36	5.36	5.36
2001	5.04	5.30	5.24	5.10	5.24	5.17	5.17
2002	5.33	5.33	5.33	5.33	5.33	5.33	5.33
2003	5.47	5.53	5.52	5.47	5.53	5.47	5.47
2004	5.55	5.67	5.67	5.60	5.66	5.60	5.60

Note: As for Table 2.

mean strategy and are the same prior to 1995 because top coding problems in the data do not affect estimation of P90/P10. Both slightly overestimate the values found in the internal data series. The *Rule-of-Thumb* series, already available to the public, yields virtually the same P90/P10 estimates as the internal series. But, once again, because top coding of total earnings in both the public use and internal data is a relatively small problem for estimation of P90/P10, any of these methods of controlling for top coding in the public use CPS results in plausible approximations of the internal CPS series.

5. Trends in Size-Adjusted Household Income Inequality for Individuals

P90/P10 estimates for the distribution of size-adjusted household income of individuals are reported in Table 4 for all seven series.[6] There are far more sources

[6]We follow common conventions in the household income inequality literature by assuming that household resources are equally shared among all members and by capturing the economies of scale in their consumption of available resources using the "square root" equivalence scale. We suppose that $Y = X/M^{0.5}$, where X is unadjusted total household income, M is the number of individuals in the household, and Y is the adjusted household income. See, e.g. Atkinson *et al.* (1995), Karoly and Burtless (1995), and Burkhauser *et al.* (2003–2004, 2004).

TABLE 4

P90/P10 Estimates for the Size-Adjusted Household Income of Individuals

Income Year	Public-Lower	Public-Upper	Public	Rule-of-Thumb	Cell-Mean	Internal-Lower	Internal-Upper
1975	6.15	6.15	6.15	6.15	6.15	6.15	6.15
1976	6.11	6.11	6.11	6.11	6.11	6.11	6.11
1977	6.23	6.24	6.23	6.23	6.23	6.23	6.24
1978	6.32	6.35	6.32	6.34	6.33	6.33	6.34
1979	6.38	6.44	6.38	6.42	6.41	6.41	6.41
1980	6.61	6.74	6.61	6.71	6.68	6.66	6.66
1981	6.84	6.84	6.84	6.84	6.84	6.84	6.84
1982	7.52	7.53	7.52	7.53	7.52	7.52	7.52
1983	7.59	7.60	7.59	7.59	7.63	7.63	7.63
1984	7.62	7.62	7.62	7.62	7.62	7.62	7.62
1985	7.67	7.68	7.67	7.67	7.67	7.67	7.68
1986	7.84	7.85	7.84	7.84	7.84	7.84	7.84
1987	7.86	7.87	7.86	7.87	7.87	7.87	7.87
1988	7.90	7.91	7.90	7.91	7.91	7.91	7.91
1989	7.70	7.75	7.70	7.74	7.74	7.73	7.75
1990	7.76	7.80	7.76	7.80	7.79	7.78	7.78
1991	7.94	8.01	7.94	8.00	8.00	7.98	8.00
1992	8.15	8.25	8.15	8.24	8.24	8.21	8.22
1993	8.55	8.69	8.55	8.65	8.65	8.62	8.62
1994	8.26	8.53	8.26	8.48	8.47	8.41	8.44
1995	8.01	8.21	8.10	8.07	8.10	8.06	8.09
1996	8.10	8.28	8.17	8.15	8.19	8.16	8.19
1997	8.23	8.48	8.32	8.28	8.33	8.29	8.31
1998	7.98	8.75	8.26	8.15	8.26	8.18	8.22
1999	7.74	8.68	8.05	7.91	8.05	7.96	7.98
2000	7.67	8.59	7.96	7.87	7.96	7.91	7.93
2001	7.78	8.80	8.07	7.96	8.08	8.02	8.04
2002	7.96	8.62	8.12	8.08	8.12	8.10	8.12
2003	8.26	9.04	8.49	8.40	8.50	8.43	8.47
2004	8.24	9.14	8.43	8.35	8.44	8.41	8.44

Notes: As for Table 2. Also, for year 1983, interest incomes are reported differently in the public and internal data files. The results reported here use numbers from the internal data file.

of household income than for total labor earnings and, because household income is assumed to be shared, the size-adjusted household income of each household member depends on the income sources of every household member. Thus censoring is likely to be a more serious problem in this literature than was the case for income from wages and salaries or from total labor earnings. Prior to 1987, 11 sources of income were reported, and the number has increased to 24 since then (see Table 1). As Table 4 shows, P90/P10 estimates derived from the public CPS data are affected by top coding problems although, prior to the 1990s, the gap between the *Public-Lower* and *Public-Upper* series is small. But the gap between the two series has risen steadily since then and especially since 1998. This is clear from Figure 1 which shows the *Public-Lower* and *Public-Upper* series.

A clue to the source of the divergence between the *Public-Lower* and *Public-Upper* series is provided by Figure 2. The top line shows, for each year, the percentage of all individuals affected by top codes in the public use CPS file. This percentage increased steadily in the early 1990s, declined a little in the mid 1990s, and then rose sharply after 1996. This is not a problem as long as censoring only
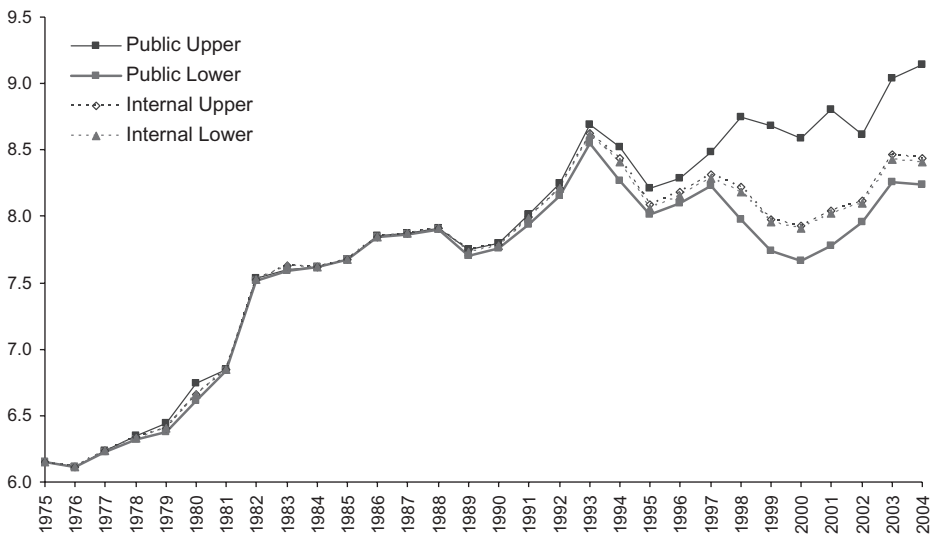
Figure 1. P90/P10 Estimates for Size-Adjusted Household Income of Individuals, by Year


Figure 2. Percentage of Individuals with Size-Adjusted Household Income Censored in the Public Use CPS File

occurs for individuals whose size-adjusted household income is above the 90th percentile of the distribution. Thus in Figure 2 we also show the percentage of all individuals who had observed size-adjusted household incomes less than the 95th percentile and whose income was affected by top coding, together with corresponding percentages for those with incomes below the 90th and 85th percentiles. Individuals with incomes below the 90th percentile began to be affected by top
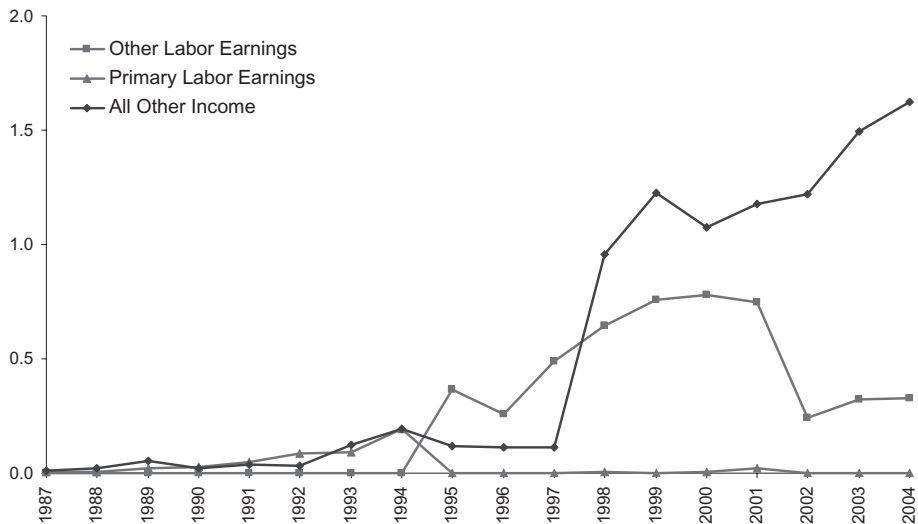
Figure 3. Percentage of Individuals with Censored Size-Adjusted Household Income Below the 90th Percentile by Income Source

coding in the early 1990s and have been more sharply affected since 1998. Note that measuring inequality in terms of the ratio of the 85th percentile to the 10th percentile rather than P90/P10 would reduce this problem somewhat but would not resolve it.

Figure 3 focuses on the post-1987 period and shows the percentage of top coded values below the 90th percentile by income source: primary labor earnings, other labor earnings, and all other income. Figure 3 shows that the jump in the gap between *Public-Lower* and *Public-Upper* estimates was primarily driven by the sharp increase in the fraction of individuals below the 90th percentile whose non-labor earnings was top coded, which rose from 0.1 percent in 1997 to 1.0 percent in 1998 and increased to 1.6 percent by 2004. In income year 1998 (corresponding to CPS survey year 1999), when the Census Bureau started to top code all non-governmental sources of non-labor income items, there was a substantial reduction in the top code values in the public use files. For example, the censoring point for interest income was $99,999 in 1997, but only $35,000 in 1998. (See appendix tables 1 and 2 of Burkhauser *et al.*, 2007 for all public use CPS censoring points.)

Hence unlike P90/P10 estimates derived from internal CPS data, P90/P10 estimates derived from public use data have been substantially affected by censoring, and this is especially the case in recent years. But, as Table 4 also shows, censoring problems are not confined to public use data. As can be seen from columns 6 and 7, *Internal-Lower* and *Internal-Upper* values are not the same in each year, although in most cases the difference is relatively small. Hence, when compared to the top coding problems in the public use CPS, the differences between the *Internal-Lower* and *Internal-Upper* series are negligible relative to the differences between the *Public-Lower* and *Public-Upper* series (see Figure 1).

For income, as for wages and salaries and total labor income, there is very little difference between the *Public* series (column 3) and the *Cell-Mean* series (column 5) from 1995 onward: compare Table 4 columns 3 and 4 with Tables 2 and 3. But the situation for income differs from the other variables before 1995. Because P90/P10 estimates of income inequality from both public use data and, to a lesser degree from internal data, are affected by censoring, our *Cell-Mean* series does a much better job of aligning P90/P10 estimates from public use data with the series estimated from internal data. In the years prior to 1995, the *Cell-Mean* series almost coincides with the internal series. But thereafter, like *Public* estimates, *Cell-Mean* estimates tend to slightly overstate P90/P10 relative to corresponding internal values. Although the *Rule-of-Thumb* estimates fall within the range provided by the *Public-Lower* and *Public-Upper* series, they now consistently fall below the range provided by the *Internal-Lower* and *Internal-Upper* series. For researchers interested in capturing long term trends in income inequality, measured using P90/P10 and estimated from public use CPS data, Table 4 shows that top coding is a problem and that our *Cell-Mean* series values do the best job of offsetting it and capturing the P90/P10 trends derived from internal CPS data.

6. Longer-Term Trends in Inequality Using Adjusted Public Use CPS Data: P90/P10 versus Gini Estimates

Researchers in the labor and income inequality literature employing public use CPS data often summarize trends in inequality using the P90/P10 measure rather than more traditional summary measures of inequality such as the Gini coefficient, Theil indices, or the coefficient of variation, because of concerns about censoring in CPS data. We have demonstrated that P90/P10 estimates are also subject to censoring problems, especially when used to measure household income inequality. But we have also shown that, by using a consistent set of cell means created from internal CPS data, one can estimate a P90/P10 series that is quite close to the P90/P10 series estimated with internal CPS data. The issue that we turn to now is whether P90/P10 estimates provide a picture of inequality trends that is robust. Does P90/P10 provide the same picture of inequality trends as a picture based on a measure that uses information about all incomes in the distribution rather than focusing only on two points?

We compare trends in inequality (of wage and salaries income, labor earnings, and the size-adjusted household income of individuals) derived from our P90/P10 *Cell-Mean* series with trends derived from Gini coefficients based on public use and on internal data. We use the Gini coefficient as it is the most commonly-estimated summary measure of inequality in the income distribution literature. We use our *Cell-Mean* series for P90/P10 because it more closely replicates the internal series than any other series currently available to general research.

We derive time-consistent Gini inequality values via a consistent top coding method that is applied to both the public use data and the internal data for the years 1975–2004. We calculate the percentage of individuals subject to top coding in every year for each income source. We determine the year in which the
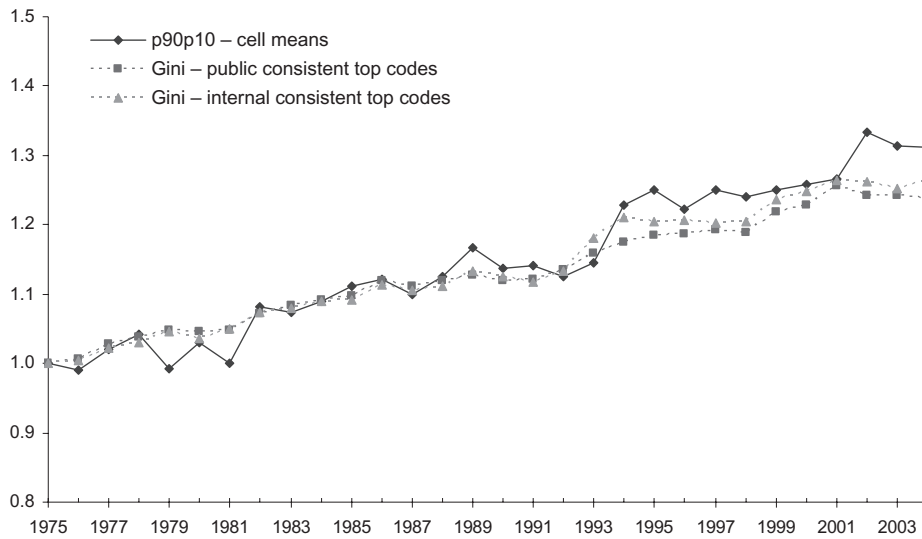
Figure 4. Trends in Consistently Top Coded Gini and Cell-Mean Adjusted P90/P10 Estimates for Wage and Salary Income of Full-Time, Full-Year Workers

Note: The Gini and P90/P10 series are each normalized by their 1975 value.

greatest percentage of the population was affected by the top code for that income source and then top code that income source for every year to yield this same percentage. This procedure ensures that a common and constant percentage of the upper values for each income source are affected in each year. In doing so, we adjust the top codes used for each subcomponent of first wage and salary earnings, then labor earnings and then household income. For a fuller discussion of this method, see Feng *et al.* (2006) for its application to labor earnings, and Burkhauser *et al.* (2008) for its application to size-adjusted household income.

We are interested in comparing trends in inequality based on our adjusted P90/P10 estimates with trends in inequality based on our consistently top coded public use and internal CPS Gini values, so all three series are normalized using year 1975 as the base. Normalized Gini coefficient and P90/P10 estimates for wages and salaries among full-time, full-year workers from 1975 to 2004 are displayed in Figure 4. The P90/P10 series shows a greater degree of variance from one year to the next. According to it, inequality increased less in the early years and more in the later years than is the case according to either of the Gini series, with the difference most pronounced in the last few years.

The estimates for the distribution of total labor earnings among full-time, full-year workers are shown in Figure 5, derived using the same methods as in Figure 4. In this case, there is a much greater difference in the relative trends. Not only is there much greater variance in P90/P10 estimates but, after the first few years, there is also a much greater rise in inequality based on the P90/P10 series over time than that produced by the estimates of the Gini coefficient from either the consistently top coded public use data or internal data.
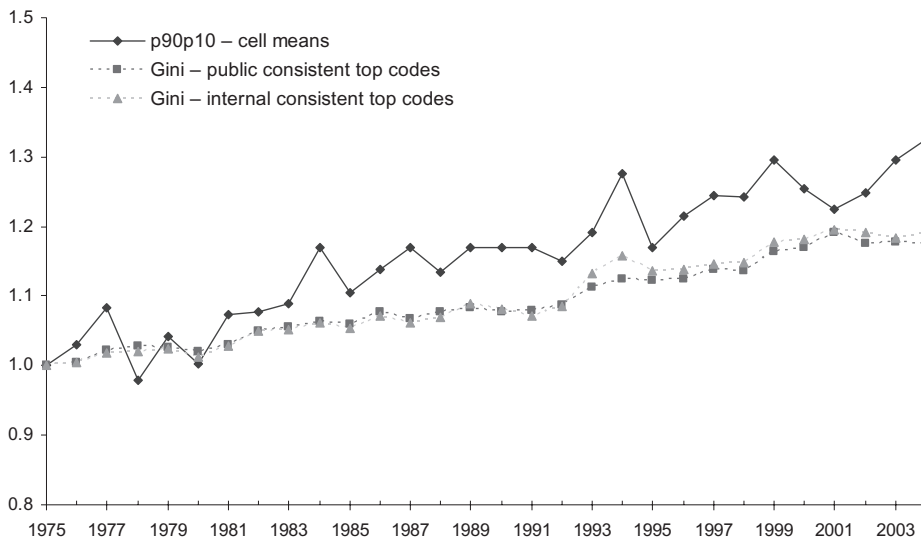
Figure 5. Trends in Consistently Top Coded Gini Coefficient and Cell-Mean Adjusted P90/P10
Estimates for the Total Labor Earnings of Full-Time, Full-Year Workers

Note: The Gini and P90/P10 series are each normalized by their 1975 value.

Estimates for the distribution of size-adjusted household income among individuals, derived using the same methods, are shown in Figure 6. Once again there is much greater variance in P90/P10 estimates over time, and there is now an even greater rise in P90/P10-measured inequality over time. The increase is much greater in magnitude than that indicated by the two Gini coefficient series.[7]

Given the differences in these trends, researchers should be cautious about using the relative position of two points in the distributions of wages and salaries, labor earnings or income to draw conclusions about how overall inequality of each of these income sources changed over the last three decades. The choice of inequality measure matters.

7. Summary and Conclusions

We investigate how P90/P10 is affected by censoring when used to measure inequality in the distribution of wages and salaries, labor earnings and household income. We do so with both public use and internal CPS data. In all cases we found that top coding is less of a problem for researchers using P90/P10 to measure inequality in wages and salaries and labor earnings than it is for those assessing inequality of size-adjusted household income. And, it is far less of a problem in the internal data than in the public use data. Except for the case of the household income distribution, estimating P90/P10 using a rule-of-thumb method to control

[7]We have tested differences in linear trends more formally using regression methods, and shown that the Gini coefficient and P90/P10 estimates pictured in Figures 4, 5, and 6, have statistically different trends. See Burkhauser *et al.* (2007).

Figure 6. Trends in Consistently Top Coded Gini Coefficient and Cell-Mean Adjusted P90/P10 Estimates for the Size-Adjusted Household Income of Individuals

Note: The Gini and P90/P10 series are each normalized by their 1975 value.

for top coding in the public use data does as good a job as using our consistently created cell mean series in estimating P90/P10 values calculated from internal data.

However, we found that the cell mean series we created for all years of public use CPS data yields superior estimates of internal data-estimates of P90/P10 than does either using no cell means or using the cell means that the Census Bureau has provided from 1995 onward.

P90/P10 is only one measure of inequality. Our comparisons of P90/P10 and Gini coefficient series derived using consistently top coded public use or internal CPS data yield large and significant differences in longer term trends for all three of the income definitions considered, but the largest differences by far were for our size-adjusted household income series. Hence researchers should be cautious about inferring longer term trends in inequality on the basis of a single inequality measure.

Furthermore, because as we have discussed above, the U.S. Census Bureau is not alone among statistical agencies in top coding income values, it is important for researchers to carefully consider the potential impact of top coding practices on their estimates of trends in inequality even if they measure inequality using P90/P10.

References

Alderson, Arthur, Jason Beckfield, and Francios Nielsen, "Exactly How Has Income Inequality Changed? Patterns of Distributional Change in Core Societies," *International Journal of Comparative Sociology*, 46(5–6), 405–23, 2005.

Atkinson, Anthony B., "The Long Run Earnings Distribution in Five Countries: 'Remarkable Stability', U, V, or W?" *The Review of Income and Wealth*, 53(1), 1–24, 2007.

Atkinson, Anthony B., Lee Rainwater, and Timothy Smeeding, "Income Distribution in OECD Countries. Evidence from the Luxembourg Income Study," *Social Policy Studies No. 18*, Organization for Economic Cooperation and Development, Paris, 1995.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Trends in U.S. Wage Inequality: Revising the Revisionists," *The Review of Economics and Statistics*, 90(2), 300–23, 2008.

Blau, Francine D. and Lawrence M. Kahn, "Do Cognitive Test Scores Explain Higher U.S. Wage Inequality?" *The Review of Economics and Statistics*, 87(1), 184–93, 2005.

Brandolini, Andrea, "Measurement of Income Distribution in Supranational Entities: The Case of the European Union," in Stephen P. Jenkins and John Micklewright (eds), *Inequality and Poverty Re-Examined*, Oxford University Press, Oxford, 2007.

Burkhauser, Richard V., Kenneth A. Couch, Andrew J. Houtenville, and Ludmila Rovba, "Income Inequality in the 1990's: Re-Forging a Lost Relationship?," *Journal of Income Distribution*, 12(3–4), 8–35, 2003–2004.

Burkhauser, Richard V., J. S. Butler, Shuaizhang Feng, and Andrew Houtenville, "Long-Term Trends in Earnings Inequality: What the CPS Can Tell Us," *Economics Letters*, 82, 295–9, 2004.

Burkhauser, Richard V., Shuaizhang Feng, and Stephen P. Jenkins, "Using the P90/P10 Index to Measure US Inequality Trends with Current Population Survey Data: A View from Inside the Census Bureau Vaults," *ISER Working Paper 2007-14*, Institute for Social and Economic Research, University of Essex, Colchester UK, 2007, http://www.iser.essex.ac.uk/pubs/workpaps/pdf/2007-14.pdf

Burkhauser, Richard V., Takashi Oshio, and Ludmila Rovba, "How the Distribution of After-Tax Income Changed Over the 1990s Business Cycle: A Comparison of the United States, Great Britain, Germany, and Japan," *Journal of Income Distribution*, 17(1), 87–109, 2008.

Daly, Mary C. and Robert G. Valletta, "Inequality and Poverty in United States: The Effects of Rising Dispersion of Men's Earnings and Changing Family Behaviour," *Economica*, 73(289), 75–98, 2006.

Danziger, Sheldon and Peter Gottschalk (eds), *Uneven Tides: Rising Inequality in America*, Russell Sage Foundation, New York, 1993.

DiNardo, John, Nicole Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semi-Parametric Approach," *Econometrica*, 64(5), 1001–44, 1996.

Feng, Shuaizhang, Richard V. Burkhauser, and J. S. Butler, "Levels and Long-Term Trends in Earnings Inequality: Overcoming Current Population Survey Censoring Problems Using the GB2 Distribution," *Journal of Business and Economic Statistics*, 24(1), 57–62, 2006.

Fichtenbaum, Rudy and Hushang Shahidi, "Truncation Bias and the Measurement of Income Inequality," *Journal of Business and Economic Statistics*, 6(3), 335–7, 1988.

Gottschalk, Peter and Sheldon Danziger, "Inequality of Wage Rates, Earnings and Family Income in the United States, 1975–2002," *The Review of Income and Wealth*, 51(2), 231–54, 2005.

Gottschalk, Peter and Mary Joyce, "Cross-National Differences in the Rise in Earnings Inequality: Market and Institutional Factors," *The Review of Economics and Statistics*, 80(4), 489–502, 1998.

Gottschalk, Peter and Timothy M. Smeeding, "Cross-National Comparisons of Earnings and Income Inequality," *Journal of Economic Literature*, 35(2), 633–87, 1997.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce, "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, 101(3), 410–42, 1993.

Karoly, Lynn A. and Gary Burtless, "Demographic Changes, Rising Earnings Inequality, and the Distribution of Personal Well-Being, 1959–1989," *Demography*, 32(3), 379–405, 1995.

Katz, Lawrence F. and David H. Autor, "Changes in the Wage Structure and Earnings Inequality," in O. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*, Vol. 3A, North-Holland, Amsterdam, 1463–555, 1999.

Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963–87: Supply and Demand Factors," *Quarterly Journal of Economics*, 107(1), 35–78, 1992.

Larrimore, Jeff, Richard V. Burkhauser, Shuaizhang Feng, and Laura Zayatz, "Consistent Cell Means for Topcoded Incomes in the Public Use March CPS (1976–2007)," *Journal of Economic and Social Measurement*, forthcoming.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *The American Economic Review*, 96(3), 461–98, 2006.

———, "The Changing Nature of Wage Inequality," *Journal of Population Economics*, 21, 21–48, 2008.

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

LIS (Luxemburg Income Study), Income Inequality Measures, http://www.lisproject.org/keyfigures/ineqtable.htm (accessed August 1, 2008).

Manski, Charles F., "The Selection Problem," in Christopher A. Sims (ed.), *Advances in Econometrics, Sixth World Congress*, Cambridge University Press, Cambridge, 143–70, 1994.

Pencavel, John, "A Life Cycle Perspective on Changes in Earnings Inequality among Married Men and Women," *The Review of Economics and Statistics*, 88(2), 232–42, 2006.

Prus, Steven and Robert Brown, "Income Inequality over the Later-life Course: A Comparative Analysis of Seven OECD Countries," *Working Paper 435*, Luxembourg Income Study, Luxembourg, 2006, http://www.lisproject.org/publications/liswps/435.pdf.

Smeeding, Timothy, M., "Twenty Years of Research on Income Inequality, Poverty, and Redistribution in the Developed World," *Socio-Economic Review*, 2(2), 149–63, 2004.

Trejo, Stephen J., "Why Do Mexican Americans Earn Low Wages?" *Journal of Political Economy*, 105(6), 1235–68, 1997.

UNU-WIDER, *World Income Inequality Database, version 2.0c (May)*, United Nations University–World Institute for Development Economics Research, Helsinki, http://www.wider.unu.edu/research/Database/en_GB/database (accessed August 1, 2008).

Welniak, Edward J., "Measuring Household Income Inequality Using the CPS," in James Dalton and Beth Kilss (eds), *Special Studies in Federal Tax Statistics 2003*, Statistics of Income Directorate, Internal Revenue Service, Washington DC, 2003.