Measuring Trends in the Distribution of Annual Earnings

David Pattison[1]

Office of Economic Analysis and Comparative Studies, Office of Research, Evaluation, and Statistics, Office of Retirement and Disability Policy, Social Security Administration.

Proceedings of the 2018 Federal Committee on Statistical Methodology (FCSM) Research Conference

Abstract

The Social Security annual earnings data includes many observations of very low earnings, either from persons who work very sporadically or from workers who begin or end a multi-year earnings spell part way through the year. If the prevalence of these low earnings varies over time, the variations can introduce trends in measures of average earnings or the median and other percentiles of the distribution, trends that might not be representative of actual changes at middle or high earnings.

A further complication is the changing age and sex composition of the workforce. The aging of the baby boom cohort through the prime working years, the large increase in female labor force participation, and recent declines in male participation at younger ages will all have affected the earnings distribution.

Several techniques are available for adjusting the measurement of trends in the earnings to protect against these variations, including the exclusion of earnings below some threshold like the federal minimum wage or a fixed percentage of the national average wage, exclusion of earnings at the beginning or end of an earnings spell, and the age-sex adjustment of earnings. The sensitivity to choice of technique needs to be checked.

This paper explores these issues using the Social Security administrative data, which have uncensored annual earnings records back to 1978. Two age-sex adjustments and several trimming thresholds are applied, alone and in combination, and the effects on trends in means and variances are analyzed.

[1] This paper was prepared for conference presentation and has not undergone full review. The findings and conclusions expressed in them are the author's and do not necessarily represent the views of the Social Security Administration.

I. Introduction

Average per-capita wage earnings rise over time with increases in prices and productivity. If this rise in measured average earnings were shared broadly and equally across the earnings distribution, then by dividing all earnings in a year by the average earnings we would arrive at a constant distributional shape.

This does not happen. The most familiar departure from a constant shape is the faster growth of earnings at the higher end, with the result that average earnings over the whole distribution rise faster than median earnings. If the earnings distribution were adjusted as suggested above, dividing by the measured average earnings, it would exhibit a widening standard deviation, and if were further standardized by dividing by the measured standard deviation it would still show significant changes in shape.

There are many reasons the shape does not stay constant. Among them:

- The entrance of the baby boom cohorts into the work force, and now their exit, has pushed the age distribution around, first increasing the frequency of younger workers at lower pay, then shifting the distribution toward higher workers with more pay, and now retiring those workers.
- The entrance of large numbers of women into the participating labor force will have altered the distribution.
- More subtle changes in labor force participation, like the increasing participation of older workers or the declining participation of younger male workers, can be expected to have effects on the distribution of earnings. It is unlikely that segments of the work force that enter or leave employment have the same distribution as the more stably employed segments.
- Income shifting – changes in the classification of personal income into or out of wage and self-employment earnings -- will change the distribution of observed wage earnings.
- The administrative data has many very low earnings, and these earnings might be particularly prone to changes in the age distribution or changes in participation among very young or very old workers. These earnings are of little importance in themselves, but when earnings are transformed to their logarithms, as is standard in studies of earnings dynamics, shifts in the frequency distribution in the lower tail have large effects.

The focus of this paper will be the first two points – the effects of the changes in the age-sex composition as the baby boomers moved through the work force and as women entered it – and the final point – the effects of very low earnings. The changing age-sex composition can be addressed with what are in effect price-indexing techniques that combine observations at detailed age-sex levels of prices (earnings) and quantities (number of workers). The low earnings are commonly addressed by simply excluding them from the sample – setting a low earnings threshold for including annual earnings in the analysis.

Low earnings can exist in the data for a variety of reasons. The very low earnings (such as earnings below $100), might arise from reporting anomalies that vary from year, particularly for the very young or for workers at the retirement ages, and it might be desirable to exclude the noise from these anomalies. Even if they are not reporting errors, the low earnings are probably not reflective of trends and dynamics in the central portion of the earnings distribution, and for analysis of earnings trends and earnings dynamics it is often desirable to suppress the noise from the lower tail the influence from these low earnings needs to be suppressed. For retirement studies we would like to filter out those workers who would consider themselves retired but still earn small amounts either regularly or sporadically.

Figure 1: Distribution of 2013 annual earnings

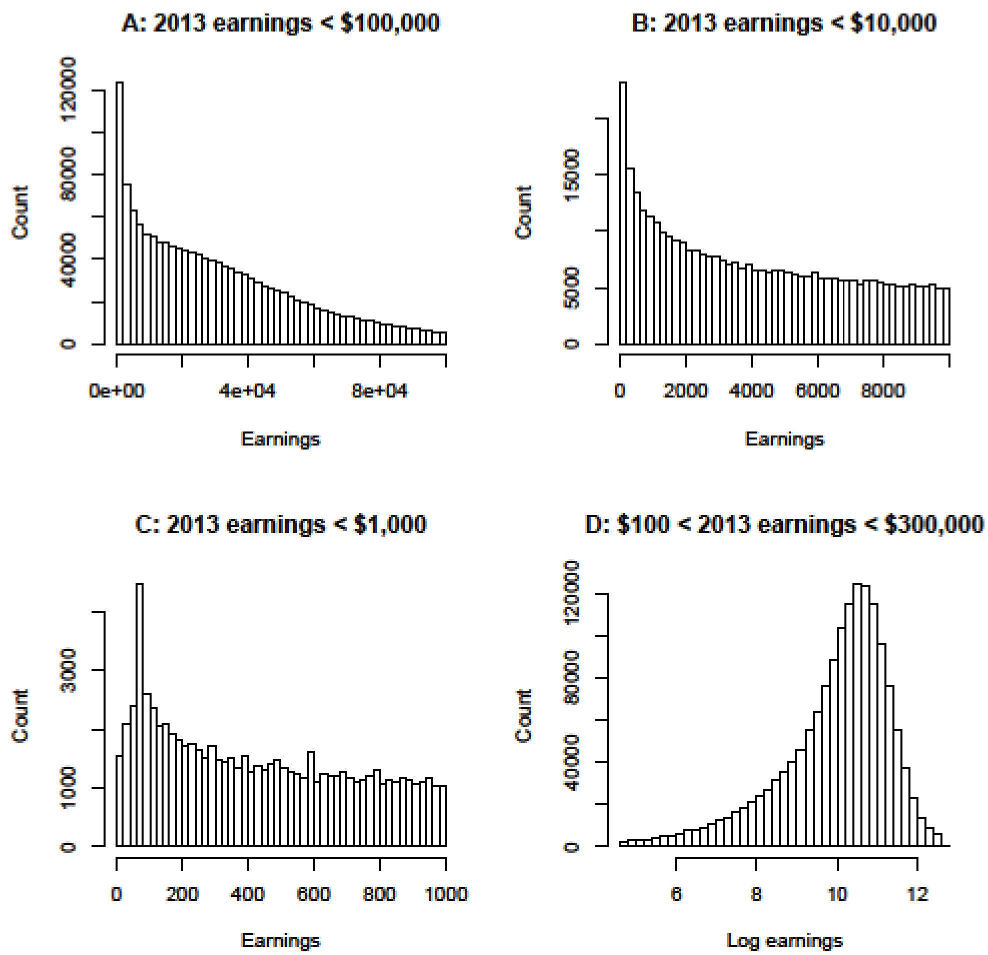

Note: Panels A - C: nominal earnings below selected upper thresholds. Panel DL Log of annual earnings. Source: Tabulation from Master Earnings File

Figure 1A-C gives an indication of the problem. These are histograms of the distribution of 2013 non-zero annual earnings, using SSA data to be discussed below. The mean of the 2013 sample is about $43,200. The median is about $28,200. Most of us will be familiar with the long upper tail of the earnings distribution. The maximum value in the 2013 sample was too high make a useful histogram of the whole earnings distribution. With a cap applied at $100,000 (Figure 1A), the most striking feature is at the other end: the modal value is at the bottom.

Successively lower caps bring this lower tail into focus (Figures 1B-C) The modal $100 interval of the nominal earnings distribution is actually the bottom one.[2]

Figure 1D gives the histogram of log earnings. The mode, if transformed back into nominal dollars, is in the vicinity of $30,000 dollars. The central part of the distribution is closer to the symmetrical bell shape of the normal

[2] The most common value in single dollar intervals is the interval just under $80. I have not learned what these are. The workers with these values are typically over 65. My current data extract does not allow me to distinguish earnings proper from deferred compensation payments.

distribution. There is now, however, a long lower tail. To make the histogram manageable, the lower tail has been trimmed at $100, above the modal interval for the nominal earnings.

The mean log earnings, if transformed back to nominal earnings, is $23,400. If the log earnings distribution were exactly symmetrical this value would match the median earnings value of $28,200. The lower tail, however, brings this value down. (If the log earnings distribution were not just symmetrical but were normal, the mean plus half the variance, transformed into nominal earnings, $62,300, would match the mean nominal earnings or $43,200. The long lower tail, by increasing the variance, drives this value upward.)

Whatever the reason, the very low earnings are trimmed by most researchers. Typically a year is chosen, a threshold is set for that year, and earnings in other years are indexed to that year. The analysis might be carried out, for example, on $2015 earnings using as a threshold 1,000 hours at the 2015 federal minimum wage.

There is very little guidance, however, for the choice of threshold. Some studies carry out a sensitivity analysis, changing the threshold level and checking that results do not change, at least qualitatively. The choice of threshold in these studies is actually a double choice: the level of the threshold and the index that is used to bring the earnings and the threshold together across the years. Indexing, for example, 1985 earnings to 2015 to compare with the 2015 threshold is equivalent to indexing the 2015 threshold back to 1985 to compare with the 1985 nominal earnings, and there is no guarantee that a price-indexed threshold is the appropriate threshold to use when the main body of earnings has grown at a different rate.

With luck, analyses of the earnings data won't be much affected by the level of the trimming threshold or its indexing. But it is something that needs to be checked. If the lower tail of the log earnings distribution is the more variable, and if the chosen index is such that the trimming threshold drifts down relative to the body of earnings, the measured variability will increase even if the variance in the earnings distribution has not actually increased. Such artifacts of the choice of measure are inevitable; the hope is that they will be small.

That is the question explored in this paper. Over a long span of years, even the choice of a price index or wage index will not be innocuous. There is an obvious difficulty in defining an index stable with respect to a shifting distribution when the goal is the measurement of the shifting distribution itself. One has to start somewhere, and this paper starts by indexing measures relative the growth in the median wage and comparing them to measures using average wages or the PCE or CPI price indexes.

The shifting age-sex composition of the population presents an additional hurdle for the analysis. There is a component of the trends in the mean and variance of the earnings distribution that is attributable to this shifting composition. The shifting composition interacts with a trimming threshold: when the population is relatively concentrated in the higher earnings years, relatively smaller proportions of workers will be affected by any trimming of low earnings. An attempt will be made to check for this aspect of the measurement of trends as well.

The focus here is on the measurement of trends in the whole earnings distribution over long spans of time. Many studies of earnings dynamics focus on subgroups of the population, such as males in their thirties, and further restrict themselves to full-time workers, with some evidence of strong attachment to the labor force. The issues dealt with here are somewhat ameliorated for these narrower groups, particularly the problem of shifting age-sex composition. But they don't go away entirely, and most studies use a low earnings filter of some sort, often as part of the definition of full-time workers.

II. Data

The data are derived from Social Security's 1 percent research sample from the "detail segment" of Master Earnings File (MEF), with annual earnings reports from 1978 on. The detail segment contains, among other items, earnings and deferred compensation amounts as reported to the IRS on the W-2 forms. (The MEF also contains self-employment earnings reports, which won't be used in this study.) The detail earnings reports are used each year to calculate the national average wage index (NAW). The OASDI creditable earnings (earnings in covered employment up to the taxable maximum) are posted to the "summary segment" in the MEF and used as the basis for the calculation of insured status and benefits. Because benefits depend only on the summary earnings, there is more emphasis on correcting those earnings if there were errors in the initial reports. Although the detail segment tends to be corrected as well, there may have been fewer such corrections in the early years of the detail segment postings. Nevertheless, the detail segment earnings are considered fairly accurate reports of earnings after about 1982. Although the summary earnings might be slightly more accurate, for the study of earnings trends they are deficient in two regards. First, they are only contain earnings from jobs in OASDI-covered employment. Second, for those jobs they are capped at the taxable maximum in the year of the earnings.

Some of the rare but large errors in the detail segment are limited by a series of edits that reduce implausibly large reported earnings. In order to avoid the years with the largest frequency of apparent errors, the span in this study is limited to earnings from 1982 on. The last year of earnings in the files used here was 2013.

The earnings in this study are linked to SSA administrative data on sex, year of birth, place of birth, year of issue of SSN, and year of death. Observations before the year turning age 16 or after the year turning age 80 are excluded. Persons born in the U.S. are tabulated as in the population for the full observable span. Persons born outside the U.S. are tabulated as in the population from the year the SSN was issued. If the year of death occurs in or before 2013, the final year is included.

The linkage to birth and death data is made even for SSNs for whom there are no recorded earnings. The total observation count thus gives an estimate of the age 16 to 80 population each year. To the extent that some persons aged 16 or more have not yet been issued an SSN, this count will be an underestimate. (With registration at birth this is not significant for more recent cohorts. For earlier cohorts many non-working spouses did not receive an SSN until they applied for spouse or widow benefits; thus the undercount for these earlier cohorts will have shrunk as the cohorts moved into the retirement ages.) More likely, the population counts are over estimated, for two reasons. First, the Social Security data does not indicate if a person emigrates from the U.S., so such persons will continue to be counted as in the labor force. Second, although death reporting of beneficiaries is accurate, Social Security does not always receive reports of deaths of workers who are not yet receiving benefits.

Although adjustments can be made to the data to approximately correct for these errors in counting the population, these adjustments are not made here. Most of the study will focus on earnings among those with earnings, for whom the population undercount is not problem. The estimates of employment rates, however, which use the population counts as a denominator, will be misestimated to an unknown extent.

The annual earnings data cannot distinguish part-year or part-time workers. The thresholds here are therefore not strictly comparable to Census survey studies which limit themselves to full-year full-time workers.

III. Measures of the earnings distribution

The measures selected for this study are the mean log earnings, the percent of the population employed, and the variance of log earnings, with the variance split into within and between components among the single-year age-sex cells.

Mean log earnings

The logarithm of earnings (or "log earnings") is commonly used in earnings studies. As already discussed, the log earnings distribution, at least its central portion, is much closer to a symmetrical bell-shaped distribution than its unlogged counterpart. For studies of earnings dynamics over long periods, the log earnings provides a better base for studies of the variance of earnings: as earnings rise over time the variance can be expected to rise, but the variance of log earnings could rise, fall, or remain steady, and the null hypothesis of an unchanging log earnings distribution can plausibly be considered.

Figure 2: Age profiles of 1982-2013 earnings

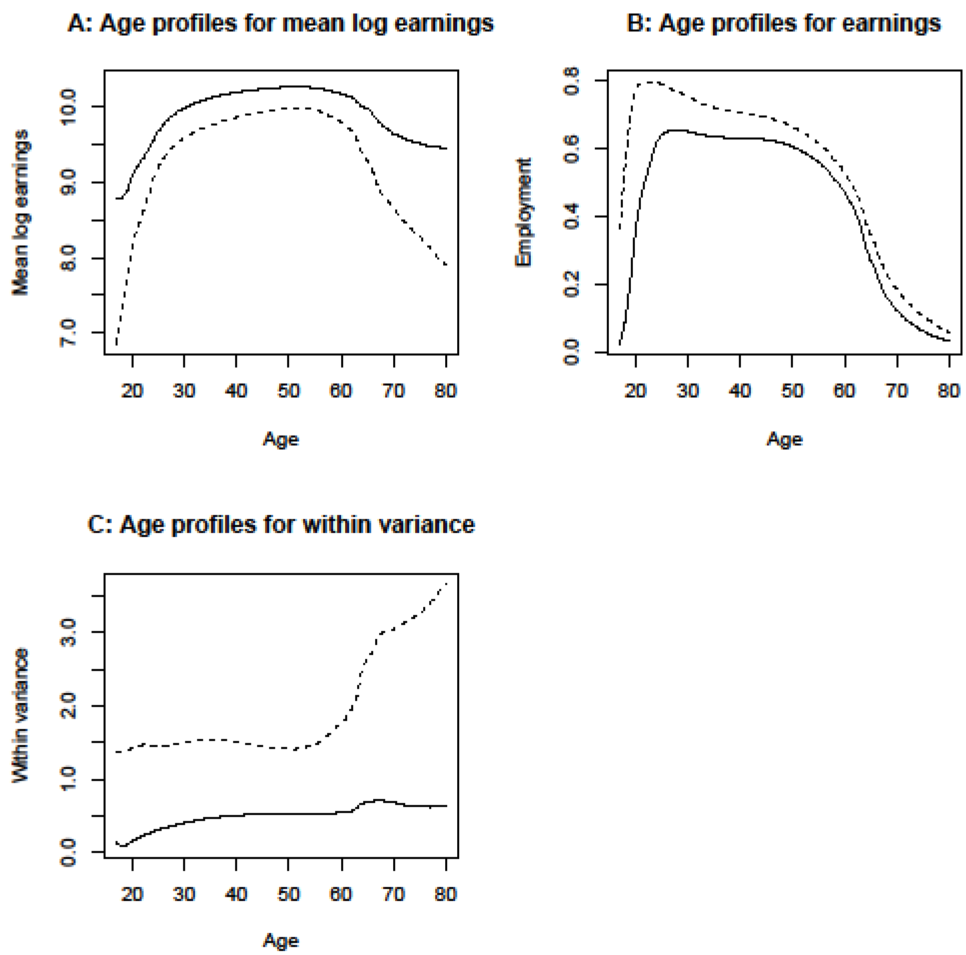

Note: Profiles are the age coefficients from regression on single-year age and year dummies. Dashed line includes all non-zero earnings. Solid line includes only earnings above 25% of median non-zero earnings. Source: Authors calculations on MEF earnings file.

The shifting location of the center of the earnings distribution could also be described with the median, rather than the mean log earnings. An advantage of the median is that it is less affected by changes at the extremes. If, however, changes in the lower tail are attributable to movements into or out of the labor force rather than changes in earnings among those in the labor force, the median will be affected. The examination of the behavior of medians and percentiles when the age-sex composition is changing has been postponed to a subsequent study that will look as well at the marginal earnings of those who shift into or out of the labor force. For this study, the focus will be on the mean log earnings.

There is a distinct age pattern to earnings, shown in Figure 2A, which displays the age coefficients from a regression of log earnings on single-year age and year dummies over the earnings sample (combining both sexes). Average earnings are low among the youngest workers and the oldest workers. This is particularly true when all nonzero earnings are included in the regression (dashed line) but it remains true when the sample is restricted to those above 25 percent of the median earnings each year (solid line). This large variation in the age profile generates shifting composition effects as the baby boom cohorts move into and out of the high earnings years.

Employment

The employment rate for an age-sex cell is defined here as the number of persons in that cell with earnings divided by the population in that cell. As already discussed, the denominator in this rate, the population count, will be overestimated to the extent that some of the apparent nonworkers might have exited the population through emigration or unrecorded death. For our purposes the misestimate is probably small enough that our qualitative analysis is not impaired.

There is a distinct age pattern to employment (Figure 2B), with low employment rates among the young and the old. If all workers are counted as employed (dashed line), the employment rates are higher than if only those with earnings above 25 percent of the median are counted (solid line). These patterns will generate shifting compositional effects as the baby boom ages.

Variance of log earnings

The variance of log earnings is the average squared deviation from the mean log earnings. When the earnings in a year can be grouped into age-sex cells, the variance can be decomposed into a “between” and a “within” component. Each age-sex cell has its own mean log earnings, and the between component of the variance is the variance of these age-sex mean log earnings among the cells. This component of variance accounts for differences like those, for example, between the average earnings among 45-year-old men and the average among 45-year-old women and the differences between 45-year-old men and 25-year-old men. The within component is the remainder of the variance, equal to the variance of the log earnings after each observation’s appropriate age-sex mean has been subtracted. The within and between decomposition is common in earnings studies, used along many dimensions besides those of age and sex. A study might compare the differences in earnings between states with those within states, or the differences between firms or industries with those within firms or industries. As is common in such studies, the age-sex within variance is much larger than the between variance: there is more variation in the earnings among, for example, 45-year-old men than there is between the average earnings of 45-year-old men and the average earnings of other age/sex groups.

Figure 3C plots the regression-estimated age profile for the within component of variance, combining both sexes. Here the difference between the total non-zero sample (dashed line) and the sample above the 25 percent of median level (solid line) is particularly large, especially at the retirement ages, an indication of the outsized effects on variance measures of the very low earnings in the lower tail of log earnings.[3]

IV. Adjustments to the earnings sample

This paper considers two sets of adjustments designed to remove from estimates of earnings trends the effects of changes in the composition of earnings. Adjustment for changes in the age-sex composition removes the effects from the aging of the baby boom and the entry of large numbers of women into employment. A trimming threshold for low earnings removes the statistical noise associated the presence in the administrative data of very low earnings.

Adjustments for changing age-sex composition

Over the span of this study, the age-sex composition of the work force changed both because of increasing labor force participation by women and because the baby-boom bulge moved through the prime working years and toward retirement. Given the differences by age exhibited in the age profiles, plus the differences by sex not illustrated here, these changes can be expected to have had some effect on overall averages and variances.

With the individual data, tabulations can be made by year, single year of age, and sex of population counts, number of workers with non-zero earnings, average log earnings, and the variance or standard deviation of the log earnings.

[3] Some of the low earnings at older ages is due to retirement during the year. But even when the sample is restricted to observations in which the following year is not a year of zero earnings, the large variance from low earnings, although reduced, is still present.

These tabulations form the raw data for age-sex adjusted statistics. (Parallel tabulations are made at each trimming threshold of the number of workers above the threshold, their average log earnings, and the variance.)

The unadjusted measures can be tabulated directly from each year's data without the detailed age-sex tabulations, but exactly the same measures can be calculated from the age-sex tabulations, a calculation that provides a parallel with the age-sex adjusted calculations. The unadjusted measure of, for example, mean log earnings for a year is simply the sum over all the age-sex cells for that year of the mean log earnings within each cell, multiplied by the number of workers in each cell, with the sum then divided by the sum of the workers in all the age-sex cells. This is a weighted average, with the weight for a cell being the number of workers in that cell that year divided by the number of workers in all the cells for that year.

The first form of age-sex adjustment is also a weighted average of the cell values for average log earnings, but the weights are held constant at those observed in a particular year, instead of varying with observation year. For the 1982-2013 sample span, there are 32 possible weight distributions that could be used. Each has the advantage that the calculated time series using the fixed weights will free of any effects from changing age-sex composition. Each has the disadvantage, however, of focusing on earnings changes in some particular part of the age distribution. If the weights are fixed to a year in which 30-year-olds were prominent in the labor force, the resulting average earnings series will give more weight to growth in earnings among 30 year olds. If they are fixed to a year in which 60 years olds are prominent, the time series will give more weight to growth in earnings among 60-year-olds.

Rather than select a particular year or average over several years, the approach taken here is to simply calculate and plot all 32 fixed weight time series. This approach has the virtue of indicating the uncertainty associated with choice of a particular index.[4]

The second form of age-sex adjustment uses chained price index techniques to eliminate the effects of year-to-year changes in age-sex composition while allowing the weights to vary over time. Each age-sex cell has a rate of growth from one year to the next, and by applying appropriate weights to these cell-growth rates the growth in aggregate earnings can be neatly divided into a growth in the number of workers and the growth in average earnings.[5] The resulting growth in average log earnings will tend to lie somewhere within the bundle of growth 32 growth rates calculated under the previous technique, although not necessarily.

An alternative approach for age sex adjustment would be to use the period profile generated by regressions on single-year age and period dummies. (The age profiles were shown earlier.) The regressions used here did not separate the sexes; nevertheless in a few comparisons (not shown here) the regression period profile moved in parallel with the chained estimates in most years.

The first technique – the fixed-weight calculation -- was used for all the earnings measures calculated here: percent of workers employed, average non-zero log earnings, variance of log earnings, within variance of log earnings, and between variance of log earnings.

The second technique – chain weighting - was used for percent of workers employed, average log earnings, and the within variance. The technique cannot be used for the between component of the variance or for the total variance.

To get a sense of interactions between age-sex adjustment and the use of trimming thresholds the age-sex adjusted earnings were calculated not only for all non-zero earnings but also for a trimming threshold set at 25 percent of median earnings.

---

[4] It is possible, using reweighting calculations, to make the equivalent of fixed-weight age-sex adjustments for medians and other percentiles, including measures of variance like the 90/10 percentile ratio. These are not reorted in this study.

[5] Technically, the index used here is the Törnqvist index, in which the cell growth rates are weighted by the cells' shares of aggregate earnings over the growth period. There are several options for the appropriate weights; the ones used in this study were the Sato-Vartia weights.

Figure 3: Cumulative growth of indexes commonly used for trimming thresholds

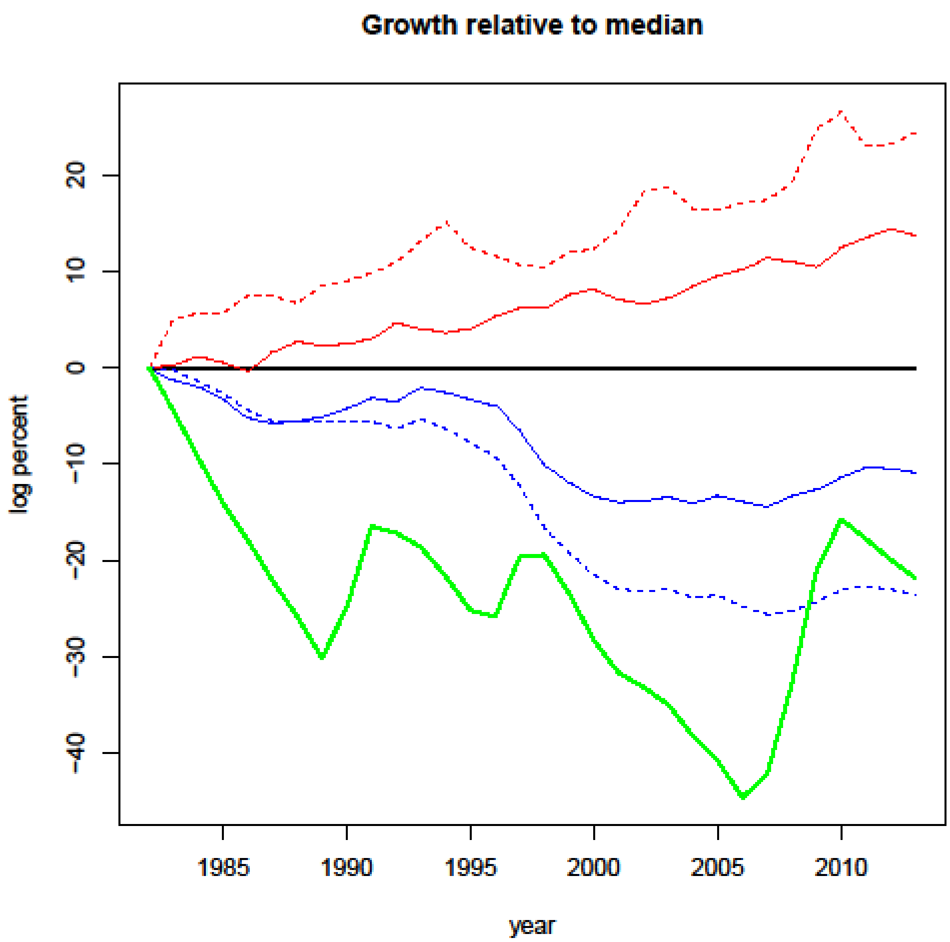

Note: Cumulative growth (in log percent) minus cumulative growth in median (in log percent). The median is thus represented by the horizontal black line. The wage index is in red, solid for the unlagged NAW and dashed for the two-year-lagged NAW. The price indexes are in blue, solid for the CPI and dashed for the PCE. The FMW is in green.

Trimming thresholds for low earnings

This section contains a non-systematic sampling of existing earnings thresholds, either those that might be taken from program provisions or those that have been applied in the literature. The purpose is to get a range of thresholds to be tested. The thresholds that will be used in this paper, ranging from 0 percent to 35 percent of the median wage and indexed by the national average wage (NAW) index, the personal consumption expenditure (PCE) deflator, and the consumer price index (CPI), were derived from this survey. [6]

[6] Many of the studies surveyed actually apply two thresholds: a dollar threshold like those studied here and a limitation to full-time full-year workers, using the hours and weeks information available on many survey files. This additional filter is not available for the SSA earnings data, although a later stage of the study will consider limiting the annual earnings sample to years in which the sampled worker had also earnings in the preceding and following year or years, excluding observations from the years in which workers entered or left the work force and excluding some less attached workers.

Most commonly the earnings analysis is carried out with earnings from all years converted to real earnings in a particular base year like 2007, using a price index like the CPI or PCE. A single threshold value, like half the 2007 FMW, is calculated, and all real earnings below that threshold are deleted from the sample.

The choice of the index, as it turns out, is often more critical to the compositional adjustments than the choice of the level in a particular year. The choice is almost always made without discussion (some studies even neglect to mention the index used to calculate the real earnings.) If the selected index grows faster or slower than the growth in median earnings, the threshold will cut off a different portion when applied to the 1982 earnings distribution than when applied to the 2007 distribution.

Over the study span the differences in the cumulative growth among the indexes can be large. Figure 3 gives the cumulative growth of the indexes used here.[7] The fastest growing have been the NAW and the lagged NAW. (The Social Security program amounts like the maximum taxable earnings are indexed to the NAW with a two-year lag, and that lagged value has been used in some studies. In a time of slowing growth the cumulative lagged NAW will rise more than the unlagged, but the difference shown here is mainly due to the divergence at the start, reflecting changes in the unlagged NAW just before the start of the study span.) The CPI grew more than the median wage at the beginning of the period, but then slowed down. The PCE and the FMW declined over the period relative to the median wage. Both price indexes, CPI and PCE, have tended to grow at the same pace as the median wage except for occasional shifts, particularly in the PCE in 1995-2000.

Figure 4 gives a sampling of the thresholds that have been applied. All except one used the method just described: setting a threshold in a given year (at the level marked by a circle), then indexing by a chosen index. The chosen index is color coded in Figure 4: red for the NAW, which has tended to rise relative to the median wage, blue for the price indexes (solid for the CPI and dashed for the PCE), both of which have tended to fall relative to the median wage, and green for the FMW. The green line is the exception that simply used one-eighth the FMW for the trimming threshold. (The NAW-indexed thresholds also distinguish between those that are in effect based on a two-year lagged NAW, dashed, and those based on the current NAW, solid.)

In my tabulations I index the threshold from the base year to the year of the earnings before applying the threshold, but most commonly the earnings are indexed to the base year before applying a constant base year threshold. For variances of log earnings or for ratios of percentiles the two approaches give identical results. For the level of earnings the typical approach will remove some or all of the growth in nominal earnings, while the approach I use here will leave that growth in.

An appendix gives citations and more details on the approaches surveyed. Here (and in Figure 4) they are summarized without citation.

I found no examples of an absolute dollar cutoff, such as $1 or $100, which would have the problem of shrinking relative to the wage distribution as wages grew.[8] Although the threshold is sometimes defined by a dollar amount, the indexing of the earnings to the base year for that dollar amount means that the threshold effectively grows with the index. One study, for example, using earnings indexed to 1979, deleting from the sample any (hourly) earnings less than $1. (That $1 hourly earnings threshold is not as low as it might seem. Working for 40 weeks at 35 hours would give annual earnings of $1,700, which was 14 percent of the tabulated median in that year.)

---

[7] The percentages here are "log percents" – 100 times the log of the value relative to its 1982 value. Log percents, which will be used throughout this paper, have the advantage of being additive: with regular percentages, a 20 percent increase followed by a 20 percent increase is a 44 percent increase rather than a 40 percent increase, but with log percentages a 20 log percent increase followed by a 20 percent log percent increase is a 40 percent log percent increase.

[8] Some studies, although very few, used all positive earnings, effectively a threshold of zero. A zero threshold is a constant absolute threshold, but it is also a constant relative threshold.

Figure 4: Trimming thresholds used in the literature

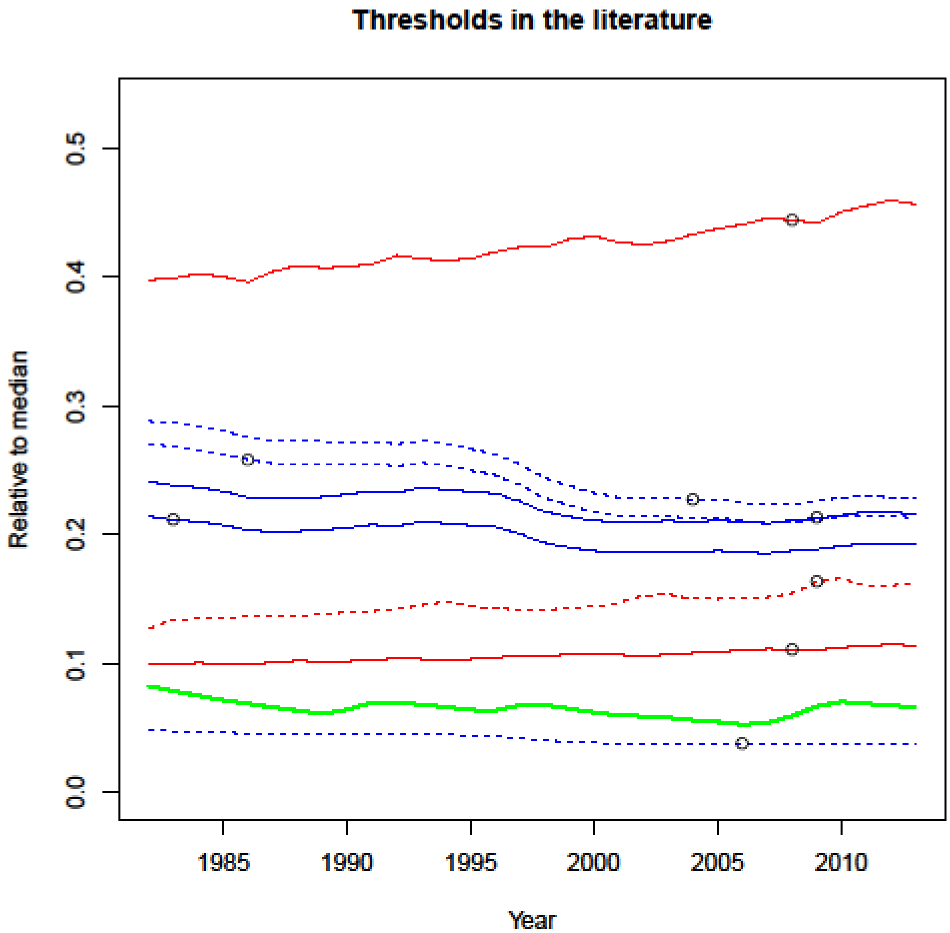

Note: The small circles represent a base year and level in that year for the trimming the threshold and the associated lines indicate the effect of indexing from other years using the index chosen for that study: NAW(red), NAW lagged two years(dashed red), CPI (blue), PCE (dashed blue), and FMW (green).

The most commonly used basis for the thresholds level, at least in my casual survey, is half-time at the federal minimum wage (FMW). Many of the CPS studies limited their sample to full-time full-year workers but added an additional FMW-based threshold to filter out workers with reported amounts believed to be plausibly low for a full-year full-time worker. The most common choice is an estimate of earnings for a half-time FMW worker. The exact definition of half-time worker varies. 1,000 hours appears to be the most common. Some studies are framed in terms of weekly wages, with half-time at 20 hours, which works out to 800 hours for a 40-week year of 1,040 hours for a 52-week year.

These thresholds, even when using the same price index, are at different levels because of the different base years, but the resulting thresholds have tended to fall between 15 percent and 25 percent of median wages

Some researchers have based their threshold on the FMW, but use a lower value than half-time. One study used 500 hours at the FMW, representing quarter-time FMW. Another used 1/6 of the FMW, applied to hourly wages.[9]

[9] This study argued that the ½ FMW threshold was too high because many women in the 1960s who were high school dropouts worked as household or farm workers and who were plausibly being paid very low wages. They

Another used 260 hours, representing one-quarter of full-time work at half the minimum wage, or about one-eighth time. At the other extreme, the study that used 500 hours at the FMW also used 2,000 hours at the FMW as a sensitivity check. The latter represents the high end of thresholds that have been applied. The lowest of these thresholds, at 1/8 or 1/6 of FMW, have been in the range of 5 to 10 percent of the median wage.

Other studies have set their threshold with reference to Social Security program provisions. The amount needed for one quarter of coverage (1QC) provides a very low threshold, at about the level of 3 percent of the median earnings, although because it is indexed to the NAW it has tended to rise slowly over time relative to the median wage. The 1QC is a lower threshold that any study has used for its main threshold, but some studies have used the four quarters of coverage amount (4QC), which has been at around 12 percent of median earnings and rising. The Social Security retirement earnings (RET) amount, also indexed to the NAW, has been at about 45 percent of median earnings, higher than the thresholds selected in most of the studies sampled here.

An approach which won't be examined here is to trim the earnings at specified quantiles, like the 1st and 99th or the 10th and 90th. These symmetrical trims are typically applied in earnings dynamics studies of log wages, and have been applied within age groups before calculation a wage regression, or have been applied to the residuals after calculating a wage regression. I hope to examine this approach in later work, since it by design is protected against shifts in the shape of the earnings distribution and shifts in the age-sex composition of the population. One study referred to sensitivity tests indicating that the amount of trimming didn't make any qualitative difference.

This survey provided a range of thresholds to be checked, from about 5 percent of the median wage to about 35 percent, with studies most typically around 25 percent of the medium wage. To cover that range I use thresholds at 0 percent, 5 percent, 15 percent, 25 percent, and 35 percent of the overall median wage.

The effect of indexing will be checked using the 25 percent of median threshold for 1998, indexed before and after that by each of the indexes found in the survey: the growth in the national average wage index (NAW), the same indexed lagged two years, as is used to derive the QC and RET thresholds; two price indices: the PCE and the CPI; and the Federal minimum wage (FMW). Thresholds defined as a fraction of the median wage are effectively indexed by the growth in median wages.[10] So there are really six indexes studied here: the PCE, the CPI, the NAW, the lagged NAW, the FMW, and the median wage. These are the indexes whose cumulative growth was plotted in Figure 3, and the same color coding will be used in subsequent figures that vary the threshold index.

The set of thresholds for study is shown in Figure 5 as a fraction of the median wage.

---

noted that "the conclusions regarding the evolution of female inequality are very sensitive to how one treats these very low observations of weekly wages. If these low observations were dropped from the sample, then there would be no fall in female inequality in the 60's."

[10] Note the recursive problem: if the median is the median of earnings above some threshold, and the threshold is defined relative to the median, exact calculations – more exact than will be attempted here – would require iteration to convergence. I haven't tried to do this; preliminary indications were that it would not make much difference. The median wage series used here is that tabulated from all positive earnings before applying a threshold and without adjustment for age-sex composition.

Figure 5: Trimming thresholds used in the study

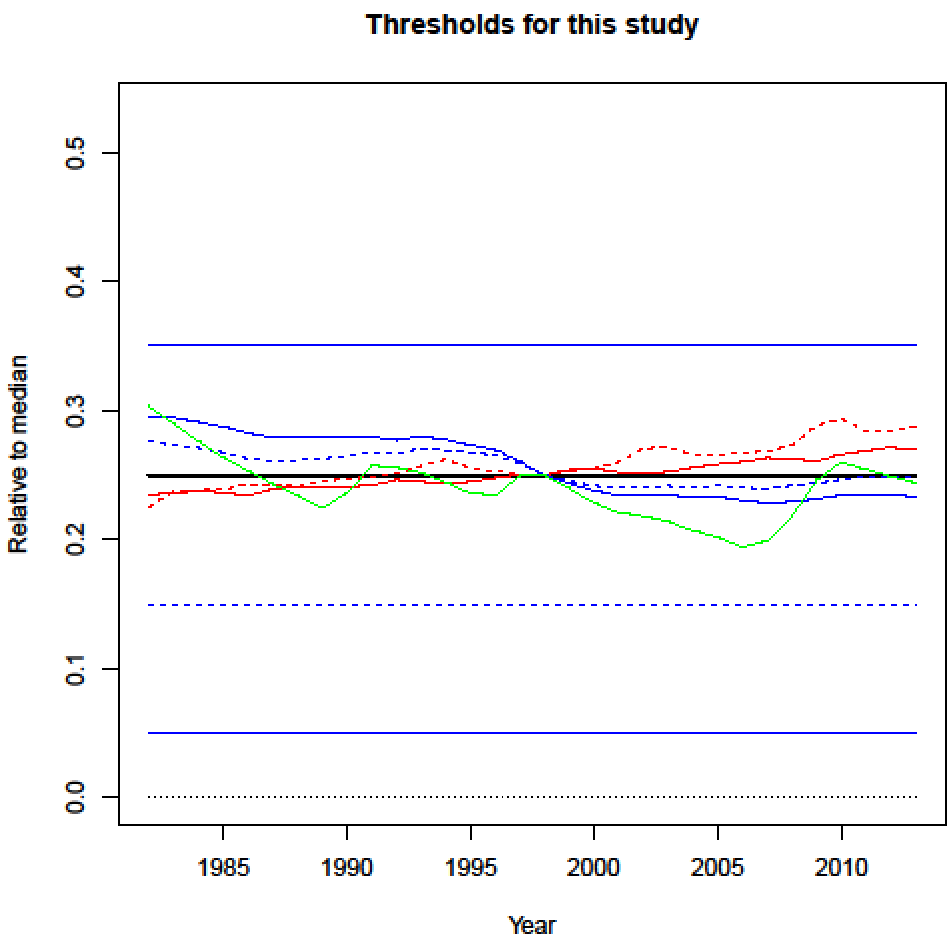

Note: The horizontal lines represent thresholds held at a fixed percentage of median wage, from bottom to top, 0 percent (dotted black), 5 percent (solid blue), a5 percent (dashed blue), 25 percent (solid black), and 35 percent (solid blue). The other lines represent an index set at 25 percent of the 1998 median wage and indexed by the NAW (solid red), the NAW lagged two years (dashed red), the CPI (solid blue), the PCE (dashed blue) and the FMW (green).

As can be seen in Figures 4 and 5, the range of levels in both the sampled literature and the here in the simulated thresholds appears considerably larger than the range of indexes. These large shifts in the threshold levels, however, while they can be expected to have substantial effects on the measures of the mean and variance of earnings, won't necessarily have much effect on the growth rates of these measures, especially so if the shape of the earnings relative to the median earnings stays roughly constant. The indexes that grow differently from the median earnings, on the other hand, will shrink or enlarge the portion of the distribution above the threshold even if the distribution relative to the median is not changing, and therefore have potentially larger effects on the measured distribution of earnings over the threshold.

This paper will attempt to measure how much difference the choice of level and index makes for measures of trends in the growth and variance of earnings. It will not attempt to determined how important that difference is. A few of the studies surveyed here did sensitivity tests for the choice of level (but not the choice of index), and typically conclude that there is no "qualitative difference." As earnings studies become more common, and as the microscope is turned up, what was a qualitative similarity for early research might become a significant difference for later studies, and it will become more important to pay attention to the issues covered in this project.

V. Results

This section presents the effects, measure by measure, of the age-sex adjustments and trimming thresholds. The focus is graphical. For each measure a pair of baselines is calculated: that for all non-zero earnings and that for all earnings above 25 percent of median earnings. The first graph for each measure is that for the levels. As expected, the trimmed version of each of the measures can differ considerably from the untrimmed measure. The interest here, however, is on whether the trimming and the adjustment have large effects on the growth rates, and the remaining graphs for each present the effects on cumulative growth. For the median log earnings, which grew considerably over this period, the earnings are detrended for the cumulative growth analysis; for the other measures the non-detrended growth is presented.

A common color scheme is used for the figures in this section:

- The baseline series are in thick black: solid for earnings above 25 percent of median, dotted for all non-zero earnings.
- For the figures presenting age-sex adjustments:
  - The adjusted versions of all non-zero earnings are in blue: a dark blue dotted line for the chain-weighted version and a set of 32 pale blue lines, usually showing as a pale blue band, for the fixed weight versions.
  - The age-sex adjusted versions for earnings above 25 percent of median are in green: a dark green line for the chain-weighted version, and a set of 32 pale green lines, usually showing as a pale green band, for the fixed-weight version.
- For the figures presenting trimming thresholds:
  - The constant percentage of median thresholds (other than the 0 percent and 25 percent baselines) are in blue: solid for 5 percent and 35 percent of median, dashed for 15 percent of median.
  - The indexed thresholds are color coded: red for wage indexes (solid for the unlagged NAW, dashed for the NAW lagged two years); blue for the price indexes (solid for the CPI and dashed for the PCE); and green for the Federal Minimum Wage.

Not all the series will be easily distinguishable in all the figures, but this arrangement is sufficient to give a notion of the range of effects and to distinguish, for example, the faster growing wage index series from the slower growing price index series.

V.A. Mean log earnings

Mean log earnings over the study period grew considerably over the study period (Figure 6A), whether measured over all non-zero earnings (think black dots in lower band of plotted series) or only over earnings above 25 percent of median (thick black line in top band of plotted earnings). Although the 25 percent of median trim raises the measures considerable, the trimmed and the untrimmed series move in parallel, so that in both cases the cumulative growth over the period is close to 100 log percent (Figure 6B). Figure 6 also includes the age-sex adjusted series, blue for the non-trimmed earnings and green for the trimmed earnings. In Figure 6A it can be seen that although these series grow largely in parallel with the unadjusted series, they grow slightly more slowly, starting above the unadjusted series and ending below them. The difference is more clearly visible in the cumulative growth in Figure 6B.

Figures 6 and 7, but not later figures, also include the median log earnings, a thin black line (just below the upper band of series in Figure 6A and just above the green band in 6B). This also grows in parallel with the mean log earnings series

The differences in cumulative growth are overshadowed by the overall growth in Figures 6B and 6B. To show the differences more clearly, the series are shown in Figure 6C after subtracting a constant average growth rate that equals growth in median log earnings over the period, so that the actual median earnings, shown again as a thin black line, after growing faster than the constant average in the first half and slower than the constant average in the second half, ends at a cumulative detrended growth of zero. In Figure 6C it is clear that although mean log earnings, trimmed or untrimmed, grew about 5 percent more than median earnings over the period, the age-sex adjusted non-trimmed earnings (blue) grew 5 to 10 percent less than median earnings and the adjusted trimmed earnings (green) grew 0 to 5 percent less, or 5 to 10 percent less than the mean log earnings

Figure 6: Mean log earnings, level and cumulative growth

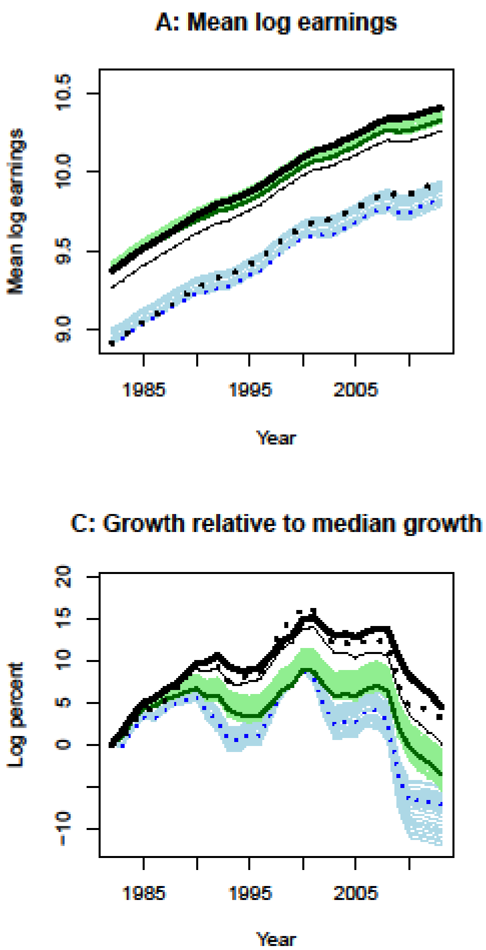

Note: Panel A – The upper set of lines gives measures of mean log earnings for earnings above 25 percent of median: the thick black line is the unadjusted measure, the dark green line is age-sex adjusted using chain weighting, and the green band gives the 32 fixed-weight age-sex adjusted measures. The remaining lines are for all non-zero earnings. The thick black lines are the unadjusted mean, the dotted blue is the chain-weighted age-sex adjusted, and the pale blue band represents the 32 fixed weight adjustments. (Also shown is the unadjusted median earnings, a thin black line just below the upper set of lines.) Panel B: The same series, using the same colors as Panel A, for the cumulative growth from 1982. Panel C – The same series relative to a fixed growth rate equal to the average growth in median earnings 1982-2013.

Figure 7 gives the results for the simulation with different trimming thresholds. The effects of different levels relative to the median (5, 15, and 35 percent, in addition to the 0 and 25 percent baselines) are shown in the left two panels (7A and 7C) for the levels and detrended cumulative change (corresponding to Figures 6A and 6C, and plotted on the same scales for comparability). Although successively higher trimming thresholds lead to higher mean log earnings (6A), the cumulative growth remains tightly clustered (6C).

The effects of different trimming threshold indexes are shown in the right two panels (7B and 7D). Referring to Figure 7D, where the differences are clearest, the wage-indexed thresholds yield a cumulative change on the order of 5 percentage points larger than the mean log earnings, while the price-indexed series yield a cumulative change of the order of 5 percentage points smaller. The difference between wage adjusted thresholds and price adjusted thresholds is therefore on the order of 10 percentage points, rivaling the size of age-sex adjustments. The minimum wage indexed thresholds give even lower cumulative growth over this period, although the increases in the minimum wage before 2010 allowed it to catch up with the price indexes.

Figure 7: Mean log earnings, level and growth under trimming thresholds

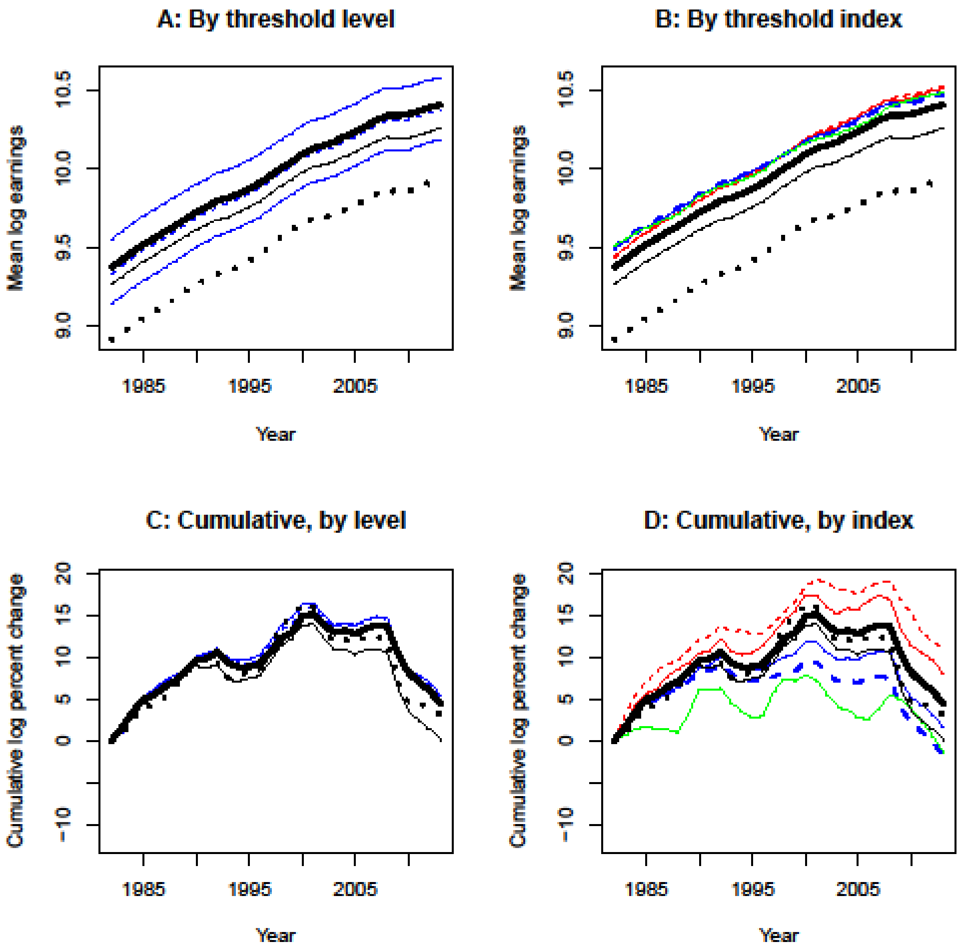

Note: The upper two panels give levels, the lower two give growth from 1982. In all panels, the thick black line represents mean log earnings for earnings above 25 percent of median, the thick black dotted line represent mean log earnings for all non-zero earnings, and the thin black line represents median log earnings for all non-zero earnings. In the two left panels the blue lines represent trimming thresholds at 5 percent of median (lowest), 15 percent (just under the black 25 percent line), and 35 percent (highest). In the right two panels the colored lines indicate indexes for the trimming threshold: NAW (red), lagged NAW (dashed red), CPI (blue), PCE (dashed blue), and FMW (green).

V.B: Employment rates

The earnings data, in conjunction with estimates of the population by age and sex, can be used to make estimates of employment rates, defined here as the number of positive earners above a threshold divided by the number of persons in the estimated population. As already discussed, the population estimates are lower to the extent that they don't include persons who have never received an SSN by the extraction date and they are higher by the number of workers whose deaths were unrecorded or who have emigrated. Because the numerator includes workers who have earnings at any time during the year, the estimates will also differ from estimates derived from monthly data. But these rates should be useful indicators of the overall shifts in participation. The rates calculated here are weighted by the number of workers, thus weighting the employment rates toward the prime working years.

Figure 8: Employment rate, level and cumulative change

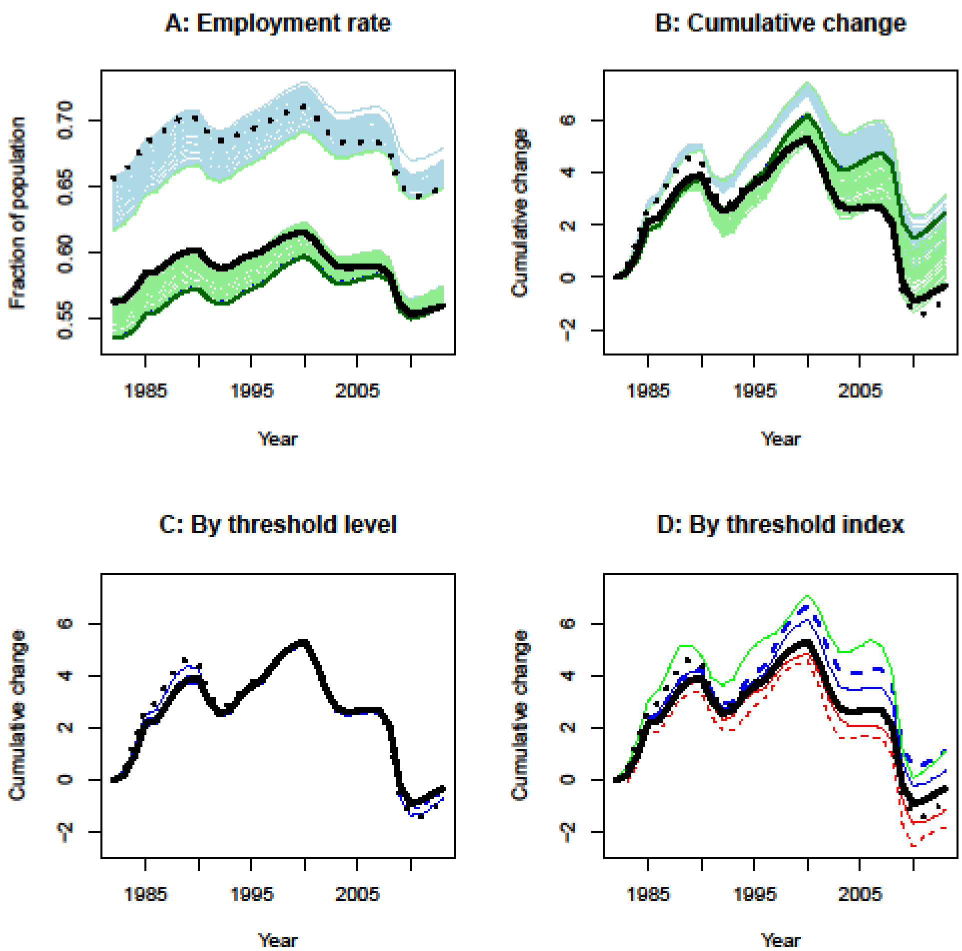

Note: Panel A gives the employment rate, and the remaining panels give the cumulative change from 1982 in percentage points. In all panels the black dotted line represents the fraction of the sample with non-zero earnings and the black solid line represents the fraction with earnings above 25 percent of median. The top two panels also show age-sex adjusted measures and the bottom two panels show effects of trimming at various levels other than 25 percent of median (Panel C) or with various indexes (Panel D). See the note to Figure 6 for the colors used.

Figure 8A gives the overall percentage of non-zero earnings (black) and the age-sex adjusted percentages (green). Trimming at 25 percent of median (black solid line), by reducing the number of workers counted as employed, reduces the employment rate compared with the untrimmed estimate (black dotted line).

The cumulative changes are given in Figure 8B. As with mean log earnings, the large differences in levels between the trimmed and on-trimmed unadjusted measures largely disappears when looking at the cumulative changes. As it happens, the employment rate during this period grew in the early part of the period and shrank back in the later point to about where it was in 1982, so that for the unadjusted measure, trimmed or not, the cumulative change over this particular period was close to zero.

As with mean log earnings, the age-adjusted measures of employment differ from the non-adjusted measures. The fixed-weight age-sex adjusted series (pale green band) ended up to 2 percentage points higher, with the chain-weighted index at the top end of that range. The age-sex adjusted untrimmed employment rates were slightly higher.

The effects of various trimming thresholds are given in 8C (for levels) and 8D (for indexes). As with mean log earnings, the effects of different threshold levels on the cumulative change are minimal, but the effects of the choice of index are on a scale that rival the effects from age-sex adjustment.

Although the effects of choice of index might arguably be considered as due to the choice of an inappropriate index for measuring a rising earnings distribution, the effects of age-sex adjustment reflect real changes that are worth noting even though the phase of the study reported in this paper won't delve into them. There were two notable trends that affected the age-sex adjusted rates. First, as the baby boomers aged, the age composition of the population shifted, with corresponding differences between the gross employment rate and the age-adjusted employment rate. Second, there was a marked increase in female employment, which took the form (not shown here) more of a rise in period employment than of a sharp rise in employment among some cohorts. There has also been a slight decrease in male employment at younger ages (not shown here).

These changes in employment are a reminder that not all changes in the earnings distribution can be read as upward or downward shifts in the earnings of workers who are in the workforce all along. Some changes are attributable to entrances or exits from the labor force that aren't evenly spread over the earnings distribution. If the decline in participation among young males came primarily from those who would have been low earners, average male earnings will have increased, even if the earnings among those who still remain have been unaffected. Similarly, if the rise in women's participation came from the increasing participation of college educated women of potentially higher than average earnings, the average female earnings, mean or median, will have increased, even if the earnings of women who would have been in the labor force already were unaffected. (And, if the new female earnings, although higher than the average already-present female earnings, are lower than the average male earnings, the increase in female participation can cause a decrease in average total earnings.)

Although the longitudinal earnings data allows the earnings distribution among those who have entered the labor force (or who will exit it) to be compared with those who remain in the labor force, these effects will not be studied in this paper, although such a study is needed for a complete understanding of measured changes in the distribution.

V.C: Variance of log earnings

As already discussed, an increase in the low earnings threshold, by eliminating part of the long lower tail, reduces the measured variance. In a particular year, the variance, when measured across all age and sex groups, can be decomposed into 1) the "within" variance that pools all the observations together after subtracting the mean log earnings for each observation's age-sex cell, and 2) the "between" variance that is the variance of those subtracted amounts. The total variance is the sum of these two components.

The components of the variance aren't as amenable to chain-weighting adjustments as measures of the mean are, because the between and total components are not unambiguously definable for the individual age-sex cells. Nevertheless, fixed-weight composition calculations can be carried out for all three components. These indicate that there are non-negligible composition effects. The "within" changes will be discussed first, followed by the "between" changes and the total variance.

Figure 9 gives the "within" variances. Figure 9A gives the levels for unadjusted within variances, both untrimmed (black solid) and trimmed at 25 percent of the median (black dotted). (The figure gives the age-sex adjusted levels as well). As expected, the trimmed variances, which cut off most of the long lower tail of log earnings, are considerably lower.

Unlike mean log earnings and the employment rate, the trimming affects the cumulative growth as well. The growth in the within variances (Figure 9B) for the trimmed earnings is both faster and smoother than that for the untrimmed earnings.

The age-sex adjustment has little effect on the within variance for untrimmed earnings (Figure 9B, blue band), but reduces the growth in the within variance for trimmed earnings (green band).

The choice of trimming threshold has potentially larger effects on the within variance than does the age-sex adjustment. Although the within variance for the 5 percent of median threshold differs little from that for untrimmed earnings, the successively higher thresholds above that yield successively higher growth in the within variance (Figure 9C). The choice of the indexing threshold has even more dramatic effects (Figure 9D), with the slower growing indexes like the price indexes or the FMW leading to still higher growth in the within variance.

Figure 9: "Within" component of variance of log earnings, level and cumulative growth

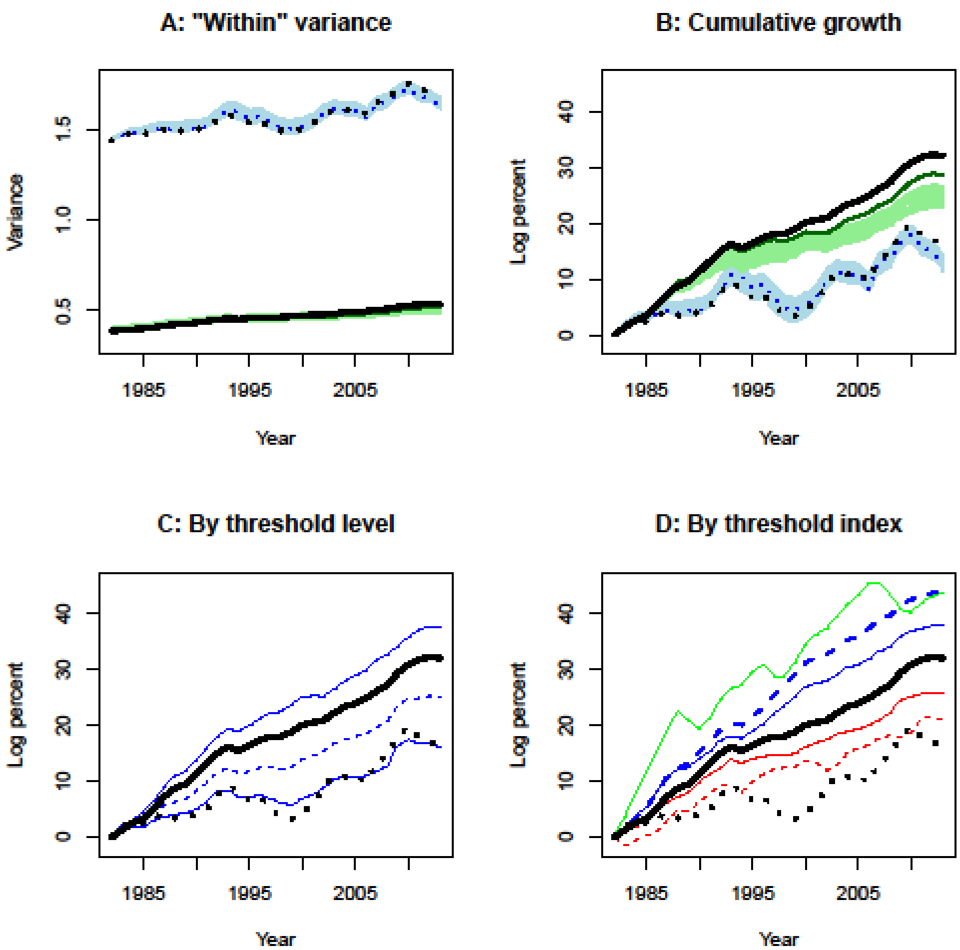

Note: Panel A gives the within variance, and the remaining panels give the cumulative change from 1982. In all panels the black dotted line represents the fraction of the sample with non-zero earnings and the black solid line represents the fraction with earnings above 25 percent of median. The top two panels also show age-sex adjusted measures and the bottom two panels show effects of trimming at various levels other than 25 percent of median (Panel C) or with various indexes (Panel D). See the note to Figure 6 for the colors used.

Figure 10 presents the corresponding effects for the "between" variances. In contrast to the within variances, the between variances have been declining over time, rather than rising (10A). The decline is more rapid for the trimmed earnings (10B, solid black line). The age-sex adjusted between variances, however, tend to show a rise or at least no decline, although the range of estimated effects is quite wide for the between variances of the trimmed earnings (10B, green band).

The choice of threshold has large effects. The higher the threshold, the faster the decline in the measured between variance (Figure 10C). The choice of index also has large effects, with the faster growing wage indexes leading to the largest declines (10D).

Figure 10: "Between" component of variance of log earnings, level and cumulative growth

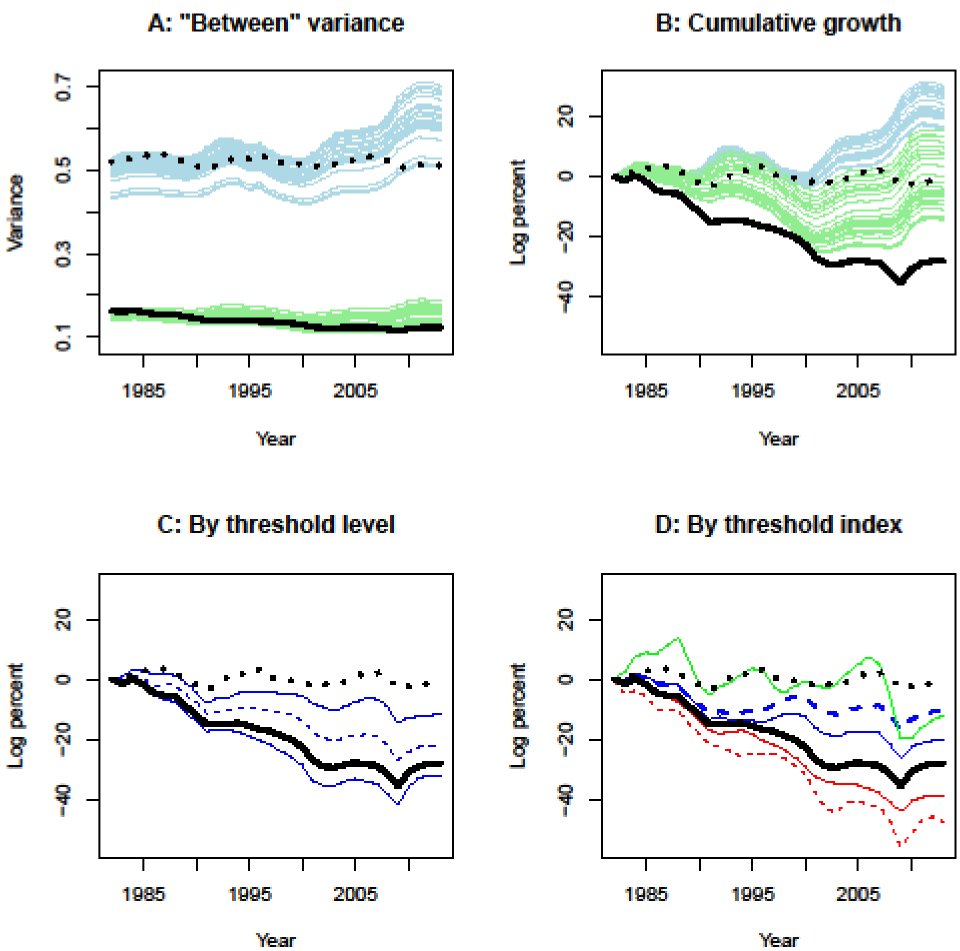

Note: Panel A gives the between variance, and the remaining panels give the cumulative change from 1982. In all panels the black dotted line represents the fraction of the sample with non-zero earnings and the black solid line represents the fraction with earnings above 25 percent of median. The top two panels also show age-sex adjusted fixed-weight measures and the bottom two panels show effects of trimming at various levels other than 25 percent of median (Panel C) or with various indexes (Panel D). See the note to Figure 6 for the colors used.

Figure 11 presents the effects for the total variance of log earnings. The decline in the between variance tends to offset the growth in the within variance, although the within variance, which provides a much larger share of the total, dominates, so that total variance grows over time. The differences between the trimmed and untrimmed earnings, and between the age-sex adjusted and the non-adjusted measures, tend to shrink when the within and between variance are combined. For the choice of threshold index, however (11D), the differences remain stark, with trimming thresholds based on a price index of the FMW rising by 20 to 30 log percent over this period while thresholds based on wage-indexed thresholds rise very little.

Figure 11: Total variance of log earnings, level and cumulative growth

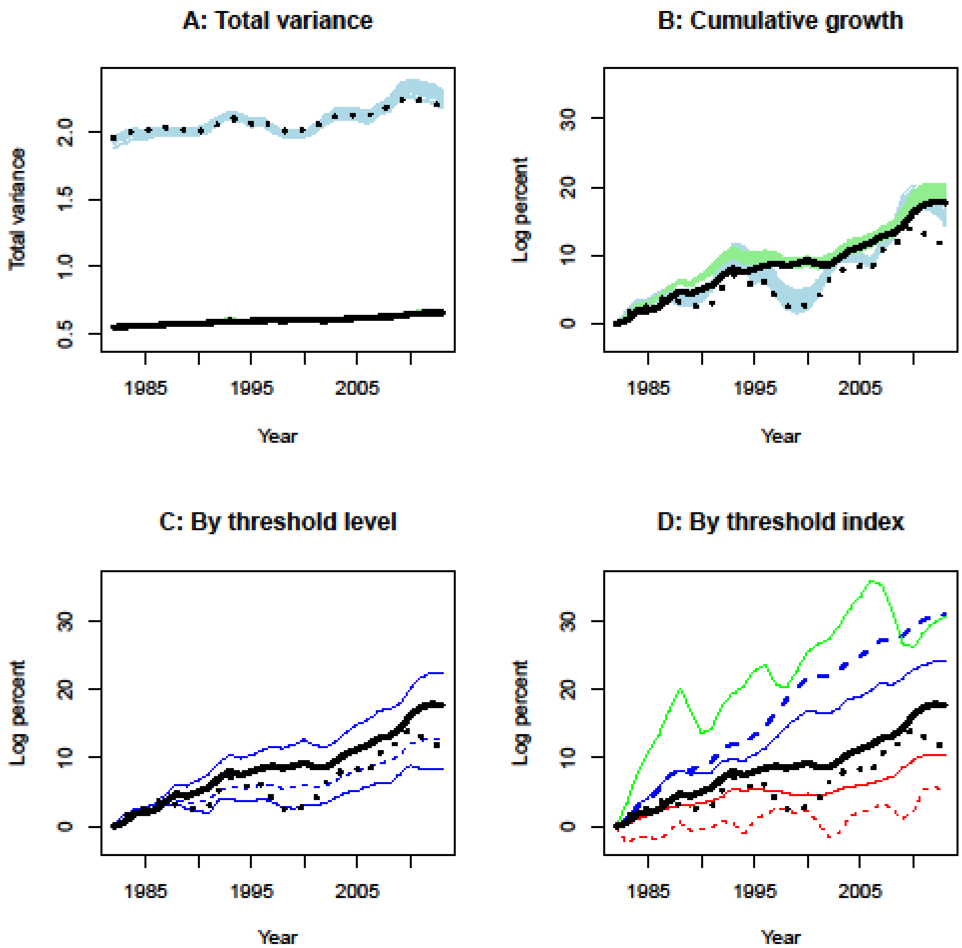

Note: Panel A gives the total variance, and the remaining panels give the cumulative change from 1982. In all panels the black dotted line represents the fraction of the sample with non-zero earnings and the black solid line represents the fraction with earnings above 25 percent of median. The top two panels also show age-sex adjusted fixed-weight measures and the bottom two panels show effects of trimming at various levels other than 25 percent of median (Panel C) or with various indexes (Panel D). See the note to Figure 6 for the colors used.

Measures of the within variance as a share of the total show perhaps the widest variation of all the measures considered here (Figure 12). Although the within-share for untrimmed earnings rose only about two percentage points (Figure 12N, black dotted line), from just under 75 percent to just over 75 percent, successively higher thresholds yield successively high increase in the within share (12C). At the 25 percent of median trimming baseline the share rises from 70 percent to over 80 percent (12A, black solid line). Age-sex adjustment reduces this latter rise considerably, although there is wide variation from the choice of the year for the fixed weights (12B, green band). Age-sex adjustment of the share for untrimmed earnings coverts the slight rise into a slightly larger decline (12B, blue band). Perhaps most surprising, given the patterns for the other measures examined here, the choice of index for the trimming thresholds has a considerably smaller effect on the within share (12D) because the percent effects on the total variance (11D) are similar to the percent effects on the within variance (9D).

Figure 12: "Within" variance as share of total, level and cumulative change

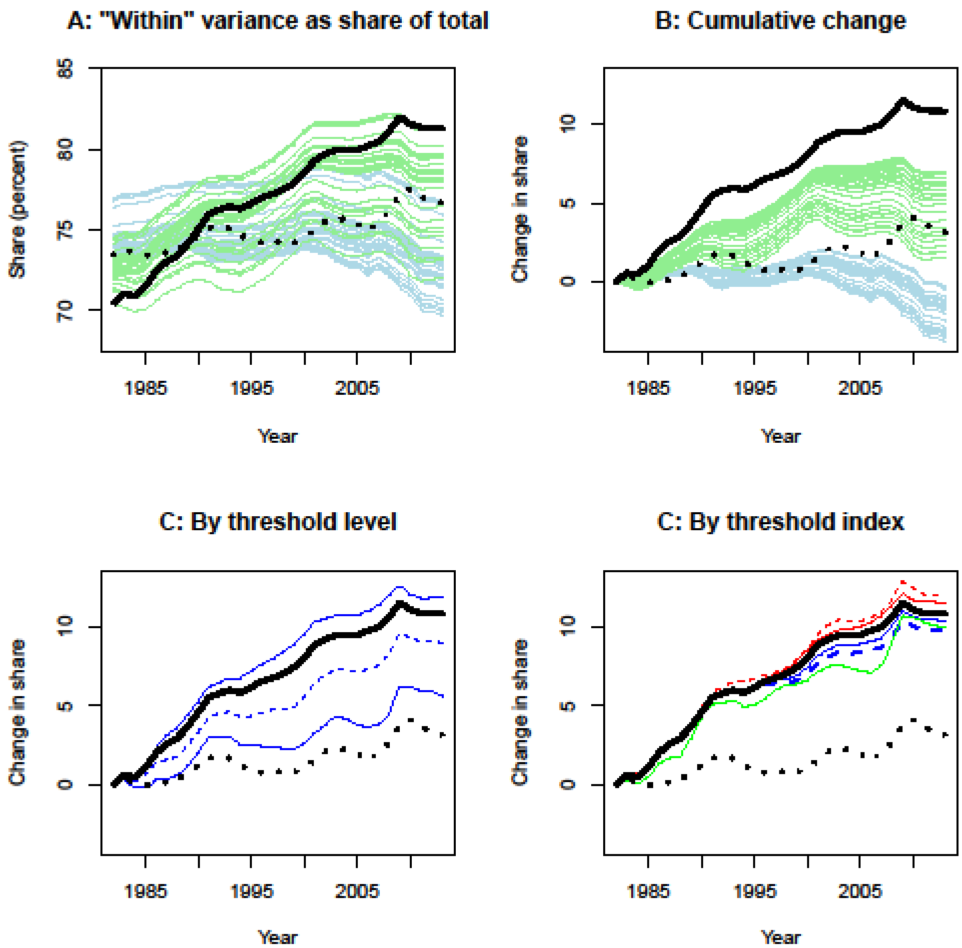

Note: Panel A gives the within variance as a share of total variance, and the remaining panels give the cumulative change from 1982. In all panels the black dotted line represents the fraction of the sample with non-zero earnings and the black solid line represents the fraction with earnings above 25 percent of median. The top two panels also show age-sex adjusted fixed-weight measures and the bottom two panels show effects of trimming at various levels other than 25 percent of median (Panel C) or with various indexes (Panel D). See the note to Figure 6 for the colors used.

VI. Discussion

Summarizing some of the results:

- Low earnings thresholds have expected effects on the mean earnings above the threshold: the higher the threshold, the higher the mean of the earnings above the threshold. (Not shown here: increasing the threshold also increases the median and other percentiles.) But although the thresholds shift the time series of mean earnings up, the growth rate is not much affected, even for the large shifts explored here (from no threshold to a threshold at 35 percent of median earnings). For mean earnings it is more important to pay attention to the effects of the shifting age-sex composition of the work force.

- The choice of the growth index for the threshold (or for the calculation of "real" wages to be compared to a fixed threshold) is for many measures more critical than the level at which the threshold is set. For measuring the growth in earnings the choice between the national average wage index and a price index like the CPI or PCE is more important than the choice of the threshold level. But the age-sex composition effects are still more important.

- For measuring the growth in the variance of log earnings, the choice of threshold and in particular the growth index becomes more important. High thresholds tend, as expected, to reduce the variance of log earnings above the threshold, but the higher thresholds, in the range of thresholds studied here, tend to *increase* the growth rate of the variance. It remains the case, however, that the choice of the threshold growth index, in the range of indexes studied here, has more effect on the growth rate of the variance than the choice of threshold level over the wide range of thresholds studied here. The fastest growing rate of those studied (the national average wage, especially in its lagged version) takes away much of the observed growth in variance.

- For the decomposition of the variance of log wages into the variance *between* mean earnings in age-sex groups and the variance *within* those groups, the picture is much the same for the *within* variance as it is for the total variance, particularly the direction of the effects, although there are differences in magnitude. For the *between* component, which has been flat or declining over this period, higher thresholds generate a faster decline, as does faster index growth, but age-sex adjustment slows or reverses the decline.

A more general conclusion is that it is difficult to formulate general rules that will hold true for all measures. Although some measures might turn out to be robust to age-sex adjustment or to the choice of a trimming threshold, many don't, and it is simpler to test each particular measure than it is to come up with a general rule for when the adjustments are necessary. It was only the logistical burden, rather than the lack of a technique, that prevented this phase of the study from extending its scope to medians and ratios of percentages. (Initial testing indicated that the median and percentile-based measures of the variance like the ratio of the 90th to the 10th percentile are also sensitive to age-sex adjustment and to the choice of trimming thresholds and indexing.)

It is worth repeating that the choice of the growth index for the threshold is a neglected aspect of choosing a trimming threshold. The choice of the threshold index is in a sense a concealed choice in many studies, concealed in the use of real earnings compared to a single year's threshold. This procedure effectively indexes the threshold to whatever price index was selected. For a study spanning several decades, the choice of the index, even the choice between a CPI and a PCE price index, can have a significant effect on the fraction of the distribution that is trimmed in years far from the indexing base year, and hence a sizeable effect on the measured mean earnings or variance.

The indications from the figures presented here are that thresholds based on the lagged NAW (such as each year's 4QC amount) yield considerably more variability in the measured variance than would an index based on an unlagged national average wage. The OASDI program is constrained to use a two year lag for the indexed provisions because the contemporaneous index is not yet available in the year the provisions apply. Researchers, however, do not face that constraint and should use the unlagged NAW. For similar reasons, thresholds based on the year to year movement in the federal minimum wage are a particularly bad choice. A threshold can be based on the level of the federal minimum wage in a particular year, but it should be indexed to other years using an index that varies more smoothly relatively to the distribution of earnings.

For studies over long time spans of changes in the broad middle of the distribution the best index is arguable one that grows along with the middle of the distribution. There has been an unspoken presumption in this study that the growth rate in median earnings is the best rate to index the threshold with (holding the threshold constant relative to mean earnings). This can be debated, but it must certainly be true that the choice of index should be given more thought than it usually receives.

There are several possible extensions of the work reported here.

Some work was done restricting the worker-year observations to those with earnings by the same worker in the preceding and following year. This restriction causes, as expected, shifts toward lower employment rates, higher average earnings, and lower variance, with particularly large shifts after age 60, but the effects of choice of threshold have yet to be explored.

The effect of thresholds on measures of the distribution of *lifetime* earnings has not yet been explored. Lifetime earnings bring in yet another dimension to the filtering: in addition to the choice of the level of threshold and the indexing of the earnings or the threshold, there is the number of years a worker must have earnings above the threshold in order to be included in the measured sample.

The more elaborate approach that has been used in some studies of earnings dynamics - trimming the earnings distributions within each age-sex cell each year at a specified percentile of that cell's distribution – has not yet been explored. This approach avoids the problem of the interaction between trimming thresholds and changing age-sex composition by defining the trimming within each age-sex cell and avoids the problem of indexing by trimming at specified percentiles rather than indexed dollar values.

A final challenge is to determine the effect that changing participation patterns have had on the distribution of earnings. If more women have come into the labor force, or if fewer younger men have participated, then unless the distribution of earnings among the marginal additions or subtractions from the labor force matches that of the remaining participants, the distribution of workers earnings will have changed in directions that are important to distinguish from changes in earnings among the continuing participants.

Appendix: Sampled Literature on Thresholds

This appendix summarizes a non-systematic sampling of the thresholds that have been applied in the literature. The purpose was to get a range of thresholds to be tested. The thresholds that will be used in this paper, ranging from 0 percent to 35 percent of the median wage and indexed by the national average wage index, the PCE deflator, the consumer price index, and the federal minimum wage, were derived from this survey.[11]

I catalog the thresholds here by the measure on which the threshold is based (most commonly the earnings from half-time federal minimum wage employment); the year chosen for that measure, such as the minimum wage in 2005; and the index used to compare earnings outside that base year, most commonly a price index like the CPI or PCE.

A few researchers used no threshold, including Utendorf (1998), and Lee (1999).

Very few have used an absolute value cutoff, such as $1 or $100, which would have the problem of shrinking relative to the wage distribution as wages grew. Some researchers have used a dollar cut off, but index all earnings to the year for which the cutoff is defined. Lemieux (2004) uses hourly wages in 1979 dollars (indexing not stated?), trimming any less than $1 or greater than $100. The $1 threshold is not as low as it might seem. Working for 40 weeks at 35 hours would give annual earnings of $1,700, which was 14 percent of the tabulated median in that year.

Several Social Security program parameters are indexed to the national average wage index (NAW), and some researchers have used these parameters as thresholds. Pattison (2014) used the one-quarter-coverage level (1 QC) for some tabulations, as well as sensitivity to higher levels (2, 3, and 4 QCs). The 1 QC level[12] works out to 3 to 4 percent of the tabulated median wage; 4 QCs accordingly would be in the range of 12 to 16 percent of the tabulated median. Sabelhaus and Song (2009) used the nominal 4QC level, noting that in 2005 this was equivalent to 715 hours at the FMW. Relative to the tabulated median, this threshold rose from about 14 percent to about 16 percent of the tabulated median.

At a higher level, higher than most thresholds that have been used in earnings analysis, is the retirement earnings test (RET62) level for Social Security beneficiaries between ages 62 and the normal retirement age. (This was the main threshold for Pattison (2014), where the interest was in how many workers remained above the RET at those older ages.) This threshold has been at about 33 percent of the wage indexing series. As a percent of the tabulated median it has ranged from 40 to 50 percent. The substantial gainful activity (SGA) level used in disability determination has been wage indexed in recent years; it lies between the half-time FMW level and the RET62 level.

My casual survey indicates that the federal minimum wage (FMW) is the most commonly used basis for a threshold. Many of the CPS studies limited their sample to full-time full-year workers but added an additional FMW-based threshold to filter out workers with reported amounts believed to be plausibly low for a full-year full-time worker. The most common choice is an estimate of earnings for a half-time (typically 1,000 hours) FMW worker.

- Autor, Katz and Kearney (2009), working with estimated weekly wages rather than annual earnings, used the 1982 FMW at half-time, equivalent to annual hours of 800 hours, if a 40-week year, or 1040 hours, if a 52 week year. Their indexing series wasn't clear.
- Bosworth, Burtless, and Zhang (2015) use 1,000 hours of the 2005 FMW, indexed to other years using the CPI.
- Leonesio and Del Bene (2011) used $5,000 in 2000, with PCE indexing. $5,000 was chosen as approximating 1000 hours at the FMW.

These thresholds, even when using the same price index, are at different levels because of the different base years, but the resulting thresholds have tended to fall between 15 percent and 25 percent of median wages.

---

[11] Many of the surveyed studies had, in addition to annual earnings, information on weeks or hours worked, and used this data to further restrict the sample to full-time full-year workers in addition to imposing a trimming threshold on annual earnings. This additional hours and earnings filter is not available for the SSA earnings data.

[12] $250 in 1978, rising to $1,150 in 2013.

Some researchers have based their threshold on the FMW, but use a lower value than half-time.

- Kopczuk, Saez, and Song used 500 hours of the 2004 FMW, or quarter-time FMW. They differ from others in using the national average wage to index the cutoff, rather than the PCE or CPI. (They also append as a sensitivity test a calculation at full-time minimum wage, 2,000 hours times the FMW, the highest series seen in this survey.)
- Eckstein and Nagycal (2004) used 1/6 of the 2002 minimum wage, PCE indexed to other years. They argued that the ½ FMW threshold was too high for many women in the 1960s who were high school dropouts worker as household or farm workers and who were plausibly being paid very low wages. They noted that "the conclusions regarding the evolution of female inequality are very sensitive to how one treats these very low observations of weekly wages. If these low observations were dropped from the sample, then there would be no fall in female inequality in the 60's."

In a series of articles Guvenen co-authors (latest: 2017) use one eighth of the FMW (expressed as one-quarter of full-time work at half the minimum wage), or 260 hours at the FMW. Their threshold is unusual in that the threshold each year is that year's nominal FMW, which has not been smoothly indexed to growth in median wages. Overall, the level is at about that of Eckstein and Nagycal, fluctuating between 5 and 10 percent of the tabulated median earnings.

An approach which won't be examined here is to trim the earnings at specified quantiles. Lochner and Shin (2014) trimmed the observations at the 1st and 99th percentiles within 10-year age cells before regressions on the log earnings. Moffit and Gottschalk (2012) also trimmed within age group cells, but only after the regression, trimming the regression residuals at the 1st and 99th percentiles. They refer to an earlier 2002 paper for sensitivity tests indicating that the trimming doesn't make any qualitative difference (although that earlier paper noted that restricting the sample to full year workers had effects that merit further investigation, and that larger or smaller amounts of trimming had effects, but no qualitative difference).

References

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*, May 2009, (90(2): 300-323.

Bosworth Barry, Gary Burtless, and Kan Zhang, "Later Retirement, Inequality in Old Age, and the Growing Gap in Longevity Between Rich and Poor," *Economic Studies at Brookings*, 2015.

Eckstein, Zvi and Eva Nagycal, "The Evolution of U.S. Earnings Inequality: 1961-2002," *Federal Reserve Bank of Minneapolis Quarterly Review*, December 2004.

Guvenen, Fatih, "Firming up Inequality," *Heller-Hurwicz Institute Policy Brief*, August 2017.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?," August 3, 2016.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song, "Earnings Inequality and Mobility in the United Stats: Evidence from Social Security Data since 1937," *The Quarterly Journal of Economics*, February 2010.

Lee, David S., "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?," *The Quarterly Journal of Economics*, Vol. 114, No. 3 (Aug. 1999), pp. 977-1023.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" May 2004.

Leonesio, Michael V. and Linda Del Bene, "The Distribution of Annual and Long-Run US Earnings, 1981-2004." *Social Security Bulletin*, Vol. 71, No. 1, 2011.

Lochner, Lance J. and Youngki Shin, "Understanding Earnings Dynamics: Identifying and Estimating the Changing Roles of Unobserved Ability, Permanent and Transitory Shocks," CIBC Centre for Human Capital and Productivity, 2014.

Moffitt, Robert A. and Peter Gottschalk, "Trends in the Transitory Variance of Earnings in the United States," *The Economic Journal*, 2002, 112(478): C68-C73.

Moffit, Robert A. and Peter Gottschalk, "Trends in the Transitory Variance if Earnings: Methods and Evidence," *Journal of Human Resources*, 47(1) (December 2012):204-226.

Pattison, David, "Trends in Retirement and Claiming in the SSA Data," December 2014.

Sabelhaus, John and Jae Song, "Earnings Volatility across Groups and Time," *National Tax Journal*, Vol. LXII, No. 2, June 2009.

Sabelhaus, John and Jae Song, "The Great Moderation in Micro Labor Earnings." Draft, Sep. 2009.
Utendorf, Kelvin U., "Recent Changes in Earnings Distributions in the United States," *Social Security Bulletin*, Vol. 61, No. 2, 1998.